Generation of realistic white matter numerical phantoms with controllable morphology for diffusion MRI simulations

Ross Callaghan

Main Supervisor: Prof. Gary Zhang

Co-Supervisor: Dr. Marco Palombo

Co-Supervisor: Prof. Danny Alexander

A thesis submitted for the degree of

Doctor of Philosophy

in

Medical Imaging

as part of

University College London

Centre for Doctoral Training in Medical Imaging

March 22, 2022

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

This work is supported by the EPSRC-funded UCL Centre for Doctoral Training in Medical Imaging (EP/L016478/1) and the Department of Health's NIHR-funded Biomedical Research Centre at UCLH

Abstract

Numerical phantoms have played a key role in the development of diffusion MRI (dMRI) techniques seeking to estimate features of the microscopic structure of tissue by providing a ground truth for simulation experiments against which we can validate and compare techniques. One common limitation of numerical phantoms which represent white matter (WM) is that they oversimplify the true complex morphology of the tissue which has been revealed through *ex vivo* studies. It is important to try to generate WM numerical phantoms that capture this realistic complexity in order to understand how it impacts the dMRI signal.

This thesis presents work towards improving the realism of WM numerical phantoms by generating fibres mimicking natural fibre genesis. A novel phantom generator is presented which was developed over two works, resulting in **Con**textual **Fi**bre **G**rowth (ConFiG). ConFiG grows fibres one-by-one, following simple rules motivated by real axonal guidance mechanisms. These simple rules enable ConFiG to generate phantoms with tuneable microstructural features by growing fibres while attempting to meet morphological targets such as user-specified density and orientation distribution. We compare ConFiG to the state-of-the-art approach based on packing fibres together by generating phantoms in a range of fibre configurations including crossing fibre bundles and orientation dispersion. Results demonstrate that ConFiG produces phantoms with up to 20% higher densities than the state-of-the-art, particularly in complex configurations with crossing fibres. We additionally show that the microstructural morphology of ConFiG phantoms is comparable to real tissue, producing diameter and orientation distributions close to electron microscopy estimates from real tissue as well as capturing complex fibre cross sections. ConFiG is applied to investigate the intra-axonal diffusivity and probe assumptions in a family of dMRI modelling techniques based on spherical deconvolution (SD), demonstrating that the microscopic variations in fibres' shapes affects the diffusion within axons. This leads to variations in the per-fibre signal contrary to the assumptions inherent in SD which may have a knock-on effect in popular techniques such as tractography.

Impact Statement

The focus of this thesis is the development of a new tool to generate realistic synthetic white matter (WM) phantoms which we can use in simulation experiments to test and develop existing and new diffusion MRI (dMRI) models. This tool, which we call ConFiG, works by mimicking natural fibre growth and enables us to generate synthetic models of white matter that are more realistic than previously achievable, opening a door into investigation of the dMRI process in more detail than ever before. This tool has the potential to impact not just our immediate academic community, but the wider academic community and society as a whole as outlined in the remainder of this statement.

Immediate academic community ConFiG produces highly realistic WM phantoms which outperform the previous state-of-the-art, producing microscopically realistic structure in the generated axons. An immediate of impact of this the recently published NeuroImage paper outlining the method and demonstrating its effectiveness. Further, in this thesis we demonstrate that ConFiG can be used to probe assumptions in dMRI models in a way that was previously infeasible, showing that some of the assumptions may not hold true which may have a downstream effect on popular dMRI techniques. This is merely the tip of the iceberg of what ConFiG could potentially bring to the dMRI community, offering the potential to study the diffusion process in greater detail than ever and generate new dMRI models capable of more accurately quantifying microscopic features of tissue non-invasively.

Wider academic community The idea behind ConFiG, to grow cells mimicking nature, could be extended to produce phantoms for other tissues including branching neuronal cells in grey matter and non-brain tissues (for instance the complex micro-environment of a tumour), potentially creating a whole suite of realistic tissues to build new models for all sorts of applications. On top of this, the phantoms ConFiG produces are stored as 3D meshes, which could be applied to many modalities outside of dMRI in which it is possible to simulate the signal using a mesh. For instance the realistic microstructure could be used with an electron microscopy (EM) simulator to generate realistic EM images to help train and test models designed to segment WM axons.

Beyond academia The work presented in this thesis is the first step on a journey towards potentially exciting clinical applications including more accurate imaging of diseases affecting the WM such as multiple sclerosis (MS). Here we demonstrate that ConFiG can generate realistic WM microstructure, and from that realistic dMRI signal which lays the groundwork for this to be used to develop new dMRI models to estimate microstructural information from WM. As an example of this, ConFiG could be used to help us more effectively differentiate axonal loss and demyelination in MS, something which is valuable information for differential diagnosis but is currently very difficult to estimate from dMRI.

Contents

Contents 7

List of Figures 11

List of Acronyms 13

I Introduction and background 17

1 Introduction 19
1.1 Motivation . . . 19
1.2 Problem Statement . . . 22
1.3 Project Aims and Scope . . . 22
1.4 Report Overview . . . 23

2 Background 25
2.1 MR Physics . . . 25
2.1.1 Nuclear Magnetism . . . 26
2.1.2 Magnetic Resonance . . . 26
2.1.3 The Bloch Equations . . . 28
2.1.3.1 Relaxation . . . 29
2.1.3.2 The rotating frame . . . 30
2.1.4 Detecting the MR signal . . . 32
2.1.5 Spin echoes . . . 33
2.2 Diffusion MRI . . . 35
2.2.1 The Bloch-Torrey equations . . . 38
2.3 Diffusion Simulation . . . 38
2.3.1 Numerical solutions . . . 39
2.3.2 Monte-Carlo simulations . . . 39

2.4 Diffusion Modelling . . . 41
2.4.1 Spherical deconvolution techniques . . . 41
2.4.1.1 Spherical harmonics . . . 42

3 Literature Review 43
3.1 Numerical Solutions . . . 43
3.2 Monte-Carlo - Packages . . . 45
3.3 Monte-Carlo - Numerical Phantoms . . . 45
3.3.1 Towards realistic WM numerical phantoms . . . 48
3.4 Conclusions . . . 49

II Methods Development 51

4 Fibre Growth for WM numerical phantom generation 53
4.1 Introduction . . . 53
4.2 A preliminary fibre growth algorithm . . . 55
4.2.1 STEP 1: Initial growth configuration . . . 55
4.2.2 STEP 2: Fibre growth . . . 56
4.2.3 STEP 3: Meshing . . . 58
4.3 Experiments and Results . . . 60
4.3.1 Effect of choice of growth network . . . 60
4.3.2 Demonstration of preFiG . . . 61
4.3.3 Comparison with brute-force approach . . . 62
4.4 Discussion and Conclusion . . . 64

5 ConFiG: Contextual Fibre Growth 65
5.1 Introduction . . . 66
5.2 Methods . . . 67
5.2.1 Overview of the ConFiG algorithm . . . 67
5.2.2 Biological motivation for ConFiG . . . 68
5.2.2.1 Chemoattraction . . . 68
5.2.2.2 Fibre collapse . . . 70
5.2.2.3 Dynamic growth network . . . 70
5.2.2.4 Axon fasciculation . . . 71
5.2.2.5 Global optimisation . . . 72
5.2.3 Summary of ConFiG input parameters . . . 73
5.3 Experiments . . . 74
5.3.1 Testing the performance of ConFiG . . . 74
5.3.1.1 Interaction of ConFiG parameters . . . 75
5.3.2 Diffusion MRI simulation . . . 75
5.3.3 3D signal visualisation . . . 76
5.4 Results . . . 77
5.4.1 Impact of biological mechanisms . . . 77
5.4.2 Interaction of ConFiG parameters . . . 78
5.4.3 Diffusion magnetic resonance imaging (MRI) simulation . . . 79
5.5 Discussion . . . 83

5.6 Conclusion . . . 86

6 Microstructural Evaluation 87
6.1 Introduction . . . 88
6.2 Microstructural measurements . . . 88
6.2.1 Virtual histology and 2D morphological measures . . . 89
6.3 Experiments . . . 91
6.3.1 Relationship between input and output morphology . . . 91
6.3.2 Packing induced microstructural complexity . . . 92
6.3.3 Variability of repeated instances of contextual fibre growth (ConFiG) . . . 92
6.4 Results . . . 93
6.4.1 Microstructural measures and virtual histology . . . 93
6.4.2 Relationship between input and output morphology . . . 95
6.4.3 Packing induced microstructural complexity . . . 96
6.4.4 Variability of repeated instances of ConFiG . . . 96
6.5 Discussion . . . 99
6.5.1 Limitations and future Work . . . 99
6.6 Conclusion . . . 100

III Applications 101

7 Impact of realistic axonal geometry on intracellular diffusion 103
7.1 Introduction . . . 103
7.2 Method . . . 105
7.2.1 Impact of beading and undulation . . . 105
7.2.2 A phenomenological model of undulation and beading . . . 106
7.2.3 Realistic axonal morphology . . . 106
7.2.4 Impact on estimated $d_{\|}$ . . . 108
7.3 Results . . . 109
7.3.1 Impact of beading and undulation . . . 109
7.3.2 Phenomenological model of undulation and beading . . . 109
7.3.3 Realistic axonal morphology . . . 109
7.3.4 Impact on estimated $d_{\|}$ . . . 109
7.4 Discussion . . . 114
7.4.1 Limitations and future work . . . 116
7.5 Conclusion . . . 117

8 Impact of within-voxel heterogeneity in fibre geometry on spherical deconvolution 119
8.1 Introduction . . . 120
8.2 Method . . . 121
8.2.1 Phantom Generation . . . 121
8.2.1.1 Real white matter (WM) fibres from electron microscopy (EM) . . . 122

8.2.1.2 Gold standard fibre orientation distribution function (fODF) extraction from microstructure 122
8.2.1.3 fODF metrics 124
8.2.2 Experiments 125
8.2.2.1 Per-fibre response heterogeneity (Experiment 1) 125
8.2.2.2 Impact on fODF estimation (Experiment 2) 127
8.3 Results 131
8.3.1 Per-fibre response heterogeneity 131
8.3.2 Impact on fODF estimation 131
8.4 Discussion 136
8.4.1 Limitations and future work 137
8.5 Conclusion 139

IV Conclusions **141**

9 Conclusions **143**
9.1 Contributions 143
9.2 Future directions 144

Publications **147**

References **151**

List of Figures

1.1 Overview of the main goal of the thesis 22

2.1 Illustration of the precessional motion of a single spin and production of magnetisation by many spins 27
2.2 Nutation motion of an on-resonance spin the presence of an RF field. . . . 28
2.3 Motion of a spin in the presence of a B_1 RF field in the rotating frame . . 31
2.4 The effective field, B_e, produced due to an off-resonance frequency Ω. . . 32
2.5 The free induction decay described by Equation (2.32) 33
2.6 The spin echo sequence and the evolution of spin under a spin echo sequence. 34
2.7 The standard pulsed gradient spin echo sequence used in dMRI. 36

4.1 Inputs to the preliminary fibre growth (preFiG) algorithm 55
4.2 Overview of the preFiG growth algorithm 57
4.3 preFiG meshing procedure 59
4.4 Axial diffusivity for different metaball meshing thresholds 60
4.5 Fibres generated using uniform grid (orange) and pseudo-random (blue) network nodes for increasing numbers of nodes. 61
4.6 Examples of generate WM numerical phantoms and simulated dMRI signals 62
4.7 Comparison of brute-force growth and proposed fibre growth algorithm . 63

5.1 Illustration of the chemoattraction process 68
5.2 Illustration of the contact guidance mechanism 70
5.3 Illustration of the fasciculation process 71
5.4 Impact of biological mechanisms in ConFiG 78
5.5 Impact of biological mechanisms on virtual histology 78
5.6 Impact of biological mechanisms on 3D phantoms 79
5.7 Impact of improvements on simulated diffusion MRI (dMRI) signal . . . 80
5.8 Interaction between target density and OD. 80
5.9 dMRI signals from HCP subject and ConFiG phantoms 81
5.10 3D visualisation of simulated and real dMRI signals 82
5.11 Proof-of-concept ConFiG phantoms with real cells 85

6.1 Mesh centre line extraction 89
6.2 Orientation distribution extraction from microstructure 90
6.3 Diameter distribution extraction from microstructure 91
6.4 Comparison of real and virtual histology 93
6.5 Slice-wise microstructural measurements 94
6.6 Slice-wise microstructural measurements with colourmap 95
6.7 Diameter distributions of real and ConFiG axons 96
6.8 Orientation distributions from real WM and ConFiG phantoms 97
6.9 Packing induced microstructural complexity 98
6.10 Variability of ConFiG phantoms under repetition 98

7.1 Example fibres used to investigate impact of undulation (a) and beading (b). 106
7.2 The method used to generate simulated intracellular signals 107
7.3 Intracellular axial diffusivity in beading and undulation 111
7.4 Intracellular axial diffusivity in realistic phantoms 112
7.5 Impact of microstructure on $d_{\parallel}$ 113
7.6 Estimated $d_{\parallel}$ in cylinders with varying OD 113
7.7 Estimated $d_{\parallel}$ in cylinders at SNR=30 116

8.1 Gold standard fODF estimation from microstructure of WM numerical phantoms 124
8.2 Steps 1 and 2 in the experiment pipeline 128
8.3 The third step in the experiment is to compute voxel signals and perform constrained spherical deconvolution (CSD). 129
8.4 Illustration of the method to produce voxel signals 130
8.5 Variability in fibre responses within a voxel at $b = 3000\,\mathrm{s\,mm^{-2}}$ at SNR=30 along with geometrical variation in fibres responsible for median, 10th and 90th percentile response. 132
8.6 Per-fibre response function at $b = 1000,\ 2000,\ 3000\ s/mm^2$ (left-to-right) for (a) the EM fibres, (b) the $\kappa = 2$ phantom and (c) the three-crossing phantom. 133
8.7 Variation in fODF estimated using the range of fibre response functions (FRFs) in Experiment 2 at b = $3000\,\mathrm{s\,mm^{-2}}$ for the single fibre bundle voxels 134
8.8 Variation in fODF estimated using the range of FRFs in Experiment 2 at b = $3000\,\mathrm{s\,mm^{-2}}$ for the crossing fibre bundle voxels. 135
8.9 Variation in fODF properties when estimated using different FRFs. 135
8.10 An example of spurious peaks in the fODF of an Human Connectome Project (HCP) subject 138

List of Acronyms

ADC Apparent diffusion coefficient

BT Bloch-Torrey

BVH Bounding volume hierarchy

CAM Cell adhesion molecule

CC Corpus callosum

CHARMED Composited hindered and restricted models of diffusion

ConFiG Contextual fibre growth

CPU Central processing unit

CSD Constrained spherical deconvolution

CSF Cerebrospinal fluid

dMRI Diffusion MRI

DKI Diffusion kurtosis imaging

DMS Diffusion Microscopist Simulator

DSI Diffusion spectrum imaging

DTI Diffusion tensor imaging

EM Electron microscopy

ESAG Elliptically symmetric angular Gaussian

FA Fractional anisotropy

FDM Finite difference method

FEM Finite element method

FID Free induction decay

fODF Fibre orientation distribution function

FRF Fibre response function

HCP Human Connectome Project

IC Internal capsule

MC Monte-Carlo

MEDUSA Microstructure environment designer with unified sphere atoms

MRI Magnetic resonance imaging

MRS Magnetic resonance spectroscopy

MS Multiple sclerosis

MSD Mean squared displacement

NMR Nuclear magnetic resonance

NN Neural network

NODDI Neurite orientation dispersion and density imaging

OD Orientation dispersion

OGSE Oscillating gradient spin echo

PDE Partial differential equation

PDF Probability density function

PGSE Pulsed gradient spin echo

preFiG Preliminary fibre growth

RF Radiofrequency

RMSE Root mean square error

RWS Random Walk Simulator

SBEM Serial block-face scanning electron microscopy

SD Spherical deconvolution

SEM Scanning electron microscopy

SGP Short gradient pulse

SH Spherical harmonic

SNR Signal-to-noise ratio

TBI Traumatic brain injury

TC Three crossing

TE Echo time

TEM Transmission electron microscopy

WM White matter

Part I

Introduction and background

1 Introduction

Chapter Contents

1.1 Motivation 19

1.2 Problem Statement 22

1.3 Project Aims and Scope 22

1.4 Report Overview 23

Chapter Summary

The purpose of this chapter is to introduce the main motivations behind this project, from the historical relevance of diffusion MRI to the use of modern simulations for validating and developing new diffusion MRI models. This chapter also sets out the problem that this project aims to address and the specific aims for addressing various aspects of it.

1.1 Motivation

In 1992, James Watson, co-discoverer of DNA, said "The brain is the last and grandest biological frontier, the most complex thing we have yet discovered in our universe." [1]. One year later, Francis Crick, fellow discoverer of DNA, and Edward Jones published a commentary in Nature lamenting how little was understood about human neuroanatomy, saying that new techniques were needed beyond the contemporary tracer studies in non-human primates [2].[1]

[1] This introduction was inspired by Richard Passingham's very nice foreword to Diffusion MRI by Berg and Behrens [3]

Just one year after that, in 1994, Basser et al. [4] showed that it was possible to use magnetic resonance imaging (MRI) to measure the movement of water along axons, providing the basis for exactly the kind of new technique Crick and Jones had felt was needed. This technique of using MRI for measuring the movement of water molecules is known as diffusion MRI (dMRI).

In the 25 years since the work of Basser et al., diffusion MRI has grown into a major MRI research field, generating thousands of publications per year. Diffusion MRI has found extensive use for imaging the brain, generating new techniques such as tractography, which attempts to map out the connections in the brain *in vivo*, and microstructure imaging, in which measurements of the diffusion of water in tissue are used to infer information about the structure of the tissue on a scale much lower than the size of an MRI voxel.

These techniques work because dMRI sensitises the MRI signal to the diffusive motion of water molecules on the micrometer scale - the further the water diffuses, the more strongly the MR signal is attenuated. The environment in which the water molecules move restricts the motion of the molecules and so will affect how strongly the dMRI signal is attenuated. This dependency of the dMRI signal on the environment in which water molecules diffuse can be exploited to infer information about the environment solely from the dMRI signal.

Microstructure imaging attempts to do exactly this, infer information about the microstructural environment such as cell size and density from the dMRI signal. Since the microstructure is on a much smaller scale than the dMRI voxels, we cannot directly image this structure, so mathematical models are used to relate the dMRI signal from a voxel to microstructural features.

The validation of these microstructural models can be difficult since ground truth microstructural features are typically inaccessible *in vivo* and classical histological validation techniques have limitations such as disruption due to tissue extraction and preparation.

One approach commonly taken for the validation of new models is simulation of the dMRI signal using well defined and controllable ground truth microstructural environments known as numerical phantoms which allow us to compare model fitting results to the ground truth from the numerical phantoms. Recent studies have used this approach to investigate grey matter [5], however numerical phantoms have more commonly been used to study dMRI in white matter (WM) [6–14].

Another application of numerical phantoms which is growing in popularity is in the direct computational modelling of microstructure. These techniques use machine learning or fingerprinting-style techniques to match simulated signals and the corresponding ground truth microstructure of the numerical phantom to the measured signal in order to directly estimate microstructural features without using an explicit analytical model [15–18].

As well as affecting the dMRI signal, tissue microstructure also affects other MR techniques such as susceptibility-weighted imaging [19, 20]. For instance, simulations have been used to show that using realistic WM axonal models affect the susceptibility weighted MR image differently to the commonly assumed straight cylinders [21]. It is, therefore, important to the wider MRI community, not just the dMRI community, to generate realistic WM numerical phantoms which accurately capture microstructural features in order to get realistic simulated signal.

Typically, however, there is a mismatch between the true complexity of brain tissue and the numerical phantoms that are used in simulations, outlined in Figure 1.1. For instance, WM fibres are commonly represented using simple geometric constructs such as straight, parallel cylinders [14, 22–26]. Real WM fibres, however, have complex shapes with morphological features such as undulation and diameter variation, as revealed by electron microscopy (EM) studies [27, 28].

It is very difficult to know exactly how this complex microstructure impacts on the dMRI signal and to understand which of these complex morphological features can be detected using dMRI. Developing new numerical phantoms which allow us to generate realistic WM microstructure would enable us to study these effects with more control and detail as well as enabling us to validate new and existing dMRI microstructure models more comprehensively.

Additionally, realistic WM numerical phantoms would enable us to drive forward a new generation of computational models of WM, using machine learning to link simulated signals to the realistic WM microstructure which could potentially improve the accuracy of microstructure models. For instance, it may be possible to disentangle axonal loss and demyelination using these computational models which can give us valuable information for differential diagnosis in diseases such as multiple sclerosis (MS).

The first step on the road to developing these new computational models, however, is to develop a method to generate realistic WM numerical phantoms with controllable microstructure. That is the main purpose of the work presented in this thesis, the development of a new tool to generate more realistic WM numerical phantoms than was previously possible.

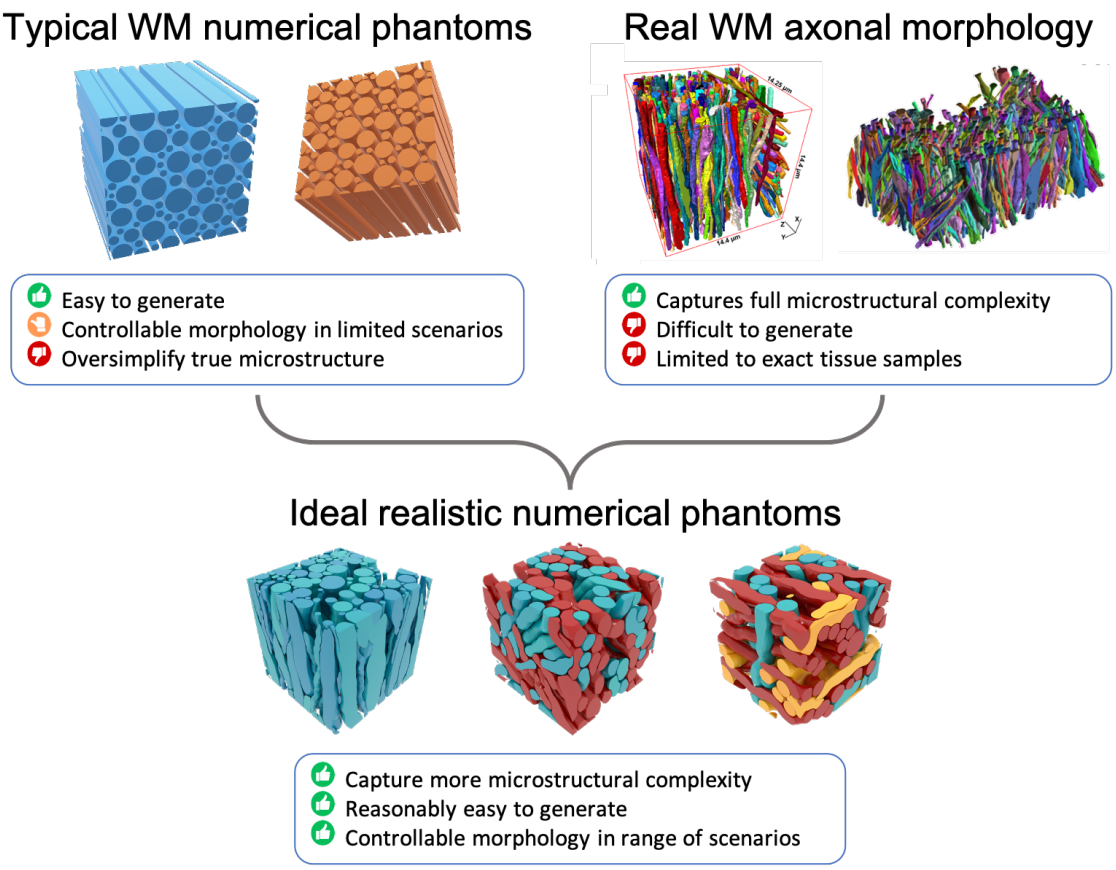

Figure 1.1: The work presented in this abstract aims to bridge the gap between the true complexity of white matter axonal morphology and the simplistic representation currently used as numerical phantoms.

1.2 Problem Statement

There is a need to be able to generate numerical phantoms that realistically represent WM microstructure, with controllable microstructural parameters such as axon packing density and orientation dispersion.

1.3 Project Aims and Scope

This report summarises work towards improving the realism of simulations of diffusion MRI. Realistic simulations allow models of the MRI signal to be validated using controllable and well known ground truth.

The main aims of this work are as follows:

1. Develop a method for generating realistic WM numerical phantoms with realistic and controllable morphology including axon packing densities and orientation dispersion

2. Test the realism of these numerical phantoms by comparing microstructure with real microstructure reconstructed from electron microscopy and by comparing simulated dMRI signal to real dMRI signals from WM
3. Demonstrate the usefulness of the phantom generator by using WM numerical phantoms to probe the intra-axonal diffusion and assumptions in spherical deconvolution based dMRI modelling techniques.

In an effort to achieve these aims, a method called ConFiG (**Con**textual **Fi**bre **G**rowth) is presented which 'grows' fibres densely, mimicking natural fibre-genesis while attempting to respect some morphological priors to generate realistic WM numerical phantoms.

1.4 Report Overview

This report is arranged into four parts and is based on a number of publications produced over the course of this project[2]:

I - Introduction and Background This part includes the current chapter, outlining the motivation behind this project and setting out the aims. Chapter 2 outlines some background for the project, explaining the physics of diffusion MRI and how we simulate it while Chapter 3 reviews current literature on dMRI simulations and numerical phantoms to establish the state-of-the-art and identify some limitations in existing numerical phantom generation techniques - primarily that they over-simplify real axonal geometry.

II - Methods Development This part describes the methodological improvements for phantom generation developed during this project, tackling aims 1 & 2 outlined above. Working towards these aims, we developed a novel phantom generation technique based on growing digital fibres, trying to mimic natural fibre growth. A preliminary version of the phantom generator, preliminary fibre growth (preFiG), is presented in Chapter 4 and is based on work published in *Callaghan et al., IPMI, 2019*, *Callaghan et al., ISMRM, 2019* and *Callaghan et al., OHBM, 2019*.

The preFiG approach produced at-that-time state-of-the-art numerical phantoms, however the maximum achievable fibre density was too low to apply to many WM regions. To improve upon the preFiG phantom generator, we introduced new growth mechanisms inspired by real biological axonal growth mechanisms. This resulted in the final phantom generator, ConFiG, which is presented in Chapter 5, with validation of the realism of ConFiG phantoms presented in Chapter 6 by comparing ConFiG phantoms to real axons. Chapters 5 and 6 are an expanded version of the work published in *Callaghan et al., NeuroImage, 2020*.

[2]For a full list of publications produced during this project, see the **Publications** chapter at the end of this report (page 147).

III - Applications This part contains a series of experiments performed using ConFiG phantoms to probe diffusion dynamics and modelling assumptions in realistic WM phantoms, addressing the third aim outlined above. Chapter 7 presents experiments investigating the intra-axonal diffusivity within realistic axons, demonstrating that undulation and beading (variable diameter along axons) affect the axial diffusivity in different diffusion time regimes. We further investigate what impact realistic axonal morphology and orientation dispersion can have on the estimation of the axial diffusivity from dMRI. This chapter is based in part on results presented in *Callaghan et al., ISMRM, 2019*.

Chapter 8 presents an investigation into how variable axonal morphology within a voxel can impact a family of dMRI techniques based on spherical deconvolution (SD). We demonstrate that variable axon morphology causes the dMRI signal to differ across axons and that this can impact SD techniques which assume all axons give the same signal. This chapter is based on work that has been submitted for publication as a journal article and is under review.

IV - Conclusions The final part presents some concluding remarks, with Chapter 9 discussing the contributions to the field made by the work done during this project and future directions that are enabled by ConFiG that we can take to push the microstructure imaging field forward

2 Background

Chapter Contents

2.1 MR Physics 25

2.2 Diffusion MRI 35

2.3 Diffusion Simulation 38

2.4 Diffusion Modelling 41

Chapter Summary

This chapter introduces the physics behind dMRI and some techniques used for the simulation and modelling of dMRI. Firstly, general MR physics is introduced, building up from a single proton to the generation of the spin echo signal. The second section discusses how the diffusion of water molecules impacts the MR signal, leading to the diffusion signal attenuation. The third section introduces some simulation approaches which can be used to generate synthetic dMRI data. The final section gives a brief overview of the kinds of dMRI models typically used in the brain with a focus on spherical deconvolution techniques which will be used later in the thesis.

2.1 MR Physics

All forms of *in vivo* magnetic resonance have their origins in the 1940s when Purcell, Torrey and Pound independently and almost simultaneously with Bloch, Hansen and Packard detected radio frequency signals from nuclei in ordinary matter [29–32]. This discovery gave birth to the field of nuclear magnetic resonance (NMR) which has become widespread, with applications in a number of areas including chemistry, biology, materials science and medical imaging [30, 33].

Within the medical imaging context, NMR typically finds two uses, magnetic resonance imaging (MRI) and magnetic resonance spectroscopy (MRS)[1]. Both of these uses are closely related: MRI is typically concerned with building images of internal structures in the body, whilst MRS is concerned with identifying the chemical composition of tissues in the body.

The theory behind NMR concerns the interaction between nuclei and magnetic fields. This section briefly introduces the NMR physics relevant to dMRI.

2.1.1 Nuclear Magnetism

The most important property of a nucleus for the application of NMR is nuclear spin. Spin is a property inherent to all subatomic particles and whilst it is a purely quantum effect, it can be thought of loosely as the particle spinning around its axis - much like a tiny planet [29]. A planet spinning about its axis will have an angular momentum associated with that rotation and similarly the spin of a particle behaves like an angular momentum. Unlike the angular momentum of a rotating planet, however, the spin is an intrinsic property of the particle itself and not a result of its motion [29].

MRI typically relies on NMR of ^{1}H nuclei in water molecules, which simply consist of a single proton. A proton carries a positive electric charge. Just as classically a rotating charge with angular momentum, $\mathbf{L}$, will produce a magnetic moment, $\boldsymbol{\mu} = \gamma\mathbf{L}$, the intrinsic spin angular momentum, $\mathbf{S}$ of a proton will produce a magnetic moment

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

where γ is the gyromagnetic ratio [29]. For a proton, $\gamma = 2.675 \times 10^8$ rad s^{-1} T^{-1}.

The fact that a proton has an intrinsic magnetic moment means that it will interact with magnetic fields and it is understanding this interaction that underpins NMR theory.

2.1.2 Magnetic Resonance

Classically, a magnetic moment, $\boldsymbol{\mu}$, placed in an external magnetic field, $\mathbf{B}_0$, will feel a torque, $\boldsymbol{\tau}$, given by [34]

$$\boldsymbol{\tau} = \boldsymbol{\mu} \times \mathbf{B}_0\,. \tag{2.2}$$

At the same time, classical mechanics gives a relationship between the change in angular momentum and the torque as [34]

$$\frac{d\mathbf{L}}{dt} = \boldsymbol{\tau}\,. \tag{2.3}$$

For a proton in its rest state, the only angular momentum is the intrinsic spin angular momentum, $\mathbf{S}$, so combining Equations (2.2) and (2.3) gives the equation of motion for a spin in an external magnetic field

$$\frac{d\mathbf{S}}{dt} = \boldsymbol{\mu} \times \mathbf{B}_0\,. \tag{2.4}$$

[1]The 'N' from NMR is dropped in the medical imaging context to avoid confusion with nuclear medicine and general squeamishness around the word nuclear

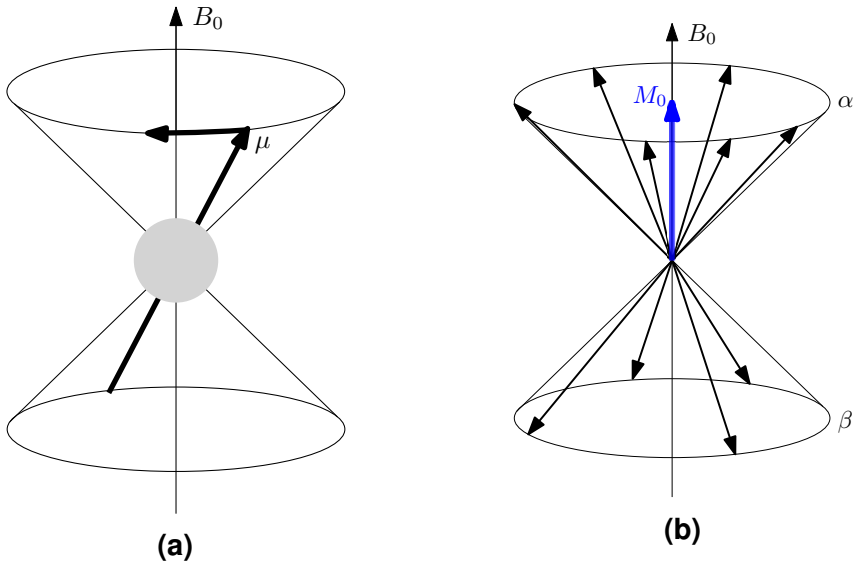

Figure 2.1: a) Illustration of the precessional motion of a single spin with magnetic moment μ in the presence of an external magnetic field B_0. b) Many individual spins in an external magnetic field precess around the external field with random phase producing a net magnetisation in the direction of the B_0 field.

Since S is equivalent to $\boldsymbol{\mu}/\gamma$ (Equation (2.1)), this becomes

$$\frac{d\boldsymbol{\mu}}{dt} = \gamma\boldsymbol{\mu} \times \mathbf{B}_0 \,. \tag{2.5}$$

This equation of motion can be solved in a few ways for the case of constant external magnetic field, with the result being that the magnetic moment precesses about the magnetic field with a frequency, ω_0, given by [29]

$$\omega_0 = \gamma B_0 \,, \tag{2.6}$$

with B_0 being the external field strength. This precessional motion is illustrated in Figure 2.1a. The frequency ω_0 is known as the Larmor frequency and lies in the radiofrequency (RF) range for typical field strengths found in MRI machines (1.5 - 7 T).

In practice, it is not possible to observe the magnetic moment of a single spin *in vivo*. The quantity observed is rather the sum of the magnetic moments from many spins together, this is known as the net magnetisation.

In a sample, slightly more protons will align with the $\mathbf{B}_0$ field than against it, meaning that the net magnetisation will be parallel to $\mathbf{B}_0$. Figure 2.1b shows a pictorial representation of the system of many spins producing a net magnetisation, $\mathbf{M}_0$ aligned with $\mathbf{B}_0$.

Each of the spins will still be precessing about the magnetic field at the Larmor frequency but since they are out of phase with one another, all transverse components of the magnetisation cancel out when they combine and all that is left is a static longitudinal component.

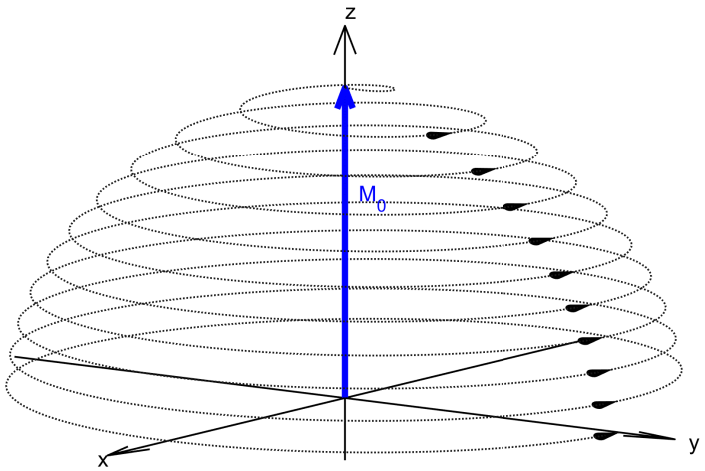

Figure 2.2: Nutation motion of an on-resonance spin in the presence of an RF field. Precession about the B_0 and B_1 fields create the spiralling motion in the laboratory frame.

In order to make measurements, the net longitudinal magnetisation needs to be 'flipped' into the transverse plane where it can be detected. This is achieved by applying a second magnetic field, $\mathbf{B}_1$, oscillating in the transverse plane. In much the same way as with $\mathbf{B}_0$, the magnetisation feels a torque from $\mathbf{B}_1$ and begins to rotate about $\mathbf{B}_1$, away from the longitudinal axis. The two external fields act simultaneously on $\mathbf{M}_0$ so the magnetisation will tip away from the z axis whilst still precessing about z with a frequency ω_0. This kind of motion is known as nutation and is illustrated in Figure 2.2.

2.1.3 The Bloch Equations

The interaction between the magnetisation and magnetic fields is described by the Bloch equations - an empirical set of equations describing the evolution of magnetisation introduced by Felix Bloch in 1946 [31].

The magnetisation arises from a sum of independent magnetic moments, meaning that we can represent the magnetisation as

$$\mathbf{M} = \sum_i \boldsymbol{\mu}_i \,, \tag{2.7}$$

with i indicating a sum over all the spins in the sample. This definition for $\mathbf{M}$ can be combined with the equation of motion for a single spin, Equation (2.5), to give [34]

$$\frac{d\mathbf{M}}{dt} = \gamma \mathbf{M} \times \mathbf{B} \,. \tag{2.8}$$

In the presence of the main external field, $\mathbf{B}_0$, the magnetisation will be static and aligned along the z axis. The x and y components of the magnetisation will have random orientations and precess about $\mathbf{B}_0$ at the Larmor frequency with a mean amplitude of zero. This will give the components of Equation (2.8) as [35]

$$\frac{d\mathbf{M}_x(t)}{dt} = \gamma \mathbf{M}_y \mathbf{B}_0 \,, \tag{2.9}$$

$$\frac{d\mathbf{M}_y(t)}{dt} = -\gamma \mathbf{M}_x \mathbf{B}_0 \,, \tag{2.10}$$

$$\frac{d\mathbf{M}_z(t)}{dt} = 0 \,. \tag{2.11}$$

To understand the interaction of the magnetisation with the $\mathbf{B}_1$ field, the oscillation of the field in the transverse plane needs to be described. Usually, the $\mathbf{B}_1$ field is circularly polarised to oscillate in the transverse plane so that the field can be described as

$$\mathbf{B}_1(t) = B_1 \cos(\omega t)\hat{\mathbf{x}} - B_1 \sin(\omega t)\hat{\mathbf{y}} \,, \tag{2.12}$$

where $\hat{\mathbf{x}}$ and $\hat{\mathbf{y}}$ are unit vectors in the x and y directions respectively.

The combined effect of the $\mathbf{B}_0$ and $\mathbf{B}_1$ fields can be seen from Equation (2.8) to get [35]

$$\frac{d\mathbf{M}_x(t)}{dt} = \gamma \left(\mathbf{M}_y(t)\mathbf{B}_0 - \mathbf{M}_z(t)\mathbf{B}_1 \sin(\omega t) \right) \,, \tag{2.13}$$

$$\frac{d\mathbf{M}_y(t)}{dt} = \gamma \left(\mathbf{M}_z(t)\mathbf{B}_1 \cos(\omega t) - \mathbf{M}_x(t)\mathbf{B}_0 \right) \,, \tag{2.14}$$

$$\frac{d\mathbf{M}_z(t)}{dt} = \gamma \left(\mathbf{M}_x(t)\mathbf{B}_1 \sin(\omega t) - \mathbf{M}_y(t)\mathbf{B}_1 \cos(\omega t) \right) \,. \tag{2.15}$$

These are the equations of motion of the magnetisation in the laboratory frame under the influence of the $\mathbf{B}_0$ and $\mathbf{B}_1$ and describe the kind of motion seen in Figure 2.2.

2.1.3.1 Relaxation

In order to get to the full Bloch Equations the concept of relaxation must be introduced. Relaxation is a term used to describe the way in which a spin system will return to equilibrium after being perturbed. The components of $\mathbf{M}$ that are parallel to the B_0 magnetic field relax differently to those perpendicular to the magnetic field leading to two relaxation terms being introduced into Equations (2.13) to (2.15).

The relaxation processes are exponential and described by two time constants, T_1 and T_2. T_1 is the longitudinal relaxation time and describes the rate at which longitudinal magnetisation regrows after a perturbation. T_2 is the transverse relaxation time and describes the rate at which transverse magnetisation decays after a perturbation. T_2 is always shorter than T_1 since all the effects which contribute to T_1 also contribute to T_2

relaxation, however T_2 relaxation is also affected by the spins going out of phase with one another. The relaxation process can be written as [35]

$$\frac{d\mathbf{M}_x(t)}{dt} = -\frac{\mathbf{M}_x(t)}{\mathrm{T}_2}\,, \tag{2.16}$$

$$\frac{d\mathbf{M}_y(t)}{dt} = -\frac{\mathbf{M}_y(t)}{\mathrm{T}_2}\,, \tag{2.17}$$

$$\frac{d\mathbf{M}_z(t)}{dt} = -\frac{\mathbf{M}_z(t) - \mathbf{M}_0}{\mathrm{T}_1}\,. \tag{2.18}$$

Combining Equations (2.13) to (2.15) and Equations (2.16) to (2.18) gives the full Bloch equations

$$\frac{d\mathbf{M}_x(t)}{dt} = \gamma\left(\mathbf{M}_y(t)\mathbf{B}_0 - \mathbf{M}_z(t)\mathbf{B}_1\sin(\omega t)\right) - \frac{\mathbf{M}_x(t)}{\mathrm{T}_2}\,, \tag{2.19}$$

$$\frac{d\mathbf{M}_y(t)}{dt} = \gamma\left(\mathbf{M}_z(t)\mathbf{B}_1\cos(\omega t) - \mathbf{M}_x(t)\mathbf{B}_0\right) - \frac{\mathbf{M}_y(t)}{\mathrm{T}_2}\,, \tag{2.20}$$

$$\frac{d\mathbf{M}_z(t)}{dt} = \gamma\left(\mathbf{M}_x(t)\mathbf{B}_1\sin(\omega t) - \mathbf{M}_y(t)\mathbf{B}_1\cos(\omega t)\right) - \frac{\mathbf{M}_z(t) - \mathbf{M}_0}{\mathrm{T}_1}\,. \tag{2.21}$$

T_2 is used to refer to relaxation due to intrinsic spin-spin interactions which cause spins to accrue phase relative to one another and thus the magnitude of the net transverse magnetisation is reduced when taking the sum in Equation (2.7) . Other effects can also contribute to the loss of transverse magnetisation, such as magnetic field inhomogeneities which can add to the T_2 relaxation. This is referred to as T_2^*, with

$$\frac{1}{\mathrm{T}_2^*} = \frac{1}{\mathrm{T}_2} + \frac{1}{\mathrm{T}_2'}\,, \tag{2.22}$$

where T_2' is the relaxation time associated with these external sources and T_2 is the intrinsic spin-spin relaxation time. The T_2' effect can be negated using special MR pulse sequences which will be covered in Section 2.1.5, however the intrinsic T_2 relaxation cannot be avoided.

2.1.3.2 The rotating frame

To this point, everything has been described in a static Cartesian frame known as the laboratory frame. The lab frame is not the most convenient reference frame to analyse the NMR experiment in, however. Moving to a frame which is rotating about $\mathbf{B}_0$ (i.e. the z-axis) at a frequency ω matching the $\mathbf{B}_1$ field oscillation simplifies the maths of the system. The axes of this rotating frame will be referred to as x', y' and z'.

The components of the magnetisation in the rotating frame can be calculated from the lab frame components as [35]

$$\mathbf{M}_x' = \mathbf{M}_x\cos(\omega t) - \mathbf{M}_y\sin(\omega t)\,, \tag{2.23}$$

$$\mathbf{M}_y' = \mathbf{M}_x\sin(\omega t) - \mathbf{M}_x\cos(\omega t)\,, \tag{2.24}$$

$$\mathbf{M}_z' = \mathbf{M}_z\,. \tag{2.25}$$

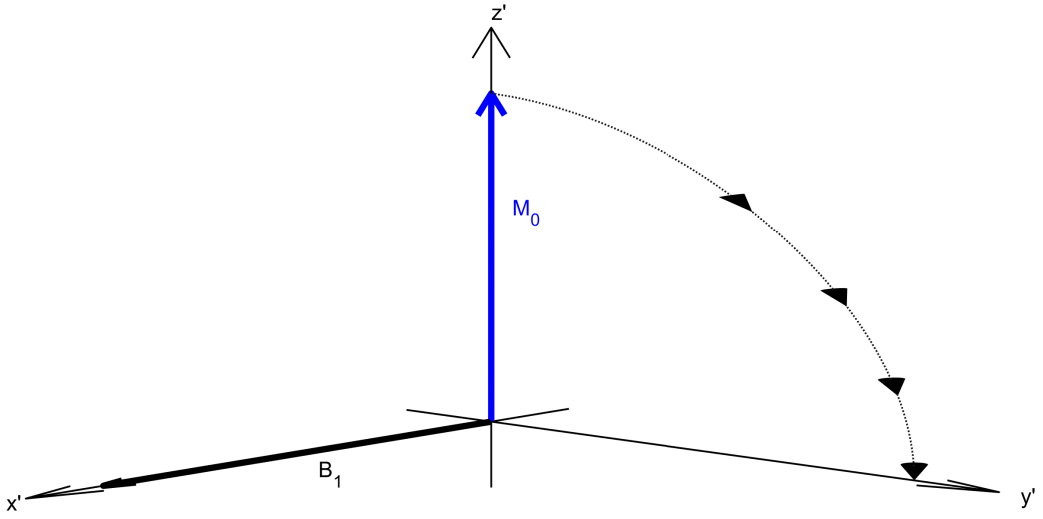

Figure 2.3: Motion of a spin in the presence of a B_1 RF field in the rotating frame. This is identical to the nutation in Figure 2.2, however viewing from the rotating frame simplifies the motion.

The rotating frame Bloch equations can be calculated by combining these rotating frame magnetisation components with the lab frame Bloch equations [35]

$$\frac{d\mathbf{M}'_x(t)}{dt} = \Omega\mathbf{M}'_y(t) - \frac{\mathbf{M}'_x(t)}{\mathrm{T}_2}\,, \tag{2.26}$$

$$\frac{d\mathbf{M}'_y(t)}{dt} = -\Omega\mathbf{M}'_x(t) + \gamma\mathbf{B}_1\mathbf{M}'_z(t) - \frac{\mathbf{M}'_y(t)}{\mathrm{T}_2}\,, \tag{2.27}$$

$$\frac{d\mathbf{M}'_z(t)}{dt} = -\gamma\mathbf{B}_1\mathbf{M}'_y(t) - \frac{\mathbf{M}'_z(t) - \mathbf{M}_0}{\mathrm{T}_1}\,, \tag{2.28}$$

where $\Omega = \omega_0 - \omega$ is the offset frequency between the $\mathbf{B}_1$ field frequency and the Larmor frequency.

Since the frame is rotating with a frequency ω, the $\mathbf{B}_1$ field appears static in the rotating frame. The precessional motion that is seen in the lab frame ($\omega_0 = \gamma\mathbf{B}_0$) is reduced to a frequency Ω in the rotating frame. When $\Omega = 0$, meaning that $\mathbf{B}_1$ oscillates at the Larmor frequency, the magnetisation simply precesses about the $\mathbf{B}_1$ field towards the transverse plane as illustrated in Figure 2.3.

This situation is known as resonance - the frequency of the RF pulse matches the Larmor frequency, perfectly tipping the magnetisation away from the z' axis and into the transverse plane.

In the off-resonance case, an additional component of magnetic field with magnitude Ω/γ is produced in the z-direction. This results in an effective magnetic field, $\mathbf{B}_e$, with a magnitude [35]

$$B_e = |B_e| = \sqrt{B_1^2 + \left(\frac{\Omega}{\gamma}\right)^2}\,. \tag{2.29}$$

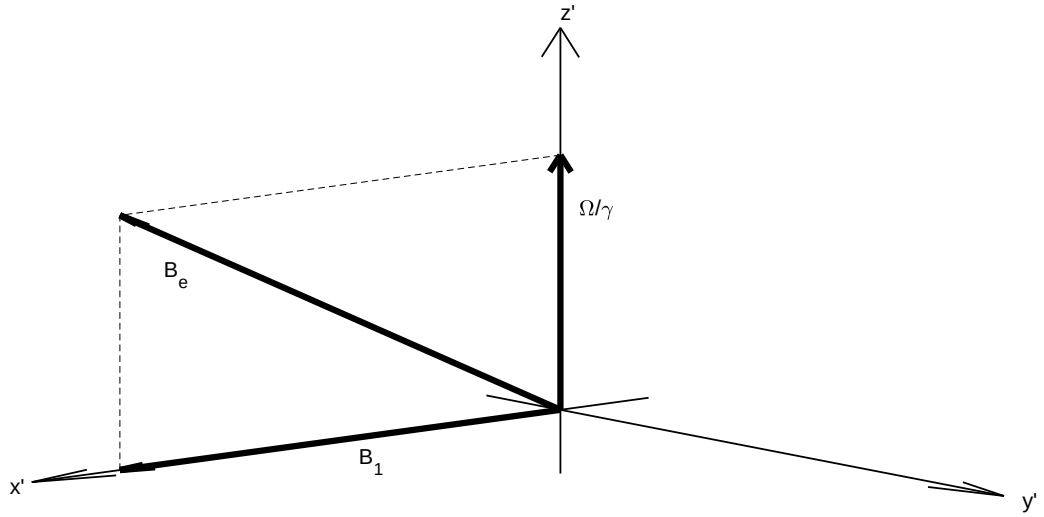

Figure 2.4: The effective field, B_e, produced due to an off-resonance frequency Ω. The off-resonance effects produce an additional component of magnetic field along the z' axis.

The effective field is illustrated in Figure 2.4 with the additional component of Ω/γ resulting in an effective field that is no longer aligned with x'. Off-resonance effects can produce unwanted results meaning the spin does not get flipped as much as expected under an RF pulse which can result in signal losses.

2.1.4 Detecting the MR signal

The detection and processing of NMR signals is a deep topic which could be the subject of its own book, however some very basic details of how a signal is formed are useful to go on from here.

The reason for flipping the magnetisation into the transverse plane using $\mathbf{B}_1$ fields is to make the magnetisation detectable. Transverse magnetisation precesses about $\mathbf{B}_0$ at the Larmor frequency, sweeping its magnetic field around $\mathbf{B}_0$. A coil of wire placed near this precessing field will feel an electromotive force induced in it according to Faraday's Law of Induction [34].

Following a pulse that flips the magnetisation from $\mathbf{M}_0$ aligned with z through an angle β towards x', the x'-component of the magnetisation will be $M_0 \sin\beta$ and (ignoring relaxation) will then precess at the offset frequency, Ω, in the rotating frame. This will give the components of the magnetisation in the transverse plane over time as

$$M_x = M_0 \sin(\beta)\cos(\Omega t) \qquad M_y = M_0 \sin(\beta)\sin(\Omega t)\,. \tag{2.30}$$

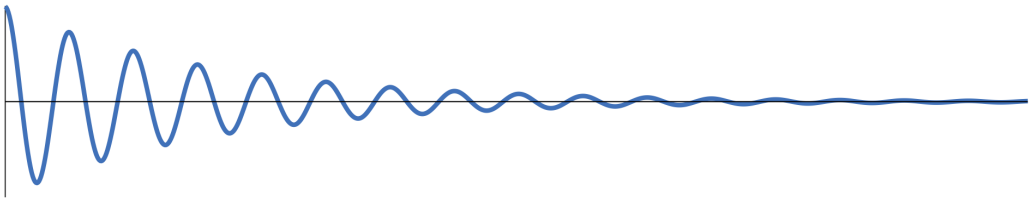

Figure 2.5: The free induction decay described by Equation (2.32). Here, just the real channel is plotted.

The signal induced into the receiver coils is proportional to M_x and M_y and so the signal will also have an oscillating form similar to Equation (2.30). From the $\sin\beta$ term, it is clear that the maximum signal will arise when $\beta = 90°$, meaning all the magnetisation is flipped into the transverse plane. Additionally, in a realistic experiment, there will be T_2^* relaxation so including this, the general form of the signal following a 90° pulse will be [35]

$$S_x = S_0 \cos(\Omega t) \exp\left(-t/\mathrm{T}_2^*\right) \qquad S_y = S_0 \sin(\Omega t) \exp\left(-t/\mathrm{T}_2^*\right) . \tag{2.31}$$

Generally NMR systems use something known as quadrature detection, meaning that both the x' and y' components of the magnetisation are measured simultaneously [29], giving the signal as a function of time as

$$\begin{aligned} S(t) &= S_x + iS_y \,, \\ &= S_0 \exp((i\Omega - 1/\mathrm{T}_2^*)t) \,. \end{aligned} \tag{2.32}$$

This time-domain signal is known as a free induction decay (FID) and has a typical form shown in Figure 2.5. If we neglect off-resonance effects, then the FID in the rotating frame will be a simple exponential decay.

$$S(t) = S_0 \exp(-t/\mathrm{T}_2^*) \tag{2.33}$$

The FID is not commonly used for dMRI for a few reasons. Firstly, magnetic field gradients need to be introduced to make the signal sensitive to diffusion. Additionally, the $\mathbf{T}_2^*$ decay is often very rapid, so sequences known as spin echo sequences are used to remove the $\mathbf{T}_2'$ relaxation.

2.1.5 Spin echoes

It is possible to undo the effects of T_2' by designing a pulse sequence to 'refocus' the spins, forming what is known as a spin echo. The first spin echo sequence was introduced by Edwin Hahn in 1950 [36]. The simplest sequence to form a spin echo consists of a 90° pulse to excite the spins followed by a 180° pulse after a delay. This sequence is shown in the diagram in Figure 2.6 along with how the signal varies during the pulse sequence.

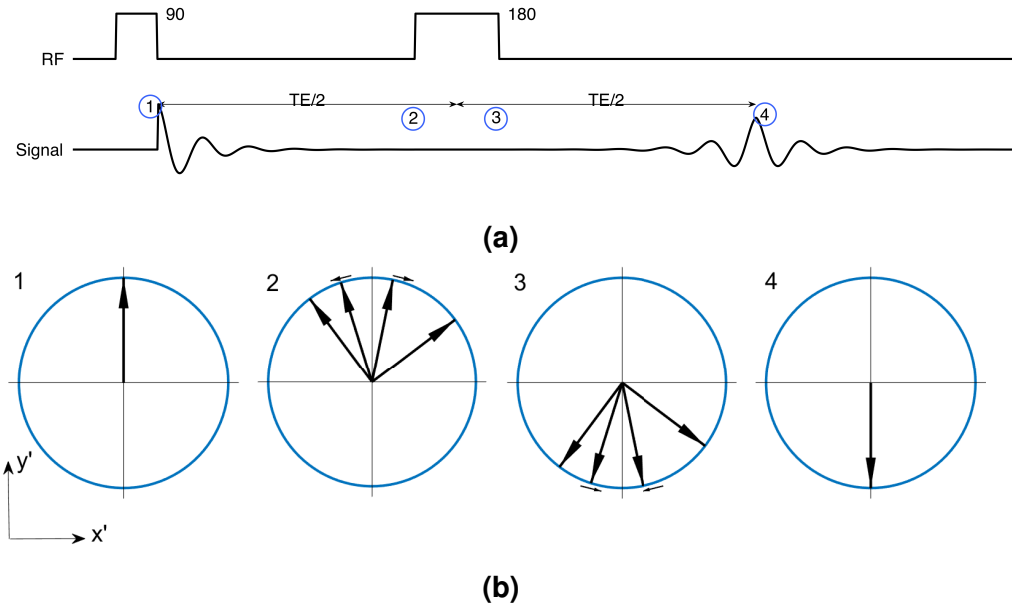

Figure 2.6: a) Spin echo sequence and (b) an indication of the evolution of spins under a spin echo sequence. This shows how the 180° refocusing pulse acts to refocus the spins after a time TE.

Figure 2.6b represents how the magnetisation evolves through the pulse sequence with the four diagrams corresponding to the points marked in Figure 2.6a. At point 1, immediately following the 90° pulse, all the magnetisation has been flipped into the transverse plane and is in phase - meaning all the magnetic moments of the spins point in the same direction in the x-y plane.

The field inhomogeneities cause the different spins to feel slightly different magnetic fields and so precess at slightly different frequencies. This causes the spins to lose phase-coherence as indicated at point 2, and so the signal decays with T_2^*.

At point 3, following the 180° pulse, the spins remain out of phase with one another, but the 180° pulse has flipped their orientations across the x'-axis. The magnetic field the spins feel is still the same, so despite their flip in orientation, they still precess in the same direction. This means that the evolution that caused the spins to dephase begins to rewind and bring the spins back into phase coherence. After a time equal to the time between the 90° and 180° pulses, at point 4, the spins will be brought completely back in phase - or, refocused - and the spin echo is formed.

The signal at point 4 will still be lower in magnitude than that at point 1 since the T_2 relaxation will still occur as it is an inherent property of the matter. The spin echo does, however, refocus the $\mathbf{B}_0$ inhomogeneities. The time between the 90° pulse and the formation of the echo is known as the echo time (TE).

The spin echo sequence shown in Figure 2.6 forms the basis of the standard pulsed gradient spin echo (PGSE) diffusion MRI sequence which is introduced in the following section, along with a description of the physics behind diffusion MRI.

2.2 Diffusion MRI

Diffusion MRI (dMRI) sensitises the MRI signal to the motion of water molecules due to diffusion. The following section describes the physics behind diffusion and how the diffusion impacts the MRI signal.

The diffusion process is driven by the Brownian motion of particles in fluids. The thermal kinetic energy of particles causes them to move around rapidly, however particles frequently collide with each other (for instance, molecules in water at room temperature experience around 60 billion collisions per second [37]) creating a very tortuous, random path.

Diffusion MRI sensitises the MR signal to this motion by exploiting the dephasing of spins as a result of magnetic field gradients.

The magnetic field will generally have a uniform component from the main B_0 field, and spatially and/or time varying components due to deliberate magnetic field gradients or typically unwanted effects such as magnetic susceptibility inhomogeneities and concomitant fields [34]. In general, $B(\mathbf{r}, t)$, the magnitude of the magnetic field at a position $\mathbf{r}$ at time t is given by

$$B(\mathbf{r}, t) = |\mathbf{B}| = |B_0\hat{\mathbf{z}} + \boldsymbol{\Delta}\mathbf{B}(\mathbf{r}, t)|\,, \tag{2.34}$$

where $\boldsymbol{\Delta}\mathbf{B}(\mathbf{r}, t)$ accounts for all of the variation in the magnetic field away from B_0. Note that $\boldsymbol{\Delta}\mathbf{B}(\mathbf{r}, t)$ is a vector quantity which may have components in the $\hat{\mathbf{x}}$ and $\hat{\mathbf{y}}$ directions.

An idealised expression for $\boldsymbol{\Delta}\mathbf{B}(\mathbf{r}, t)$ often applied to MRI assumes that all of the change in the magnetic field is due to an applied magnetic field gradient, $\mathbf{g}(\mathbf{r}, t)$, which only has a significant $\hat{\mathbf{z}}$ component. This means that Equation (2.34) can be written as

$$\begin{aligned} B(\mathbf{r}, t) &= |B_0\hat{\mathbf{z}} + (\mathbf{g}(\mathbf{r}, t) \cdot \mathbf{r})\,\hat{\mathbf{z}}|\,, \\ &= B_0 + \mathbf{g}(\mathbf{r}, t) \cdot \mathbf{r}\,. \end{aligned} \tag{2.35}$$

Magnetic field gradients introduce a deliberate variation in the magnetic field which, according to Equation (2.6), causes the Larmor frequency to vary spatially as well temporally.

Since the Larmor frequency varies spatially, spins in different locations will precess at different frequencies and accrue a phase shift relative to spins in different locations. The incremental phase, $d\phi$, accrued for a single spin, i, in an infinitesimal time, dt, is given by

$$d\phi_i = \gamma B(\mathbf{R}_i(t), t)dt\,, \tag{2.36}$$

where γ is the gyromagnetic ratio and $\mathbf{R}_i(t)$ is the position of the particle at time t.

Putting Equation (2.35) into Equation (2.36) and integrating over the time of the diffusion experiment will give the total phase accrued for a single spin:

$$\phi(\mathbf{R}_i(t)) = \gamma B_0 t + \gamma \int_0^t \mathbf{g}(\mathbf{R}_i(t'), t') \cdot \mathbf{R}_i(t')dt'\,, \tag{2.37}$$

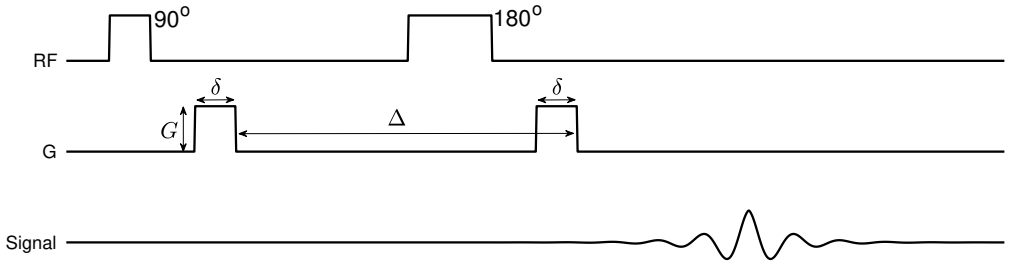

Figure 2.7: The standard pulsed gradient spin echo sequence used in dMRI.

The first term in this equation is the phase accrued due to the main magnetic field which will be the same for all spins in the system. The second term is the phased accrued due to the gradient, which will be dependent on the motion of each individual spin. The dot product here indicates that only displacement projected onto the gradient direction affects the phase, allowing the gradient direction to be used to probe the diffusion in different directions.

The first diffusion MR sequence, introduced by Stejskal and Tanner in 1965 [38], is the pulsed gradient spin echo (PGSE) sequence, shown in Figure 2.7. The PGSE sequence consists of a standard spin echo sequence with a pair of gradient pulses added either side of the refocusing pulse. In the ideal case, each pulse is rectangular with a gradient strength, G, and duration, δ and they are separated by a time, Δ.

The effect of this pulse sequence can be simplified by considering the case when $\delta \ll \Delta$. This is known as the short gradient pulse (SGP) approximation and means that the motion of spins during the pulses can be ignored.

Under the SGP approximation, the phase accrued by a spin at a position $\mathbf{R}_0$ during a pulse at a time t_0 will be

$$\phi(\mathbf{R}_0) = \gamma B_0 \delta + \gamma \delta \mathbf{g}(\mathbf{R_0}) \cdot \mathbf{R_0} \,. \tag{2.38}$$

Here, $\mathbf{R}_0$ refers to the spin position at time $t = t_0$ and is not time-dependent, because under the SGP approximation the motion of spins during the pulse is assumed to be negligible.

The 180° pulse in the PGSE sequence is crucial, since it flips the orientation of the spins' magnetic moments as in the spin echo sequence in Figure 2.6. Not only does this refocus T_2^* effects as outlined in Section 2.1.5, but similarly, it means that the phase accrued due to the first gradient pulse now has a negative sign relative to the phase that will be accrued due to the second pulse.

A spin which is at a position $\mathbf{R}_0$ during the first pulse and then diffuses to a position $\mathbf{R}_1$ during the second pulse will therefore have a relative phase shift of

$$\Delta\phi(\mathbf{R}_1 - \mathbf{R}_0) = \gamma \delta \mathbf{g} \cdot (\mathbf{R}_1 - \mathbf{R}_0) \,. \tag{2.39}$$

The B_0 term from Equation (2.38) is the same for both gradient pulses, meaning that during the subtraction in Equation (2.39) it cancels and the only relative phase shift comes from the diffusive motion.

The total MR signal comes from an ensemble of spins, each with their own random Brownian motion and thus, from Equation (2.39), their own relative phase shift. To get to the total MR signal, we need to consider the probability that a particle starts at position $\mathbf{R}_0$ (i.e. the initial spin density, $\rho(\mathbf{R}_0)$), which can generally be considered uniform within a voxel [39]) and the probability that a particle which starts at $\mathbf{R}_0$ moves to $\mathbf{R}_1$ during the time Δ, $P(\mathbf{R}_0, \mathbf{R}_1, \Delta)$. Putting these together, gives an expression for the total MR signal [38, 39]:

$$S(\mathbf{g}, \Delta) = S(\mathbf{0}, \Delta) \int\int \rho(\mathbf{R}_0) P(\mathbf{R}_0, \mathbf{R}_1, \Delta) e^{i\gamma\delta\mathbf{g}\cdot(\mathbf{R}_1 - \mathbf{R}_0)} d\mathbf{R}_0 d\mathbf{R}_1 \,. \tag{2.40}$$

This quantity, $P(\mathbf{R}_0, \mathbf{R}_1, \Delta)$, is known as the diffusion propagator and is of great interest for diffusion MRI because $P(\mathbf{R}_0, \mathbf{R}_1, \Delta)$ encodes the information about the environment in which the spins are diffusing. For diffusion in an isotropic, homogeneous medium, the diffusion propagator is a Gaussian distribution [39]:

$$P(\mathbf{R}_0, \mathbf{R}_1, t) = (4\pi D t)^{-3/2} \exp\left(-\frac{(\mathbf{R}_1 - \mathbf{R}_0)^2}{4Dt}\right) . \tag{2.41}$$

In the case of Gaussian diffusion, Equation (2.40) can be solved analytically and will give an MR signal attenuation (that is, $S(t)/S(0)$) which is also Gaussian [38, 39]

$$E(g, \Delta) = \exp(-\gamma^2 g^2 \delta^2 D \Delta) \,. \tag{2.42}$$

The general form of this expression, accounting for finite duration gradient pulses, can also be analytically derived to give the Stejskal-Tanner equation [38, 40]

$$\ln(E) = -\gamma^2 g^2 \delta^2 D(\Delta - \delta/3) \,, \tag{2.43}$$

$$= -bD \,, \tag{2.44}$$

where $b = \gamma^2 g^2 \delta^2 (\Delta - \delta/3)$ is the so-called b-value which describes the strength of the diffusion encoding.

We can formulate a more general form of Equation (2.40) without requiring any assumptions on the gradients (such as the SGP approximation used above) as [25, 39]

$$S(\mathbf{g}, t) = S(\mathbf{0}, t) \int_{-\infty}^{\infty} P(\phi, t) e^{i\phi} d\phi \,, \tag{2.45}$$

where the phase, ϕ will be given by Equation (2.37) and $P(\phi, t)$ is the probability density function of the phase distribution after a time t.

In the case of restricted diffusion (i.e. diffusion in an inhomogeneous or anisotropic environment) the form of the diffusion propagator becomes more complex and closed form solutions of Equations (2.40) and (2.45) are only possible for certain simple geometries and assumptions.

Analytical solutions can be found for some simple restricting geometries such as spheres, cylinders and parallel plates [41–43]. For more complex environments, however, an analytical solution is intractable and we must rely on simulations to approximate the diffusion MRI signal.

2.2.1 The Bloch-Torrey equations

As well as describing the dMRI signal by considering the microscopic diffusion of spins, a macroscopic formulation can be derived considering Fick's laws of diffusion. The combination of the Bloch equations (Equations (2.26) to (2.28)) with Fick's second law of diffusion leads to the Bloch-Torrey (BT) equations, proposed by H. C. Torrey in 1956 [44–46]:

$$\frac{\partial \mathbf{M}(\mathbf{r},t)}{\partial t} = \gamma \mathbf{M} \times \mathbf{B}(\mathbf{r},t) - \frac{M_x \hat{\mathbf{x}} + M_y \hat{\mathbf{y}}}{\mathrm{T}_2} - \frac{(M_z - M_0)\hat{\mathbf{z}}}{\mathrm{T}_1} + \nabla \cdot (\mathbf{D} \nabla \mathbf{M}) \ . \quad (2.46)$$

This version of the BT equation is sometimes referred to as the standard BT equation. An additional term can be added to account for the evolution of the magnetisation due to a flow, described by a velocity field $\mathbf{v}$ to get the generalised BT equation [47,48]. However, this flow term is often dropped, assuming no net flow, in the application to dMRI, leaving Equation (2.46). $\mathbf{D}$ is the diffusion tensor, a generalisation of the diffusion coefficient, D, to allow for anisotropic diffusion. In short, this means that diffusion happens at a different rate in different directions. It is the $\mathbf{D}$ term in Equation (2.46) which encodes the evolution of the magnetisation when there is diffusion.

In general, the BT equations cannot be solved analytically, apart from in some simple cases such as isotropic free diffusion. For instance, the solution to the BT equations in isotropic free diffusion can be shown to give the expected Stejskal-Tanner equation, Equation (2.43) [39].

For complex geometries, as with the case above, we must rely on computational methods to come to a solution for the BT equations.

2.3 Diffusion Simulation

Diffusion simulations attempt to evaluate Equation (2.45) or Equation (2.46) computationally. Simulation approaches broadly fall into two categories: numerical solutions of the BT equation and Monte-Carlo (MC) simulations of the diffusion dynamics. This section introduces these techniques, highlighting some of the similarities and differences between them.

At a high level, all diffusion simulations have three common components: the substrate, the diffusion dynamics and the measurement. The substrate describes the environment in which the diffusion is taking place. A common example of this is parallel cylinders representing axons in white matter. The diffusion dynamics describe our understanding of the processes underlying the diffusive motion of molecules and the measurement describes how this diffusive motion results in a synthetic dMRI signal.

2.3.1 Numerical solutions

Numerical solution approaches generally attempt to solve the Bloch-Torrey equation [44]. These approaches combine both the dynamics and the measurement components of the diffusion simulation by solving for the magnetisation in Equation (2.46). The third component, the substrate, defines boundary conditions required for the solution of the equation.

There are two typical methods for solving the partial differential equation (PDE) in Equation (2.46), finite difference methods (FDMs) and finite element methods (FEMs). Finite difference methods evaluate the PDE using a local Taylor expansion at discrete points which are generally uniformly separated in each spatial as well as the temporal dimension [49]. FDMs are an efficient method for solving PDEs when the problem can fit into a rectangular grid, however they are less effective when applied to complex geometries [49, 50].

Finite element methods subdivide the domain into small elements which are simple geometric shapes, though unlike the FDMs, they do not have to form a regular grid, but rather an arbitrary mesh. In each element, the PDE solution is approximated by simple functions such as a linear combination of polynomials. The combination of all of these local approximations can be solved to give a numerical solution of the PDE across the whole domain [51]. FEMs are generally more complex to formulate and implement than FDMs, however the added complexity can be worth the effort for more difficult problems in which FDMs may be ineffective [52].

2.3.2 Monte-Carlo simulations

Monte-Carlo (MC) methods take a different approach to the simulation of the dMRI signal. MC methods simulate the Brownian motion of a large number of particles by simulating the motion of each particle individually, along with the MR acquisition to generate the dMRI signal.

There exist many different implementations of the MC simulation of dMRI [24, 25, 53–57], however the underlying principles are similar for all of them. The following is a general description of the MC simulation process, however the specifics for each different implementation may vary.

Most early MC studies used simple, easily parameterised substrates like regularly packed cuboids [57] or cylinders [24]. As computational power has increased, so too has the capacity for more and more complex substrates. This includes cylinders with randomly distributed radii [25], undulating cylinders [54], beading cylinders [58] and meshes, both generated from high resolution microscopy of tissue [59] and computer generated cell models [60, 61].

Algorithm 2.1 Basic algorithm for taking a step in the random walk.

```
generate randomly oriented step vector
check if step crosses a barrier
while step crosses barrier do
    amend step according to barrier interaction (e.g. elastic reflection)
    repeat barrier checking on amended step
update the particle position
```

The Brownian motion of particles is typically simulated as a random walk of many independent particles. The time domain is discretised into many time points and at each time point each particle takes a random step through the substrate. One step of the random walk can be briefly summarised as shown in Algorithm 2.1.

Following the Brownian motion, each particle in the simulation will have taken many steps giving each particle a unique trajectory that it has traversed. The incremental phase, $\Delta\phi$, accrued at each step can be calculated from a discrete version of Equation (2.37). Under the assumption of uniform B_0, only the gradient term matters, giving

$$\Delta\phi = \gamma \mathbf{g}(\mathbf{R}(t), t) \cdot \mathbf{R}(t)\, \Delta t\,, \tag{2.47}$$

where Δt is the duration of the step and $\mathbf{g}(\mathbf{R}(t), t)$ and $\mathbf{R}(t)$ are the gradient and particle position during that step respectively.

The phase accumulation in Equation (2.47) for each spin in the simulation can be combined with Equation (2.45) to approximate the total signal for the dMRI acquisition as

$$S = \sum_j e^{i\Phi_j}\,, \tag{2.48}$$

where Φ_j is the total phase accrued for each spin.

Monte-Carlo simulations are a powerful tool for dMRI simulation due to their ability to handle any arbitrary substrate and MR pulse sequence. Additionally, MC simulations can be modified to account for effects that are more difficult to formulate for analytical and numerical solutions of the diffusion equation such as semi-permeable membranes, membrane-particle interactions and spatially and/or temporally varying $\mathrm{T}_1, \mathrm{T}_2$ and diffusivities.

A drawback of MC simulation, particularly for complex substrates is the need to simulate enough spins to mimic the ensemble behaviour of spins *in vivo* as well as enough discrete time points to adequately capture the dynamics through the pulse sequence. The huge number of calculations required to handle large simulations can be alleviated by exploiting the inherent parallel nature of the problem to run simulations in parallel on a central processing unit (CPU) cluster or, even more effectively, a GPU cluster.

Details on some dMRI simulation packages and numerical phantoms are presented in Chapter 3.

2.4 Diffusion Modelling

As discussed in Section 2.2, the environment in which spins diffuse can impact on the diffusion process, imparting information about this environment in the diffusion weighted MRI signal via the diffusion propagator. As dMRI research has advanced, many works have tried to take advantage of this to estimate information about the cellular environment spins are diffusing to gain biological meaningful information [7, 62, 63].

Broadly speaking, these models fall into two families: signal representations and biophysical models. Signal representations seek to express the signal from a voxel using a general expression which captures salient features of the data without making too many assumptions and by construction don't carry any particular biological meaning [62]. Some examples of representations are diffusion tensor imaging (DTI) [64], diffusion kurtosis imaging (DKI) [65] and diffusion spectrum imaging (DSI) [66]. Biophysical models, however, model tissue by breaking it down into compartments that can be described by simple analytical models such as spheres and cylinders. Some examples of this kind of model are ball-and-stick [67], composited hindered and restricted models of diffusion (CHARMED) [68] and neurite orientation dispersion and density imaging (NODDI) [13].

As different models and representations have been proposed over the past two decades, a menagerie of techniques has emerged that is too deep and varied to summarise here, however a number of articles have attempted to review the field [7, 62, 63]. The remainder of this section will briefly cover one family of models that are used in this thesis.

2.4.1 Spherical deconvolution techniques

A popular approach in dMRI modelling of brain tissue, in particular WM, is to model the signal as a combination of two components: one describing the orientational distribution of the tissue (e.g. direction of axons) and one describing the dMRI response of typical tissue (e.g. the typical signal from an axon), sometimes called the kernel. One such family of techniques, which aims to estimate the orientational distribution of fibres in WM, models this relationship as a spherical convolution of the fibre orientation distribution function (fODF) with a kernel that is the typical dMRI response of a single fibre (or bundle of fibres). This kernel is commonly referred to as the fibre response function (FRF). Mathematically, this is expressed as

$$S(\theta, \phi) = F(\theta, \phi) \otimes R(\theta)\,, \tag{2.49}$$

where S is the dMRI signal, F the fODF, R the FRF and θ and ϕ are the elevation and azimuthal angles respectively. Note here that the FRF, $R(\theta)$, is assumed to be axially symmetric.

Since the goal of these techniques is typically to estimate the fODF, the approach is to estimate the FRF *a priori* and deconvolve this out to leave the fODF. While there are many techniques based on this SD each with slightly different approaches, many such techniques express the signal, fODF and FRF in their spherical harmonic (SH) expansions to facilitate SD.

2.4.1.1 Spherical harmonics

The spherical harmonic series is a spherical analogue to the Fourier series in Cartesian space, providing an orthonormal basis of functions on the sphere which can be used to represent any spherical function. The basis functions are denoted $Y_l^m(\theta, \phi)$ where l is called the degree and m the order of the SH. Any (potentially complex valued) function, f, on the sphere can then be written as a linear combination of SHs:

$$f(\theta, \phi) = \sum_{l=0}^{\infty} \sum_{m=-l}^{l} c_l^m Y_l^m(\theta, \phi)\,, \tag{2.50}$$

where the SH coefficients, c_l^m, are given by

$$c_l^m = \int_0^{2\pi} \int_0^{\pi} f(\theta, \phi) Y_l^{m*}(\theta, \phi) sin(\theta) d\theta d\phi\,, \tag{2.51}$$

where $*$ denotes complex conjugation.

The benefit of working in SHs is that the spherical convolution in Equation (2.49) simplifies to a set of matrix multiplications in spherical harmonics:

$$\vec{S}_l = R_l \vec{F}_l\,, \tag{2.52}$$

where $\vec{S}_l$ is the $2l + 1$ vector of l^{th} degree SH coefficients of the signal, $\vec{F}_l$ is similarly the l^{th} degree SH decomposition of the fODF and R_l is the $(2l+1)(2l+1)$ matrix representing the l^{th} degree rotational harmonic decomposition of the FRF. R_l must be represented by rotational harmonics (an analogue of the Fourier series in $SO(3)$ rather than S^2 as with SH) for this to work mathematically [69,70] however much of the complexity is simplified due to the physical properties of the dMRI signal.

Firstly, the assumption that the FRF is axially symmetric means that the R_l collapses to a single scalar value for each l [70], making the solution of Equation (2.52) for $\vec{F}_l$ a simple case of the division of $\vec{S}_l$ by a scalar. Secondly, since the dMRI signal is antipodally symmetric, all odd m terms are zero and since it's real valued, the SH coefficients exhibit conjugate symmetry, that is $c_l^m = (-1)^m c_l^{-m*}$ [70, 71].

All of this means that SD based techniques are efficient methods to estimate the fODF from dMRI and as such have been used in many studies to investigate the organisation and structural connectivity of the brain [70, 72–79]. These techniques will, of course, depend on the FRF which is estimate prior to performing the spherical deconvolution, something which we investigate in Chapter 8.

3 Literature Review

Chapter Contents

3.1 Numerical Solutions 43

3.2 Monte-Carlo - Packages 45

3.3 Monte-Carlo - Numerical Phantoms 45

3.4 Conclusions 49

Chapter Summary

This chapter presents some examples of the use of numerical phantoms from literature. It represents a review of the contemporary literature in dMRI simulation, highlighting the state of the art and presenting some weaknesses which this project aims to address.

3.1 Numerical Solutions

Many of the studies presenting a numerical solution are focused on the validation of the technique and improvements to various algorithms rather than the direct use of the technique involving numerical phantoms.

One application of FDM solutions of the Bloch-Torrey equations is in the simulation of dMRI signals from histological images. For example, Chin et al. [80] simulate the signal from segmented histological images of mouse spinal cord white matter, showing that the fast and slow components of a bi-exponential decay of diffusion attenuation do not arise from a contribution from each of the intra- and extracellular components. Hwang et al. [81] extend this technique to 3D, showing that their FDM solutions agree well with analytical solutions for hexagonally packed cylinders.

Xu et al. [82] develop a matrix based FDM, also testing their solution on hexagonally packed cylinders, showing a reduction in error compared to a conventional FDM. This FDM is used by the same group to investigate the sensitivity of dMRI to intracellular structure [83]. A numerical phantom is used with cells represented as densely packed spheres with spherical nuclei at the centre of each sphere. The FDM is used to show that an oscillating gradient spin echo (OGSE) sequence is more sensitive to changes in the nucleus size than a PGSE sequence. Similarly, Xu et al. [12] use histology based FDM simulations to investigate the efficacy of an axon diameter estimation technique based on an OGSE sequence. They show that the OGSE technique is able to distinguish axons of a lower diameter than traditional PGSE techniques.

The first example of an FEM solution known to the author is presented by Hagslätt et al. [50]. In this study, rather than solving the BT equation, the FEM is used to solve for the diffusion propagator [39, 45]. From the diffusion propagator, the diffusion attenuated signal is calculated based on an assumption of infinitely narrow gradient pulses. A good agreement is shown between simulation and theoretical solutions for a range of simple geometries (parallel plates, a lamellar system and hexagonally packed cylinders).

More studies have recently begun investigating diffusion simulation using FEMs. Moroney et al. [84] present an FEM solution of the BT equation without the relaxation and flow terms for numerical analysis of dMRI experiments in the short gradient pulse limit. FEM results are compared to analytical solutions and MC simulations in simple geometries, showing that the FEM is more accurate than MC simulations, whilst taking less time to run.

Nguyen et al. [85] also present an FEM solution of the standard BT equation, showing its application to diffusion simulation with more general gradient waveforms. The FEM solution is shown to be more accurate in some simple geometries than a finite volume method, with second order accuracy in both the spatial and temporal domains. Three example applications to questions in dMRI are demonstrated using this FEM [85]. One shows that an infinitely thin membrane can be used to approximate a thick membrane. The second shows that the apparent diffusion coefficient (ADC) approaches the value predicted by mathematical homogenisation for long diffusion times. Finally, a model of a neuron is presented as a spherical body, with cylindrical axons and dendrites protruding. The ADC is shown to approach a steady state faster with a smaller neuronal body.

Beltrachini et al. [48] present a solution of the generalised BT equation, extending the FEM of Nguyen et al. [85] to include the relaxation and flow terms. This FEM improves on some of the restrictions in the FEM, making the simulations more stable through the use of an implicit scheme that is stable for a coarser discretisation without compromising the validity of the result.

3.2 Monte-Carlo - Packages

Historically, most studies utilising MC simulation used in-house developed MC simulation software [57, 86–88], however in more recent years and as the complexity of situations possible to simulate has grown, a range of MC simulation packages have been released for public use.

Hall and Alexander [25] introduced MC simulation as part of the Camino diffusion MRI toolkit [89] in the context of simulating swelling cylinders as a model of the effect of ischaemic stroke, however the MC framework is very general and can be used to simulate any arbitrary substrate from simple geometric constructs to complex 3D meshes.

Balls and Frank [56] present DiffSim, a dMRI simulation framework which embeds the MCell [90–92] cellular microphysiology simulator within an MRI simulator for synthesising the dMRI signal. DiffSim is used to simulate myelinated white matter [93], showing that an analytical solution model by Sen and Basser [94] holds for an SGP approximation or long diffusion time, however with more realistic pulse sequence parameters, the numerical simulations show lower anisotropy than the analytical model.

Landman et al. [55] developed the DW-MRI Random Walk Simulator (RWS) showing, as an example of its flexibility and reproducibility, a range of geometrical models for white matter damage, including healthy straight cylinders, bulging cylinders, crimped cylinders and broken cylinders.

Yeh et al. [53] present Diffusion Microscopist Simulator (DMS), showing a range of diffusion substrates ranging from simple parallel uniform cylinders to more complex undulating, beading or crossing arrangements of fibres, albeit at relatively low fibre densities. As of the writing of this review, the DMS software package has not been publicly released.

3.3 Monte-Carlo - Numerical Phantoms

The above packages, as well as dMRI simulation software developed in-house in various research groups, have been used to investigate the diffusion signal in many different numerical phantoms.

A common target of microstructure imaging is the estimation of axonal diameter and density. As mentioned above, FDM approaches have been used to investigate this [12,80], whilst this has been the subject of MC simulation studies as well. Alexander et al. [26] use Camino to simulate a series of numerical phantoms of parallel cylinders with radii drawn from a Gamma distribution for the validation of a technique for orientationally invariant indices of axon diameter and density.

Nilsson et al. [8], investigated the theoretical resolution limit for cylinder diameter estimation using diffusion MRI. Analytic expressions based on the Gaussian phase distribution approximation [39] were used for the intracellular signal and validated with MC simulations to determine a d_{min}, the diameter below which a cylinder cannot be differentiated from a cylinder with diameter approaching zero. The resolution limit for clinical scanners was found to be between 4 - 8 μm. This suggests a limitation on the level of microstructural detail that can be estimated using current clinical MRI machines.

Another problem commonly investigated using numerical phantoms is that of exchange between the intra and extracellular compartments of tissues. Permeability is difficult to control and vary in physical or biological phantoms, so numerical phantoms offer a unique tool with which to explore permeability and exchange models.

Nilsson et al. [14, 23] and Fieremans et al. [22] investigate the Kärger model [95], a model for exchange between two signal bearing compartments. These three studies all use similar numerical phantoms made of straight cylinders in which there is some probability that on encountering a barrier, the spin will pass through the barrier, exchanging spins between the compartments.

In their first study, Nilsson et al. [23] use simulations and experimental data to draw the conclusion that it is necessary to include exchange in a model containing two compartments, one of which is restricted. Fieremans et al. [22] show that the Kärger model is able to describe the signal for long diffusion times and sufficiently impermeable membranes, however at larger permeabilities, the Kärger model underestimates the value of the permeability. Nilsson et al. also investigate the effectiveness of the Kärger model at estimating the intracellular water fraction, showing that the Kärger model has a negative bias, underestimating the intracellular water fraction by up to 25% when there is high permeability when compared to a computational model made by building a database of simulated signals [14].

Nilsson et al. [54] also investigate the importance of axonal undulation on diffusion MRI measurements. In this experiment, numerical phantoms consisting of axons with either sinusoidal or helical undulations were used in MC simulations to investigate the impact on a range of dMRI measured parameters. Nilsson et al. show that undulation affects essentially all of the parameters they tested derived from dMRI, for instance, undulation results in an overestimation of axonal diameter when using models that assume axons are straight [54].

Budde et al. [58] use MC simulations to investigate the effect of neurite beading, showing that beading is sufficient to explain the decrease in ADC after ischaemic stroke. Numerical phantoms consisting of straight cylinders with increasing amounts of beading introduced are simulated, showing a decrease in ADC in both the intra and extracellular spaces with increased beading.

Lin et al. [96] investigate the effect of traumatic brain injury (TBI) on dMRI derived parameters. Using a numerical phantom consisting of straight cylinders representing axons and ellipsoids representing glial cells, the effects of TBI are investigated by varying various parameters such as the size of the glial cells, the permeability of the cylinders and the spacing of the cylinders [96]. Using this technique, Lin et al. conclude that the inconsistencies amongst previous dMRI based TBI studies [97–99] are due to differences in the timing between the onset of TBI and the diffusion measurement, arguing that different processes drive the TBI at different timings, leading to different dMRI characteristics.

Lin et al. [100] similarly investigate the effect of myelin water exchange on various dMRI derived parameters. In this work, their representation of WM is slightly different, choosing nested cylinders to represent the intra-axonal and myelin-water compartments. Using MC simulations on these cylinders, they show correlations between dMRI derived parameters such as ADC and fractional anisotropy (FA) and the echo time.

Lam et al. [101] produce an empirical model of the extra axonal space using a series of MC simulations based on both regularly and randomly packed cylinders. The model is based on the diffusion spectrum [102], modelling diffusion in densely packed cylinders as diffusion in a series of pores with a small chance of exchange between the pores. The empirical model agrees closely with MC simulated data.

Some studies combine analytical solutions of the Bloch-Torrey equation and MC simulations. Rensonnet et al. [103] use this combined simulation to synthesise signals for parallel and crossing cylinders. The intracellular component is modelled using an analytical solution for diffusion within a cylinder based on Grebenkov's multiple correlation function approach [104]. The extracellular compartment, which is much more complex geometrically, is simulated using Monte-Carlo simulations. This hybrid approach yields simulation results which are indistinguishable from pure MC simulation whilst being quicker and more precise than a purely MC approach [103].

In a further study, Rensonnet et al. [105] use this approach to assess the validity of the superposition approximation of crossing fascicles (i.e. that the total signal from crossing fascicles is the sum of the signal from each fascicle independently). They are able to show that the signal differences between the superposition approximation and a full simulation of interwoven fascicles is small enough compared to typical noise levels in clinical dMRI data, that the superposition approximation is sufficient to describe the signal. A drawback to this hybrid approach is that the intra and extracellular compartments are treated as distinct, non-interacting compartments, meaning that membrane permeability is not accounted for.

An emerging application for dMRI simulations in the direct computational modelling of microstructure. The first example found for this type of modelling is actually the 2010 work by Nilsson et al. [14] mentioned above for evaluating the Kärger model. In recent years however, this idea has reemerged, partially thanks to the popularity of machine learning and emergence of MR fingerprinting.

One area of application these approaches have found is in the estimation of axonal permeability. Nedjati-Gilani et al. [18] use a machine learning technique known as random forest regression to learn the relationship between dMRI signal and microstructural parameters. They simulate diffusion using Camino in a range parallel cylinder substrates with different microstructural parameters including membrane permeability to build a dataset to train the random forest regression. The random forest regression is shown to estimate membrane permeability well, performing better than the Kärger model and an application to MS is presented, showing *in vivo* results consistent with pathology.

Palombo et al. [17] verify this method using a cuprizone treated *in vivo* mouse model. Cuprizone is a well known mouse model of WM demyelination, which is important as demyelination is hypothesised to affect axonal permeability. The random forest approach achieves accurate and robust estimation of microstructural parameters which match expected microstructure changes from electron microscopy and gains more specific information that typical dMRI measures such as ADC and FA.

Hill et al. [16] extend this work to use a deep neural network (NN) in place of the random forest regression. They are able to show that the NN outperforms the random forest approach and an application in *in vivo* mice estimates microstructural parameters in the biologically plausible range.

Rensonnet et al. [15] also attempt to use computational models to directly estimate microstructural parameters using dMRI simulations. Using the hybrid approach mentioned above [103], they generate a dictionary of dMRI simulation signals for various combinations of microstructural parameters and crossing fascicles. Measured signals are then compared against this dictionary to find the entry that most closely matches the measured signal to get an estimate of the microstructural parameters. They show that their approach achieves accurate and robust estimates of microstructural parameters and shows good correspondence with histology compared to traditional closed-form models when applied to an *in vivo* mouse model of spinal cord injury.

3.3.1 Towards realistic WM numerical phantoms

One thing common to all of the dMRI simulation studies presented so far is that they represent WM axons using simplified representations. By far the most common representation of axons in as straight cylinders in parallel bundles or simple crossing arrangements [14, 15, 25, 26, 103].

True axonal morphology, however, is much more complex than this. *ex vivo* EM and high-resolution x-ray imaging studies have shown that real axons have complex morphologies with undulation, beading (variable diameter along the axon) and non-circular cross sections [27, 28, 106, 107]. Recently, a few groups have tried to tackle the challenge of generating WM phantoms with more realistic microstructure, capturing some of these features.

Rafael-Patino et al. [61] extend the phantom generator of Close et al. [108], previously used to simulate dMRI using a simple diffusion tensor model, to generate crossing, interdigitated fibre bundles. They use these models to show that previous simple representations may not be sufficient to accurately represent realistic dMRI signals [109].

Ginsberger et al. presented an extension of DMS which shows more complex white matter numerical phantom including orientation dispersion, tortuosity, beading and nodes of Ranvier [110], however the range of orientation dispersion and axon densities achieved does not reach typical *in vivo* values. They subsequently improved their phantom generator, producing "microstructure environment designer with unified sphere atoms (MEDUSA)", which works by generating a set of overlapping fibres decomposed into small segments and iteratively refining their positions to remove the overlap between them [111, 112]. MEDUSA produces phantoms with complex fibre configurations at high fibre density, but as of writing, no dMRI application has been presented.

3.4 Conclusions

There has been significant work into the use of numerical phantoms in the simulation of diffusion MRI signals. As discussed, one thing which is common throughout most of the works is a simple model of WM. Typically, WM is represented as a set of densely packed, straight, parallel cylinders, which oversimplifies the underlying complexity of the microstructure which has been revealed through EM [27, 28].

Generating WM substrates which accurately represent real microstructure is important for the kinds of model validation studies mentioned in Section 3.3 because if a model is validated using overly simplified simulations, it may not be very robust when applied to real data, something with Rafael-Patino et al. have studied [109].

Realistic WM representations are additionally important for the kind of computational models mentioned in Section 3.3 for similar reasons. Another element which is particularly important for building the training datasets or dictionaries for these kinds of computational models is controllability. The WM numerical phantoms need to be realistic, but also must be able to be generated in a controlled manner so that we can generate substrates covering the relevant ranges of microstructural parameters needed to build robust computational models.

These are the main driving reasons behind the work presented in this thesis which aims to build these realistic WM substrates with controllable morphology.

Part II

Methods Development

4 Fibre Growth for WM numerical phantom generation

Chapter Contents

4.1 Introduction 53

4.2 A preliminary fibre growth algorithm 55

4.3 Experiments and Results 60

4.4 Discussion and Conclusion 64

Chapter Summary

This chapter introduces preliminary fibre growth algorithm (preFiG), a method for generating white matter (WM) numerical phantoms with realistic orientation dispersion and packing density. The growth algorithm is introduced, describing how each WM fibre is grown individually to generate a densely packed substrate.

Some experiments are presented, one of which explores how different inputs to preFiG impacts the shape of resulting fibres. An application of the method to dMRI is demonstrated with simulations of the diffusion-weighted MR signal in three example substrates with differing orientation dispersion, packing density and permeability.

4.1 Introduction

Numerical phantoms have found much use for validating many magnetic resonance imaging (MRI) experiments. In particular, numerical phantoms are often used when developing diffusion MRI (dMRI) microstructure imaging techniques where simulations of the dMRI signal in phantoms with known microstructural properties are used in lieu of an *in vivo* ground truth measure of microstructure [63].

While recently numerical phantoms have proven useful for validating microstructure imaging of grey matter [5], they have more commonly been used for validating white matter (WM) microstructure, with many studies comparing parameter estimates from fitting their models to the known ground truth from the phantoms e.g. [6–14]. Some recent works directly estimate microstructural features using fingerprinting techniques and machine learning to match simulated signals and the corresponding ground truth microstructure of the numerical phantom to the measured signal [15–18]. As well as affecting the dMRI signal, microstructural features also influence other MR techniques such as susceptibility-weighted imaging [19, 20]. For instance, Xu et al. [21] recently used simulations to show that using realistic axonal models rather than simple circular cylinders affects the MR signal. Therefore, it is important to the MRI community to generate realistic WM numerical phantoms which accurately capture microstructural features in order to get realistic simulated signal.

Generating realistic WM numerical phantoms which accurately capture realistic microstructural features (such as dispersion, undulation, beading, etc.) at high packing densities is a major open challenge for the dMRI community. While densely packing straight, parallel, fibres is relatively easy, only a few groups have attempted to densely pack irregular, non-parallel, fibres.

The most common approach to this is the packing of fibres into densely packed configurations [61, 108, 110, 111]. The typical approach, as taken in the state-of-the-art MEDUSA algorithm [111], is to generate a set of overlapping fibres decomposed into small segments and iteratively refine their positions to remove the overlap between them. Despite their recent progress, further advance of this class of techniques may be limited, because nature does not create fibres before attempting to pack them together. Instead, real axons are guided by chemical cues and fit into available space as they grow [113,114]. Mimicking the natural fibre genesis may prove important for building more realistic phantoms.

Here, we propose a completely different strategy: rather than densely ‘packing’ irregular fibres, we ‘grow’ fibres contextually, mimicking natural fibre genesis. We propose a preliminary fibre growth algorithm (preFiG) for the generation of WM numerical phantoms with more realistic orientation dispersion and packing density. Fibres are grown one-by-one following a cost function which attempts to impose the morphological priors that are input to the algorithm.

The rest of the chapter is organised as follows: Section 4.2 describes preFiG, Section 4.3 details some experiments showing the potential of the algorithm and comparing it to a brute-force approach to fibre growth and Section 4.4 summarises the contributions and discusses limitations in the preliminary implementation and future work to improve the method.

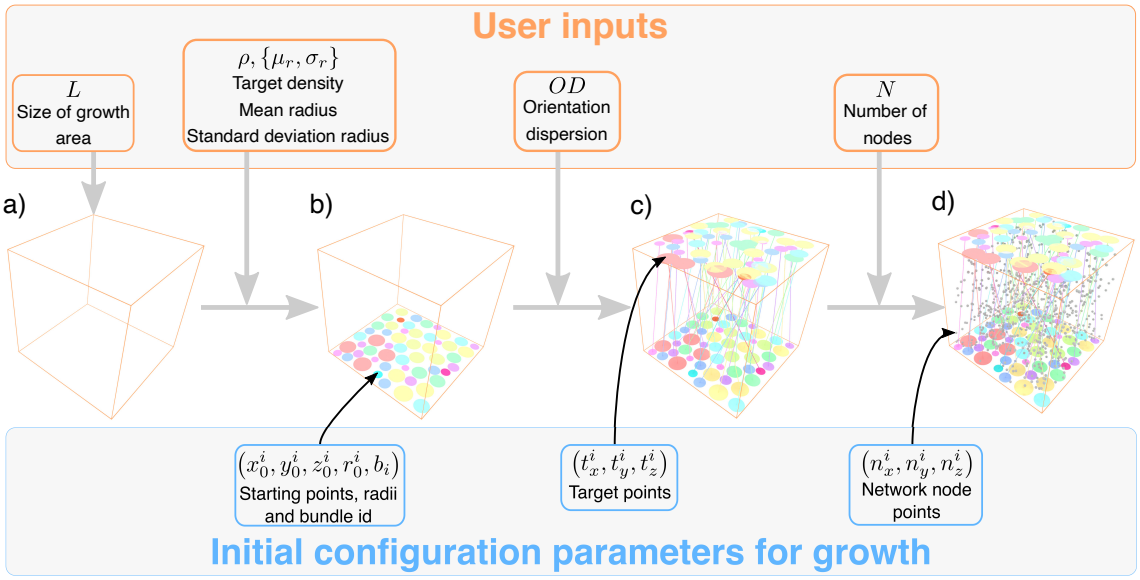

Figure 4.1: Inputs to the preFiG algorithm for the single bundle case. L defines the size of the area of that the growth will take place in. The target density and fibre radius distribution govern the generation of starting points for each fibre by packing in 2D. Orientation dispersion parameters govern the generation of target points corresponding to each starting point. N defines the number of nodes to use when generating the network. In the case of multiple bundles, starting and target points are generated for each bundle and then combined into the same space which is filled with nodes for the network.

4.2 A preliminary fibre growth algorithm

In this section we describe a preliminary fibre growth algorithm (preFiG) which grows fibres one-by-one, avoiding intersection between fibres whilst attempting to ensure that the resulting substrate has desired morphological input properties such as orientation dispersion, diameter distribution and packing density.

The generation of preFiG numerical phantoms happens in three main steps:

STEP 1 Generate initial growth configuration from user inputs

STEP 2 Grow the fibres using preFiG growth algorithm

STEP 3 Generate 3D meshes for dMRI simulation

The remainder of this section outlines these steps in further detail, breaking each step down into smaller steps and detailing each one.

4.2.1 STEP 1: Initial growth configuration

STEP 1 is broken down in to three substeps which are outlined in Figure 4.1:

STEP 1.1 Generate fibre starting points for each fibre to grow from (Figure 4.1a-b). To generate these starting points preFiG packs circles with the desired diameter distribution up to the target density (defined in terms of the desired fibre volume fraction) in 2D, following the approach taken in [25].

STEP 1.2 Generate fibre target points for each point to grow towards (Figure 4.1c). To encode the desired orientation distribution, each fibre has a direction drawn from the target distribution which gives a target point for the fibre to grow towards. As a demonstration of the framework, in this work we use the Watson distribution [115].

STEP 1.3 Generate growth nodes (Figure 4.1d). preFiG uses a set of pseudorandomly placed points (which we refer to as nodes) to sample the space and encode which regions are occupied by existing fibres. This simplifies collision checking making growth more efficient than a direct collision detection approach involving growing each fibre one small step at a time and checking collisions with existing fibres. This is demonstrated in Section 4.3.3.

4.2.2 STEP 2: Fibre growth

STEP 2, the main growth algorithm, is broken down into a series of substeps which are outlined in Figure 4.2:

STEP 2.1 Create growth network (Figure 4.2a&b). In order to encode which nodes a fibre can move to from any other node, the growth nodes are connected using the Delaunay triangulation.

STEP 2.2 Grow one fibre step (Figure 4.2c-e). Fibres grow one-by-one in a random order along this network towards their target points while avoiding existing fibres.

During growth, a fibre must choose in which direction it should grow. This direction is chosen in preFiG by following a cost function which encourages fibres to grow towards their target points (Figure 4.2d). From a starting node, s, the candidate nodes, c, that the fibre can move to are any nodes that share an edge with s. In addition to its position, each network node stores the maximum fibre diameter, d_c, that can be sustained at that node without intersecting another fibre (described further in **STEP 2.3**). The fibre will move to a candidate node according to a cost function consisting of two terms; l_t, which penalises taking very large steps or moving away from the target point, t, and l_d, which penalises moving to a position where d_c is low meaning that the fibre will have to shrink. The cost function for a fibre at a position, s, to move to a candidate node, c, given a target point, t, is

$$l = l_t + f l_d \,, \tag{4.1}$$

where

$$l_t = \frac{1}{2} \cdot \frac{\|s-c\|}{1+\|s-c\|} \cdot \left(1 - \frac{((c-s)\cdot(t-s))}{\|c-s\|\|t-s\|}\right) , \tag{4.2}$$

$$l_d = \max\left(0, \frac{1}{d_0}\left(d_0 - d_c\right)\right) . \tag{4.3}$$

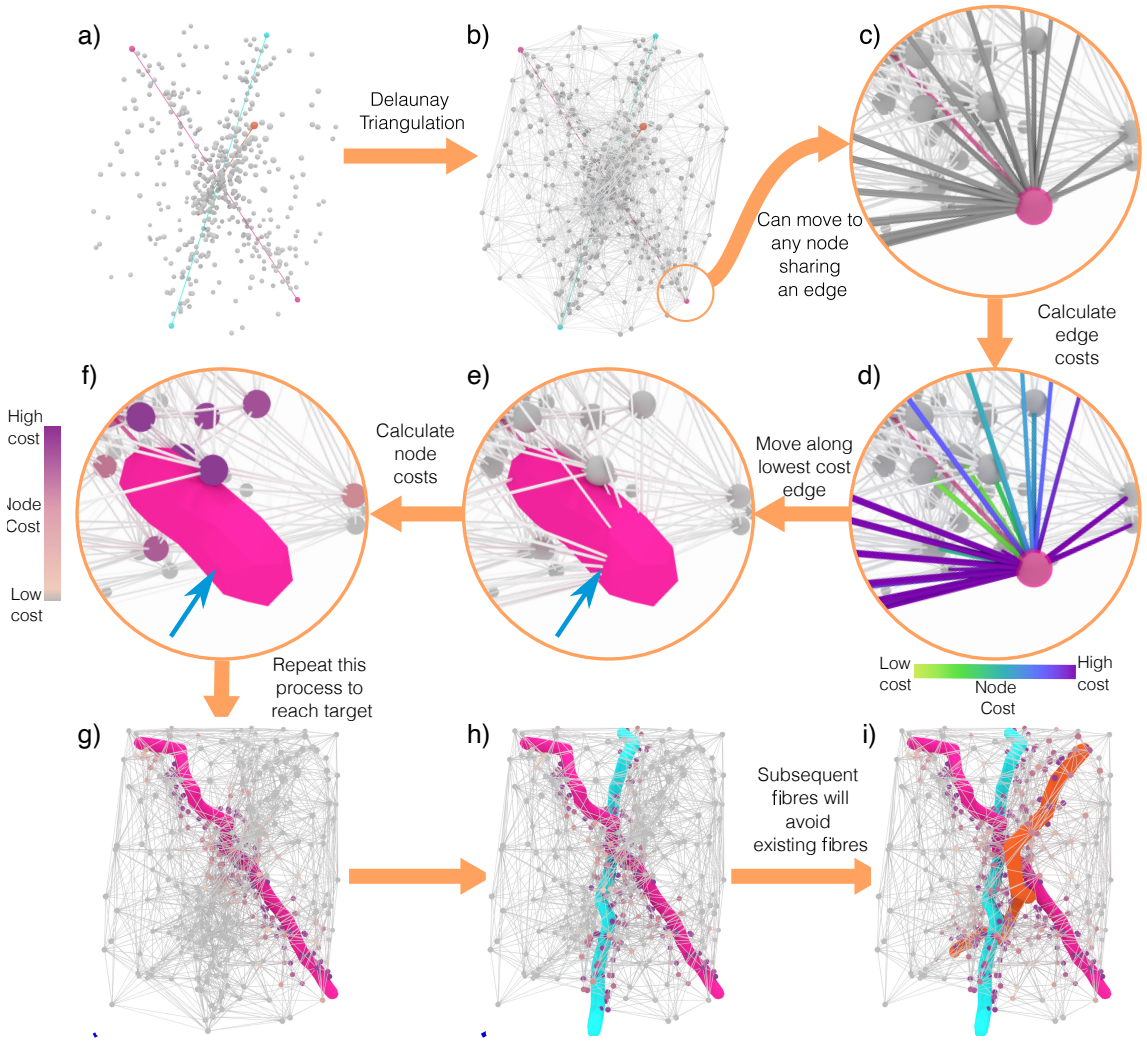

Figure 4.2: Overview of the basic growth algorithm in preFiG. In this example, three fibres are shown with a growth network that only contains relevant nodes for the sake of visualisation. From the set of nodes, a network is constructed using the Delaunay triangulation. Each fibre then grows from node to node, along any edge connected to the current node. The node moved to will be the node with the lowest cost. Once a fibre segment has grown, the network nodes are updated to store information about which nodes are occupied or near to an existing fibres. This contributes to the cost function for any future fibres, penalising moving to nodes too close to existing fibres. It is not possible to move to any node now inside a fibre as indicated by the removal of this edges from the network (pairs of blue arrows show where this is happening). The next fibres grow, now avoiding existing fibres until all fibres have finished. See Supplementary Video 1 for an animation of this algorithm.

Here, d_0 is the target diameter of the fibre and f is a weighting factor between the two terms. In this work, f is fixed to 0.2 to more strongly weight growth towards the target.

The next node for a fibre will be the candidate node which has the lowest cost according to Equation (4.1). This method of finding a path through the triangulation by choosing the lowest cost node at each position amounts to a greedy best-first

pathfinding approach with a heuristic given by Equation (4.1).

STEP 2.3 Update the network (Figure 4.2f). The growth network is updated in order to store the information about the space this fibre is occupying so that future fibres can avoid it. The way that this is done is to simply store the minimum distance from each node in the network to any existing fibre.

With the next node chosen, the value of d_c needs to be updated for other nearby nodes. All nodes have d_c set to the Euclidean distance between the node and the surface of the new section of fibre if that distance is less than the current value of d_c. This is illustrated in Figure 4.2d.

Any nodes which now lie within the fibre have d_c set to zero. Nodes with $d_c = 0$ are disallowed from future steps, meaning that once a fibre has grown, no future fibres can connect to any nodes within the fibre. This, in addition to shrinking the radius of future fibres according to d_c at each node means that the fibres grow in an almost completely non-intersecting manner. Since the value of d_c is set based on fibre-to-point distances, there can be cases in which the fibres would intersect when the closest point between two fibre sections is not at one of the fibre nodes. In order to account for this, a meshing process was developed which can deform fibres around one another. This is described in Section 4.2.3.

STEP 2.4 Repeat **STEP 2.2** and **2.3** until fibre reaches target (Figure 4.2g). By default in preFiG, each fibre will grow completely before the next one starts, meaning that **STEP 2.3** only needs to be performed once the fibre has finished growing. If fibres are allowed to grow concurrently, **STEP 2.3** must be performed after each growth step.

STEP 2.5 Repeat **STEP 2.2-2.4** for remaining fibres (Figure 4.2h-i). As noted in Figure 4.2 (e-h), as the network is updated, more and more nodes become inaccessible making the network sparser. This means that some fibres may reach a point from which they cannot grow any further and will become stuck. Currently, these fibres are simply removed from the final phantom, meaning the final phantom may have a lower density than the target density.

4.2.3 STEP 3: Meshing

Finally, **STEP 3**, the meshing procedure is described in more detail below:

STEP 3 Generate 3D fibre meshes. After the growth process, each fibre will be represented by a series of connected 3D points and corresponding diameters at each point, stored in the Stockley-Wheal-Cole (SWC) format [116]. In order to simulate diffusion MRI signals, these fibre skeleta need to be turned into 3D meshes. preFiG uses a meshing procedure designed to eliminate overlap between fibres, using Blender (`https://blender.org`), building upon on the SWC mesher add-on (`https://github.com/mcellteam/swc_mesher`).

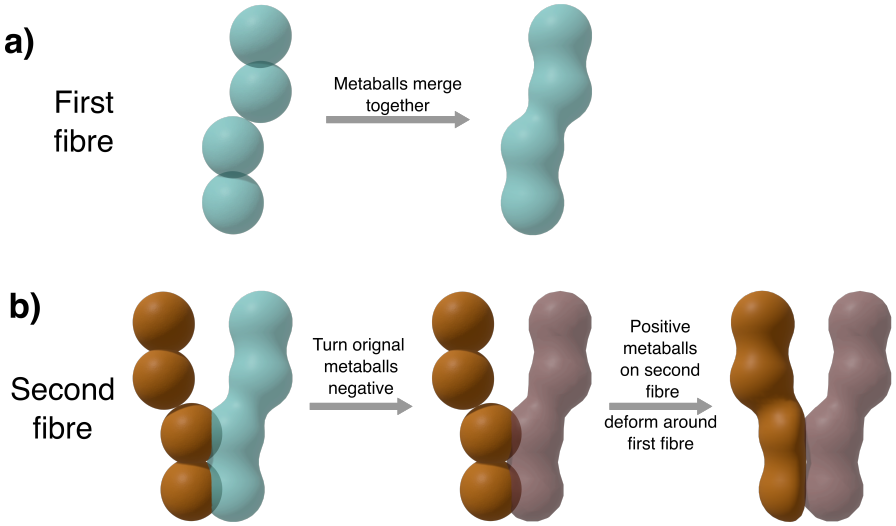

Figure 4.3: Demonstration of the meshing procedure in preFiG. The first fibre is created using metaballs to create a smooth surface. The second, and following fibres will be created using negative metaballs for any fibres that intersect in order to deform around them. Note that in practice, more spheres will be much more closely placed along the skeleton to create a smoother surface (see Figure 4.4

preFiG meshes are constructed using Blender metaballs, an implicit surface representation which is the isosurface of a function; typically a function analogous to the electric potential from a point charge. When two metaballs come close to one another, the fields combine and the surfaces will merge to form a smooth surface. By placing a series of metaballs along the skeleton of each fibre, a smooth surface is formed for each fibre one-by-one as shown in Figure 4.3a. Figure 4.3 shows an extreme example with few metaballs to demonstrate the meshing procedure, in practice however, more metaballs are placed closer together to form a smoother surface. Figure 4.4 demonstrates the surface formed in a realistic preFiG phantom and that the preFiG meshing procedure does not impact the diffusion dynamics compared to a perfectly straight cylinder.

When fibres are densely packed, the surfaces from neighbouring fibres may overlap. To account for this, a meshing procedure was developed in which fibres can deform around nearby fibres to avoid overlap. The metaball surface for one fibre is created as described above. This surface is then turned into a triangulated mesh, however the metaballs are retained. The metaball potential is then turned negative, meaning that rather than merging with any future nearby metaball surfaces, it will repel them, as shown in Figure 4.3b. This means that subsequent fibres which are meshed very close to, or overlapping with, existing fibres will deform organically to resolve the intersection, thus creating a series of completely non-intersecting fibre meshes which can be used by the dMRI simulator.

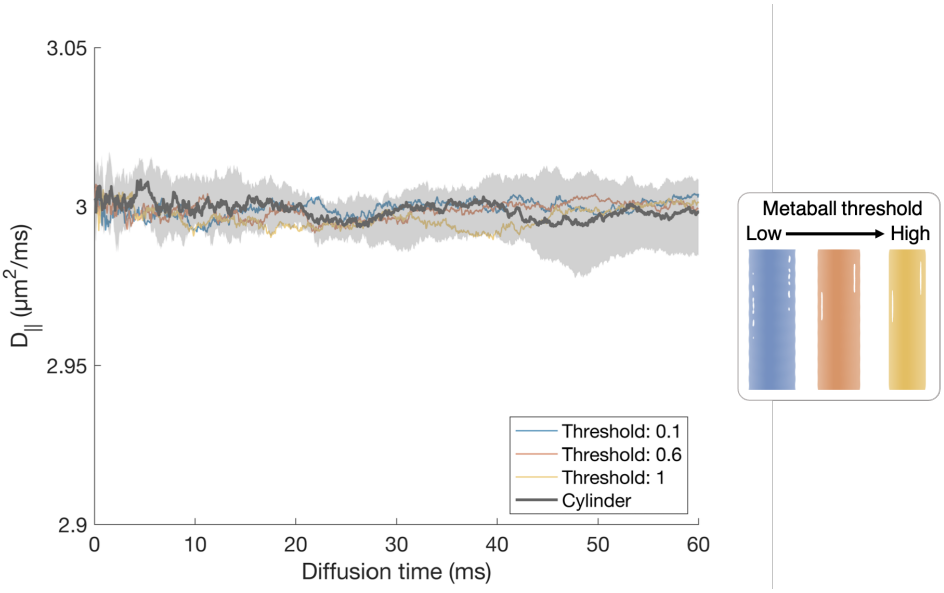

Figure 4.4: Axial diffusivity as a function of diffusion time for different metaball meshing thresholds. The threshold determines the level of the isosurface, a low threshold making the spheres behave more like hard spheres and a higher threshold smoothing nearby metaballs together. Lines show the median of 20 runs of the simulation with different seeds and the grey area shows the interquartile range of the cylinder diffusivity.

4.3 Experiments and Results

4.3.1 Effect of choice of growth network

As mentioned in Section 4.2.1, the choice of the node points in the network will affect the morphology of the resulting substrate. In order to investigate this, a qualitative experiment was performed in which a single fibre was grown on a network with either a) nodes on a uniform grid or b) pseudorandom nodes. In each case, the number of nodes was increased and the resulting fibre investigated.

The fibre was defined by a start point (20, 0, 0) μm, target point (0, 0, 50) μm and diameter 1 μm. This configuration was chosen so that the fibre would not have a path that directly followed one of the 90° or 45° lines in the uniform grid. Node points were initialised in either a uniform grid or pseudorandomly within the space [-5, -5, -5] to [25, 5, 55] to ensure coverage of the space in which the fibre would grow. The number of source points used was $N \approx 1000, 10000, 100000, 1000000, 5000000$.

The resulting fibres can be seen in Figure 4.5, where orange fibres are grown using the uniform grid and blue fibres using pseudorandom points. In both cases, as the number of nodes increases, the resulting fibre has more of a smooth, straight path between start and target. The uniform grid fibres, have a much more angular, structured path due to being forced to grow on the grid, while the pseudorandom fibres more irregular paths, which could be considered more ‘organic’ looking.

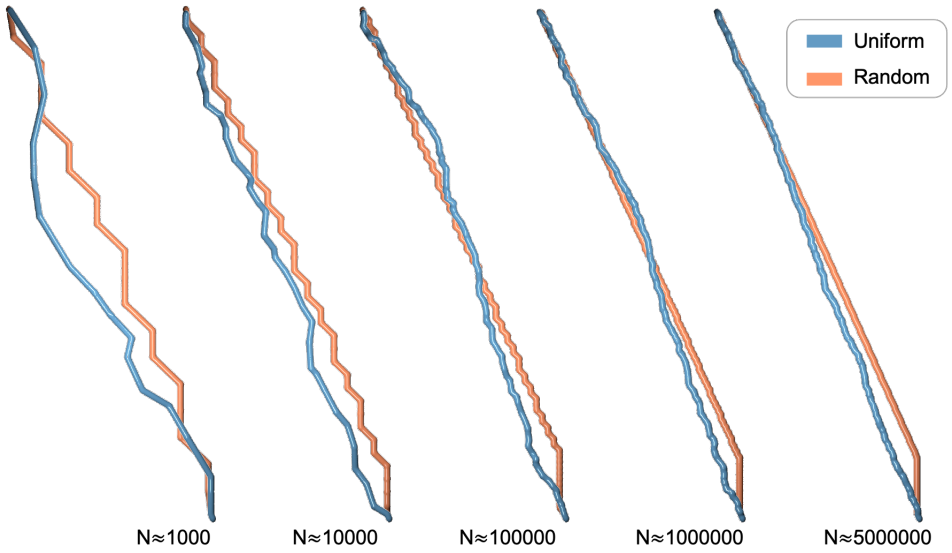

Figure 4.5: Fibres generated using uniform grid (orange) and pseudo-random (blue) network nodes for increasing numbers of nodes.

4.3.2 Demonstration of preFiG

To demonstrate the potential of preFiG, three substrates at different (dispersion, packing density) conditions were generated: (0°, 60%), (15°, 30%) and (35°, 25%), shown in Figure 4.6a. Each substrate is grown using 5×10^6 pseudo-randomly placed source nodes for the growth network, giving a network with 3.88×10^7 edges and a mean distance between any given node and its neighbours of 0.29 μm. The packing densities chosen represent the highest densities achievable using preFiG for each dispersion condition.

For the 0° dispersed substrate, initial diameters were drawn from a gamma distribution with mean $d_0 = 2$ μm and standard deviation $\sigma_d = 0.2$ μm. The 15° and 35° substrates were generated with $d_0 = 1.2$ μm and $\sigma_d = 0.2$ μm in order to show the flexibility of preFiG to generate substrates with different diameter distributions as well as orientation dispersion and packing density. Diameters were limited to be permitted to shrink to 25% of the original fibre diameter in order to fit into space.

For each substrate, the pulsed gradient spin echo (PGSE) signal was simulated in Camino [89] using 5×10^5 diffusing spins and 5×10^3 discrete time steps, uniformly distributed with bulk-diffusivity D_0=2 μm^2/ms. To show the range of simulation possibilities available, three different membrane permeabilities (κ=0, 0.0025, 0.0050 μm/ms) were also imposed. The simulated PGSE measurement parameters were: $\delta/\Delta = 1/40$ ms and 50 b-values from 0 to 9 ms/μm^2 along x-, y- and z-directions.

The corresponding direction-averaged simulated PGSE signals at different permeabilities are shown with SNR = ∞ in Figure 4.6b and SNR = 20 in Figure 4.6c. The signal decays to a lower value as the dispersion increases and density decreases, as expected.

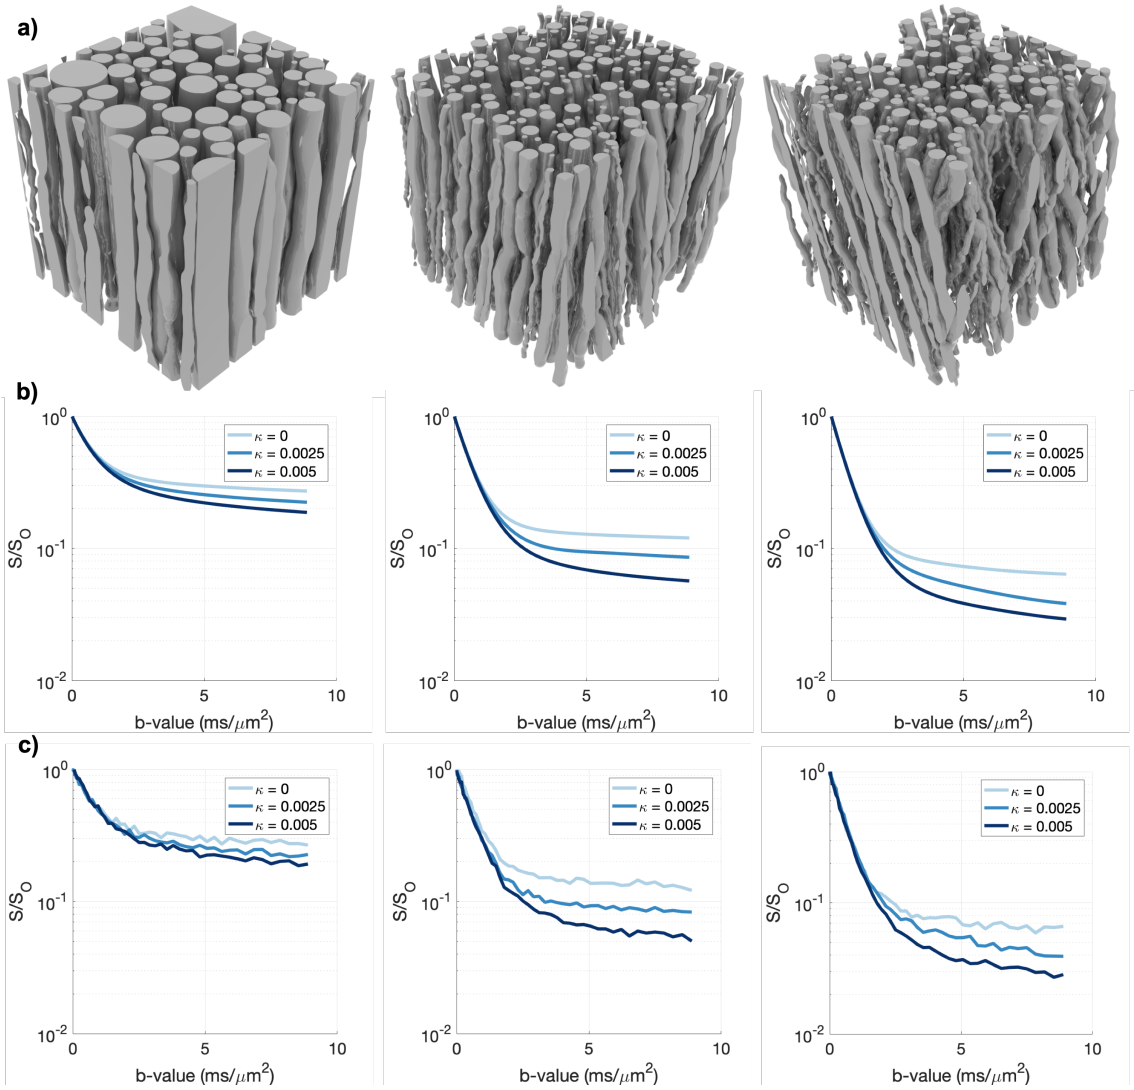

Figure 4.6: a) Example substrates (cut into 30x30x30 μm^3 cube) from the fibre growth algorithm, left to right: Zero macroscopic dispersion (60% density), 15° of macroscopic dispersion (30% density), 35° dispersed (25% density). b) Simulations for each substrate for varying permeabilities with SNR = ∞ and c) SNR = 20. Units of κ are μm/ms.

4.3.3 Comparison with brute-force approach

preFiG was compared against the naïve brute-force approach to fibre growth. The brute-force approach grows fibres one segment at a time and checks for collisions between the new segment and all existing fibres. Each new segment is chosen from one of 128 candidate directions on a cone aligned with the previous segment, with each direction being weighted according to Equation (4.2).

Substrates were grown with both the brute-force approach and preFiG using the same starting and target points and initial diameters. These initial parameters were determined by packing circles with gamma distributed radii (mean $d_0 = 2$ μm, standard deviation $\sigma = 0.6$ μm) into a 40 μm x 40 μm square up to a packing density of 60%. Target

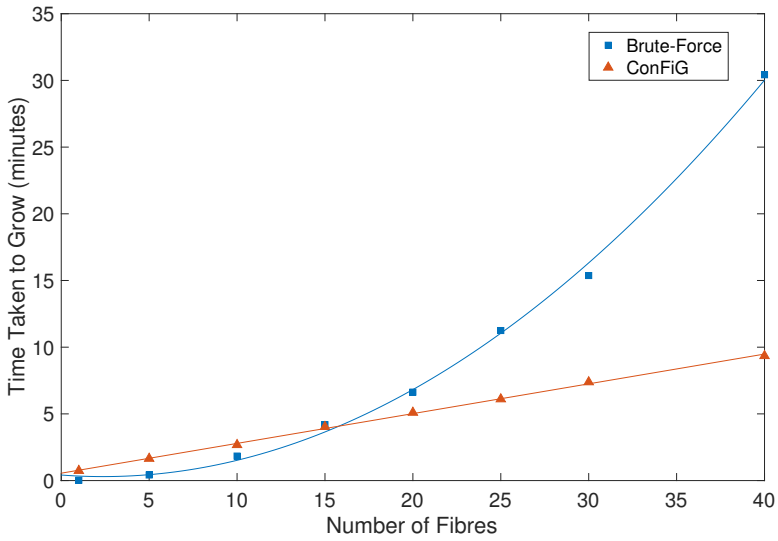

Figure 4.7: Timing of brute force growth vs. the fibre growth algorithm along with a quadratic fit (brute-force) and linear fit (fibre growth algorithm). The fibre growth algorithm is clearly linear in the number of fibres, while brute force growth fits an order n^2 well.

points were set as 40 μm directly above the starting points to define a substrate with 0° macroscopic orientation dispersion. This resulted in a substrate with a total of 54 initial fibres.

The fibre growth algorithm used 1×10^6 randomly distributed points for the Delaunay triangulation giving a mean distance between points of 0.5 μm, matching the brute force approach which used a segment length of 0.5 μm for each new fibre segment.

From these initial parameters, fibres were grown using a subset of $n = 1, 5, 10, 15, 20, 25, 30, 40$ fibres and the growth was timed. Each value of n was timed 5 times with and the mean taken to reduce single-run timing fluctuations.

Figure 4.7 shows the timing results of the brute-force approach versus the fibre growth algorithm. The fibre growth algorithm has approximately $\mathcal{O}(n)$ complexity with n being the number of fibres. Conversely, the brute-force algorithm shows $\mathcal{O}(n^2)$ complexity owing to the fact that every new segment has to check for collisions with all existing fibres.

The fibre growth algorithm has a higher $n = 0$ offset which is caused by the overhead in calculating the Delaunay triangulation for the growth network. This causes the brute-force approach to have better performance at low n, while at higher n (approaching the > 100 fibres needed for a realistic dMRI voxel) the linearity of the fibre growth algorithm gives it much faster performance.

4.4 Discussion and Conclusion

preFiG shifts the perspective from previous works attempting to pack together fibres by trying to mimic natural fibre genesis. This approach represents a major step towards very high fibre packing, enabling us to reach the highest dispersion at the highest packing density reached so far, to our knowledge. Our (15°, 30%) and (35°, 25%) represent an average ~50% and ~200% improvement, respectively, over the best previously reported results of (10°, 20%) [110][1].

The substrates presented in Figure 4.6 are just a few examples of the kinds of substrates that can be produced using our preFiG method, though by varying the setup of the morphological controls and start and target points, many different fibre configurations can be produced. Currently, fibres will attempt to grow in a straight line between the start and target points, meaning that certain configurations such as kissing bundles cannot be represented. However, the algorithm can in principle be extended to allow for series of target points, allowing the definition of a desired 'path' of a fibre.

Additionally, some input parameter settings cannot be achieved. For instance, trying to grow a substrate with both very high density and very high dispersion will result in a final substrate that does not reach the density required. The reason for this could be a combination of limitations of the algorithm in restricting growth to a discrete network and also the fact that some morphological settings are practically infeasible. This limitation, however, also applies to the fibre packing and brute force growth approaches.

One weakness of the fibre-growth algorithm is that since the fibre diameters are calculated from a fibre-to-point distance, there can still be some small amount of overlap between fibres. This is solved using the meshing process in Blender to deform the regions of slight overlap between neighbouring fibres.

Another limitation of the preFiG implementation is that while the phantoms produced improve upon the state-of-the-art, the fibre density achieved is still too low to really be applicable to white matter. Typical expected fibre density would be between 50 and 80%, so improvements need to made to enable phantoms to be generated at a higher fibre density.

To conclude, the proposed preFiG approach, using the fully connected growth network, is shown to be more efficient than a 'brute-force' growth approach. The fact that preFiG processing time is linear with the number of fibres makes it far more efficient for high numbers of fibres. For instance, a realistic voxel will need hundreds or thousands of fibres which will become impractically slow for the 'brute-force' approach, whilst remaining manageable for our algorithm. This efficiency, along with the high density and orientation dispersion achieved means that preFiG represents a promising step forward in the construction of ultra-realistic numerical phantoms of WM.

[1] The preFiG method presented here outperformed the state-of-the-art at the time it was published in 2019 [117]

5 ConFiG: Contextual Fibre Growth

Chapter Contents

5.1 Introduction 66
5.2 Methods 67
5.3 Experiments 74
5.4 Results 77
5.5 Discussion 83
5.6 Conclusion 86

Chapter Summary

This chapter introduces a series of mechanisms which were added to improve the initial implementation of preFiG described in Chapter 4 and the biological fibre growth processes which inspire these mechanisms.

These improvements are tested against the state of the art MEDUSA algorithm and simulated dMRI signals compared to real dMRI signals from the white matter of a healthy subject.

5.1 Introduction

Numerical phantoms play a valuable role in the development and validation of many MRI techniques. In particular, numerical phantoms are often used when developing dMRI microstructure imaging techniques where simulations of the dMRI signal in phantoms with known microstructural properties are used in lieu of an in vivo ground truth measure of microstructure [63]. While recently numerical phantoms have proven useful for validating microstructure imaging of grey matter [5], they have more commonly been used for validating WM microstructure, with many studies comparing parameter estimates from fitting their models to the known ground truth from the phantoms e.g. [6–14]. Some recent works directly estimate microstructural features using fingerprinting techniques and machine learning to match simulated signals and the corresponding ground truth microstructure of the numerical phantom to the measured signal [15–18].

Typically, however, there is a mismatch between the complexity of true brain tissue microstructure and the models used in simulation, with simulations simplifying the microstructure. On one hand, *ex vivo* EM studies have revealed the high complexity of real axonal morphology [27, 28, 106]. Reconstructions of axons from these studies show that real WM contains axons with complex morphologies on an individual axon basis such as undulation, beading and non-circular cross sections, as well as non-trivial configurations including orientation dispersion and crossing bundles. On the other hand, the models used in simulation studies often represent axons in WM using simplistic geometrical representations such as parallel cylinders with uniform [14,22–24] or polydisperse [25,26] radii. Some studies investigate the effect of differing configurations of fibres such as simple crossing [105, 118] and orientation-dispersed [10, 13, 119] fibre bundles. A few groups generate WM numerical phantoms with complex fibre configurations for the application to tractography [108, 120]; however realistic microstructural morphology is not the focus of these approaches.

Other studies introduce more microstructural complexity into the numerical phantoms, typically only considering one mode of morphological variation at a time; some examples of this include harmonic beading [55, 58], spines [121], undulation [54, 122] and myelination [123].

Recently, a number of groups have attempted the challenge of combining these features to generate phantoms approaching the morphological complexity and density of real tissue. The most common approach to this is the packing of fibres into densely packed configurations [61, 108, 110, 111]. The typical approach, as taken in the state-of-the-art MEDUSA algorithm [111], is to generate a set of overlapping fibres decomposed into small segments and iteratively refine their positions to remove the overlap between them. Despite their recent progress, further advance of this class of techniques may be limited, because nature does not create fibres before attempting to pack them together. Instead, real axons are guided by chemical cues and fit into available space as they grow [113,114]. Mimicking the natural fibre genesis may prove important for building more realistic phantoms.

To this end, we have proposed an approach to generate WM numerical phantoms by emulating natural fibre growth as presented in Chapter 4 and [117]. In this work we present an extension of this preliminary fibre growth algorithm, adding in further growth mechanisms which mimic a set of key biological mechanisms which govern real axonal growth. The result of this is contextual fibre growth (ConFiG), our WM numerical phantom generator. We assess the performance of ConFiG by measuring the impact of each of the biologically inspired mechanisms on the achievable phantom density and comparing against state-of-the-art MEDUSA phantoms. To test how realistic ConFiG phantoms are, we compare simulated dMRI signal in the ConFiG phantoms to real dMRI data.

The rest of the paper is organised as follows: Section 5.2 describes the ConFiG algorithm, Section 5.3 details the experiments outlined above and Sections 5.4 and 5.5 summarise the contributions and discuss future work.

5.2 Methods

In this section we describe the ConFiG algorithm, beginning with a brief outline of the main steps in growth algorithm before describing the biological mechanisms motivating ConFiG and how each of these are implemented to give the final ConFiG algorithm.

5.2.1 Overview of the ConFiG algorithm

The ConFiG growth algorithm follows the same overall structure as the preliminary implementation described in detail in Chapter 4. The generation of ConFiG numerical phantoms happens in three main steps:

STEP 1 Generate initial growth configuration from user inputs

- Set up target growth parameters for fibres and network upon which growth takes place. This is the same in ConFiG as described in Section 4.2.1.

STEP 2 Grow the fibres using ConFiG growth algorithm

- Fibres grow one-by-one attempting to meet the target distribution established in **STEP 1**. It is this step in which the majority of the ConFiG improvements take place

STEP 3 Generate 3D meshes for dMRI simulation

- Convert fibre skeleta created in **STEP 2** into 3D surface meshes. Again, this step is largely the same as described in Section 4.2.3.

The remainder of this section outlines the biological process governing real axonal growth, and how these processes motivated the final implementation of the ConFiG algorithm.

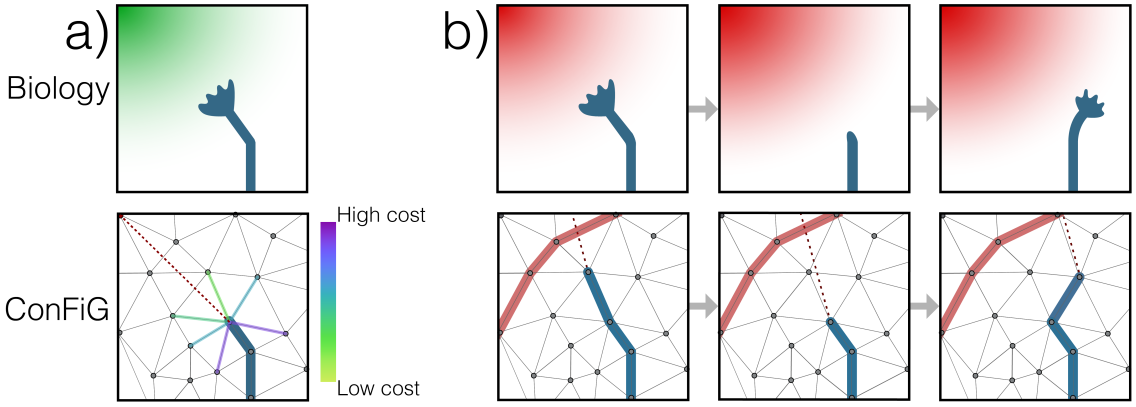

Figure 5.1: Illustration of two of the biological motivations and how they are implemented in ConFiG. a) Growth towards the target is enforced by means of a cost function encouraging growth towards the target point. b) Fibre collapse is implemented by allowing the fibre to move backwards if it reaches a node from which there are no viable steps. The biological figures are adapted from [113]

5.2.2 Biological motivation for ConFiG

In nature, axons grow following chemical cues in their environment through various mechanisms which either attract or repel fibres to guide their growth [113, 114, 124–128]. In an attempt to emulate real axonal growth, mechanisms motivated by the following guidance processes have been integrated into ConFiG:

- Chemoattraction – the process by which fibres are attracted to chemical cues in their environment [113, 125].
- Fibre collapse – a response to a chemorepulsive source whereby a fibre withdraws and regrows in a different direction [127].
- Cell adhesion molecules – chemical signals on the surface of cells which guide axons that come into contact with them [128].
- Fasciculation – the process by which multiple axons come together to form bundles [113, 129].

The following sections detail how mechanisms motivated by these biological processes are implemented in ConFiG while Figures 5.1 to 5.3 illustrate these biological processes alongside their ConFiG counterparts.

5.2.2.1 Chemoattraction

As a fibre grows it must choose in which direction it will move. One of the main processes governing the guidance of real axons is chemotropism; a process by which axons respond to diffusible chemical cues in their environment. One key chemotropic mechanism is chemoattraction, in which fibres are attracted along a chemical gradient towards a target region [113].

To approximate this chemoattractive mechanism, each fibre is encouraged to grow towards its target point (i.e. the target point acts like a chemoattractive source). This is the same guidance mechanism in the preliminary fibre growth algorithm presented in Section 4.2.2. The chemoattractive mechanism and its ConFiG counterpart are illustrated in Figure 5.1a.

From any node in the growth network, the fibre will move along an edge that takes it towards its target while avoiding existing fibres according to a cost function [117]. From a starting node, s, the candidate nodes, c, that the fibre can move to are any nodes that share an edge with s. In addition to its position, each network node stores the maximum fibre diameter, d_c, that can be sustained at that node without intersecting another fibre. The fibre will move to a candidate node according to a cost function consisting of two terms; l_t, which penalises taking very large steps or moving away from the target point, t, and l_d, which penalises moving to a position where d_c is low meaning that the fibre will have to shrink. The cost function for a fibre at a position, s, to move to a candidate node, c, given a target point, t, is

$$l = l_t + f l_d \,, \tag{5.1}$$

where

$$l_t = \frac{1}{2} \cdot \frac{\|s-c\|}{1+\|s-c\|} \cdot \left(1 - \frac{((c-s)\cdot(t-s))}{\|c-s\|\|t-s\|}\right) , \tag{5.2}$$

$$l_d = \max\left(0, \frac{1}{d_0}\left(d_0 - d_c\right)\right) . \tag{5.3}$$

Here, d_0 is the target diameter of the fibre and f is a weighting factor between the two terms. In this work, f is fixed to 0.2 to more strongly weight growth towards the target.

The next node for a fibre will be the candidate node which has the lowest cost according to Equation (5.1). This method of finding a path through the triangulation by choosing the lowest cost node at each position amounts to a greedy best-first pathfinding approach with a heuristic given by Equation (5.1).

Growing fibres along the network using just this chemoattractive mechanism is the minimal implementation of ConFiG that will generate substrates to try and meet the morphological inputs. There are some limitations to this minimal approach however; the greedy growth and the sparse sampling of the space means that fibres can grow into regions from which they cannot grow further and become stuck. Additionally, in this approach, fibres grow independently of one another, whereas real fibres grow forming bundles in the process known as fasciculation.

Sections 5.2.2.2 to 5.2.2.4 describe further mechanisms which were added to enable ConFiG to address these limitations in order to meet more complex morphological priors (e.g. high density and orientation dispersion together).

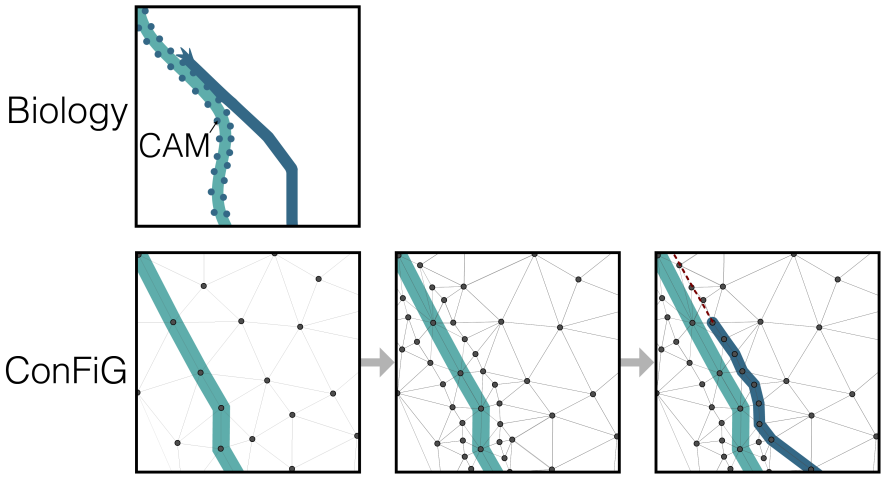

Figure 5.2: Illustration of the contact guidance axonal growth mechanism and the dynamic growth network implemented in ConFiG. The dynamic growth network is implemented as a set of points added around each fibre after growth, enable future fibres to more easily grow along/around existing fibres.

5.2.2.2 Fibre collapse

As mentioned in Section 5.2.2, in ConFiG a fibre can become stuck when there are no possible next steps because all neighbouring nodes are inaccessible. In an attempt to ameliorate this a process mimicking fibre collapse was implemented, illustrated in Figure 5.1b.

In ConFiG fibre collapse, the fibre will move back by an initial distance, g_0, and regrow from there avoiding any nodes in the route it took previously. If the fibre becomes stuck again, it will move back by a further distance, $g_0 + \delta$, where δ is the additional distance to step back. This process is repeated until the fibre reaches the target or gets stuck a user-defined maximum number of times. In this work, $g_0 = 2\,\mu\text{m}$ and $\delta = 5\,\mu\text{m}$ in an approximation of the biological fibre collapse process investigated by Rauch et al. [127] who show fibres collapsing up to $25\,\mu\text{m}$ back towards the soma. The maximum number of steps back is set to 5, meaning that the maximum step back is $27\,\mu\text{m}$, in line with real fibres. If there is no possible route after 5 attempts then the fibre will stop growing and will be removed from the phantom. This process of removing stuck fibres means that the resulting substrate may not always have the same density as the input desired fibre density.

5.2.2.3 Dynamic growth network

In the preliminary implementation of ConFiG [117], the network nodes were initialised pseudorandomly within the growth region and once initialised, the growth network was static, meaning that the nodes and edges of the network were fixed. This limited the growth to the specific instantiation of the network and it could not adapt to where fibres were once they had grown. Furthermore, as illustrated in Figure 4.2, as fibres grow, many nodes become inaccessible due to being within fibres meaning that the network becomes

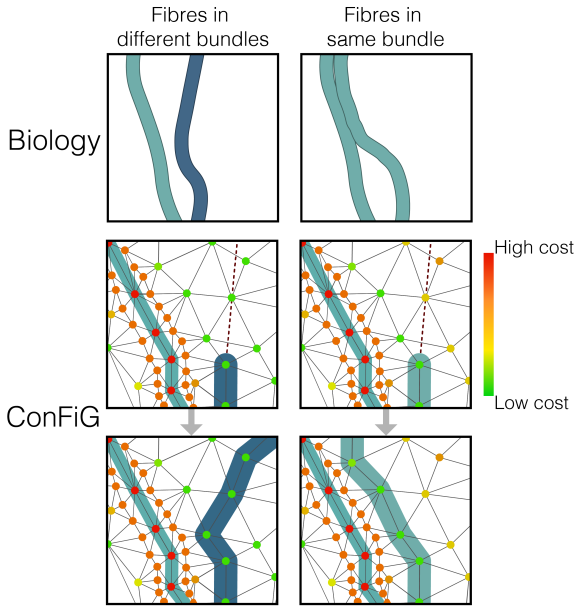

Figure 5.3: Illustration of how the labelled pathway hypothesis is expected to work in biology and its ConFiG counterpart. Fasciculation is implemented using the cost function term in Equation (5.4) which means that fibres in the same bundle are encouraged to stay close to one another.

gradually sparser.

A dynamic growth network was implemented to ameliorate these effects. Now, once a fibre has reached the target, a number of nodes, N_{added}, are generated around the path of the fibre. This gives a denser sampling of the space in regions in which fibres exist and serves to give subsequent fibres more nodes to use to grow along or around that fibre, helping to increase the achievable density by limiting the number of fibres which get stuck. In this work, where the dynamic network is used, $N_{added} = 2500$.

This is also loosely motivated by the contact guidance mechanism in which axons are attracted to or repelled by chemical cues on the surface of other cells, known as cell adhesion molecules (CAMs). Here, the added points act like CAMs meaning that a future fibre which grows can use these points near to the fibre to grow around or along it as if it were following contact guidance cues. Figure 5.2 shows how CAMs work in biological axonal growth alongside the ConFiG dynamic network, illustrating the parallels between the two.

5.2.2.4 Axon fasciculation

One particular role CAMs play is in axon fasciculation, the process in which axons follow a so-called pioneer axon closely, forming a bundle [113, 128]. To mimic the process of axon fasciculation, the term in the cost function penalising moving into regions in which the fibre had to shrink, l_d Equation (5.3), was altered to be conditional on which fibre bundle is closest.

A fibre, f, with a target diameter, d_0, moving to a candidate node, c, which has a maximum sustainable diameter d_c will now have l_d given by:

$$l_d = \begin{cases} \max\left(0, \frac{1}{d_0}\left(d_0 - d_c\right)\right) & \text{if } b_c \neq b_f \\ \text{abs}\left(\frac{1}{d_0}\left(d_0 - d_c\right)\right) & \text{if } b_c = b_f \end{cases} \tag{5.4}$$

Where b_f is an index identifying the bundle that fibre f belongs to and b_c is the index of the bundle that is closest to c (i.e. the index of the bundle of the fibre that set d_c). This means that when c is closest to the same bundle as f, the cost function penalises moving away from that bundle as well as shrinkage, whereas when the bundles differ, it only penalises shrinkage.

This new form of the cost function encourages fibres of the same bundle to stick together while still avoiding fibres of different bundles, inspired by the labelled pathway hypothesis, which states that axons join different fascicles based on different CAMs expressed on the fibres [113]. In this case, bundle indices b_c and b_f act like different identifying CAMs. Figure 5.3 shows how this fasciculation process is expected to happen in biology alongside how the improved cost function encourages a similar process in ConFiG.

5.2.2.5 Global optimisation

Since the growth of fibres in ConFiG takes place on a discrete network of points, the final positions of fibre nodes may be suboptimal for achieving the maximum density. In other words, certain fibres' nodes may be closer to other fibres than they would ideally be in order to reach their target diameter (i.e. the fibre has had to shrink its diameter at that node).

To mitigate against this, a global optimisation step was added at the end of the growth in a procedure similar to MEDUSA [111]. For each point, i, that is part of a fibre, its nearest n neighbours ($j \in NN(i)$) from other fibres are found; in this work $n = 10$. The distance to all of the neighbours is found and the point's position is updated from these distances according to the update vector, $\vec{u}_i$

$$\vec{u}_i = \sum_{j \in NN(i)} D(i,j) \cdot \left(\vec{p}_i - \vec{p}_j\right) , \tag{5.5}$$

where $\vec{p}_i$ and $\vec{p}_j$ are the locations of point i and j. $D(i,j)$ is the function determining whether the interaction is repulsive or attractive:

$$D(i,j) = \text{sgn}\left(r_i + r_j - \|\vec{p}_i - \vec{p}_j\|\right) . \tag{5.6}$$

Here, sgn is the signum function and r_i and r_j are the target radii of point i and j. The sum of these radii is the desired distance between the points since that means the fibres are just touching. $D(i,j)$ imposes that the force is repulsive if the points are closer together than the desired radius and attractive if they are further apart. The update vector is scaled such that if $\|\vec{u}_i\| > 0.2r_i$, the update vector is rescaled so that $\|\vec{u}_i\| = 0.2r_i$. This acts to prevent the update vector from becoming very large.

There is some biological evidence that this kind of interaction between fibres is important in the fasciculation process. The fasciculation process described in Section 5.2.2.4 relies on CAMs detected at the tip of a growing axon, however some studies provide evidence for fasciculation through interactions along axon shafts, known as zippering [129–131]. In zippering, nearby axon shafts attract one another to form more closely packed fascicles, which is a similar process to the global optimisation process in ConFiG.

5.2.3 Summary of ConFiG input parameters

Table 5.1 summarises the key parameters that govern the generation of ConFiG phantoms. Parameters are split into those which define the target microstructural morphology and those which define the instantiation of the growth algorithm. For each parameter, the theoretical range is reported alongside the practical range that has been tested so far. This is due to stochastic nature of the algorithm and the interdependence of the parameters. For instance a very large substrate is possible if very large fibres are chosen, but likely impossible with very small fibres since this will require a very large number of fibres and run into memory limitations.

Table 5.1: Summary of ConFiG parameters split into parameters which define the target microstructural morphology and parameters which define the instantiation of the growth algorithm. For each parameter the theoretical range is reported as well as the practical range that has been tested so far.

	Parameter	Meaning		Theoretical Range	Practical Limits Tested
Target microstructure parameters	L	Size of growth region		$\mathbb{R}_+^3$	$[0,0,0] \rightarrow [50,50,50]$ µm
	ρ	Fibre volume fraction		$[0,1]$	$[0,0.8]$
	μ_r	Mean radius		$\mathbb{R}_+$	$[0.5,2]$ µm
	σ_r	Standard deviation radius		$\mathbb{R}_+$	$[0.1,0.5]$ µm
	GAD	Global angular dispersion	Watson: κ	$\mathbb{R}_+$	$[4,100]$
			ESAG μ	$\mathbb{R}_+$	$[2,10]$
			ESAG γ	$\mathbb{R}_+^2$	$[2,2] \rightarrow [10,10]$
Growth algorithm parameters	N	Number of growth nodes		$\mathbb{N}$	$[0,10^7]$
	f	Cost function weighting term		$[0,1]$	$[0,0.5]$
	g_0, δ	Fibre collapse initial and subsequent step length		$\mathbb{R}_+$	$[1,5]$, $[1,5]$ µm
	N_{added}	Dynamic network no. nodes added		$\mathbb{N}$	$[0,5000]$

5.3 Experiments

In order to assess the performance of ConFiG, a range of experiments were performed. The first set of experiments were performed in order to explore the impact of each of the biologically inspired growth mechanisms. Another set of experiments aimed to show that ConFiG is able to generate substrates with realistic microstructure by comparing generated substrates with real tissue. Additionally, the relationship between the user-specified target morphology and the final output morphology was investigated by comparing resulting phantoms to their inputs (target density and orientation distribution). Finally, a simulation experiment was performed to assess how well ConFiG phantoms can be used to generate realistic diffusion MRI data. The rest of this section outlines these experiments.

5.3.1 Testing the performance of ConFiG

In order to test how each of the biological mechanisms proposed in Section 5.2 impacted on the resulting phantoms, an experiment was devised to measure how phantoms changed when each mechanism was introduced. Four scenarios of interest were generated using several variants of the ConFiG algorithm that included these mechanisms either one at a time or all at once, attempting to grow phantoms as densely as possible:

- one bundle of parallel fibres, target density 75%
- one bundle with Watson distributed fibres ($\kappa = 8$), target density 75%
- two perpendicular crossing bundles, intra-bundle target density 40%
- three mutually perpendicular crossing bundles, intra-bundle target density 30%

These target densities were chosen to ensure that the centre of the phantom (i.e. the crossing region for crossed bundles) had a high target density whilst ensuring that each bundle had a reasonable number of fibres to begin with (>50).

The ConFiG variants were tested by generating phantoms for each of the scenarios starting with the same initial conditions. Each phantom was generated 5 times with a different random seed and results averaged across the seeds.

To investigate the impact of the biological mechanisms on dMRI simulation, a comparison was made between real dMRI signals and simulations from ConFiG phantoms. The NODDI model [13] was fitted to a WM ROI in the corpus callosum of a Human Connectome Project (HCP) [132] subject to provide sensible input parameters (target fibre density and orientation dispersion) for ConFiG to generate phantoms. We generated phantoms using the two extreme cases: the minimal growth case only using chemoattraction, and the complete ConFiG algorithm using all mechanisms. Whilst the random nature of ConFiG means that the resulting phantom will not have morphology exactly matching the input parameters, this approach ensured that the phantoms were reasonable for this proof of concept experiment.

The dMRI signal was simulated in the phantoms using Camino [25, 89] with identical simulation conditions in both cases and the measurement scheme corresponding to the HCP dMRI sequence [133]. An important consideration when performing dMRI simulations is the size of the substrate relative to the diffusion length. The phantom should be large enough that it is bigger than the diffusion length, but not so large as to require excessive computational resources. Owing to the relatively long diffusion time ($43\,\mathrm{ms}$) in the HCP sequence, phantoms were extended with reflected copies [134, 135] to increase their effective size relative to the diffusion length scale.

All dMRI simulations in this work used a bulk diffusivity $D = 2.0\,\mu\mathrm{m}^2\,\mathrm{ms}^{-1}$ in agreement with values used in similar Monte Carlo simulations [23, 25, 105] with 10^5 spins and 2000 timesteps. Standard Camino periodic boundaries were used [25], with dMRI signal was generated from a central region 75% the size of the total phantom to avoid boundary effects [59].

5.3.1.1 Interaction of ConFiG parameters

In practice, certain combinations of ConFiG target microstructure parameters are incompatible and are also influenced by the choice of growth algorithm parameters in Table 5.1. For instance, it is unlikely to be possible to achieve very high densities with large amounts of orientation dispersion due to the complex nature of the packing that will be required.

To investigate this, a series of phantoms were generated, each attempting to reach a high fibre density of $\rho = 75\%$ (chosen to represent the upper end of what would be expected *in vivo*), with variable amounts of orientation dispersion (OD), modulated through changing the target κ ($\kappa = 2...50$) for the Watson distribution. For each ρ, κ combination, a phantom was generated using $10^4, 10^5, 10^6$ and 5×10^6 growth nodes. Each phantom was generated in a $20 \times 20 \times 20\,\mu\mathrm{m}$ region with a mean radius of $1\,\mu\mathrm{m}$ and standard deviation of radius of $0.1\,\mu\mathrm{m}$.

For each phantom, the resulting fibre density was measured, to test which combinations of κ and ρ were achievable. A similar experiment investigating the ability to recover the correct OD is presented in Chapter 6, using the microstructural evaluation methods developed therein.

5.3.2 Diffusion MRI simulation

To qualitatively verify that the simulated diffusion MRI signals from ConFiG phantoms are realistic, simulated signals from ConFiG phantoms were compared to real HCP data [132, 133].

In the real data, the fODF was fit in each voxel using constrained spherical deconvolution (CSD) in MRTrix [75, 136]. Voxels were selected in regions of interest in the midbody of the corpus callosum (CC) and the internal capsule (IC), regions in which a single bundle of fibres is found from the FOD. A third voxel was selected in which three crossing fibre populations were found from visual inspection of the FOD.

In each voxel, the diffusion tensor was fit to the signal and the principal eigenvector used to define a major direction of diffusion in the voxel, n. From this, the normalised diffusion weighted signal was plotted against $|n \cdot G|$, where G is the gradient direction. Additionally, the direction averaged signal was calculated for each b-shell.

To attempt to generate representative microstructure for each voxel using ConFiG, the NODDI model [13] was fitted to the dMRI signal to give some initial parameters for ConFiG. Most importantly, the value of κ for the Watson distribution [115] estimated using NODDI was used to initialise the orientation dispersion in the ConFiG phantoms used to represent CC ($\kappa = 6.2$) and IC ($\kappa = 5.5$) regions. To represent the three crossing (TC) voxel, a phantom generated using three mutually perpendicular crossing bundles was used.

ConFiG phantoms were grown using these initial conditions and the diffusion MRI signal simulated using the Camino Monte Carlo diffusion MRI simulator [25]. For each phantom, the same processing as with the real data was performed, finding the direction dependent and direction averaged signal per b-shell.

5.3.3 3D signal visualisation

In order to better understand how close the simulated signal is to the real signal, a new method for 3D visualisation of the signal was developed. A 6th order SH representation of the simulated signal was calculated for a given b-value and the surface plotted in 3D. On top of this, the real data for that b-value were plotted as points with a line projected along the gradient direction to the surface to show the distance between the real and simulated signals in that direction. Each point is coloured red or blue depending on whether the measured signal is above or below the simulated signal respectively.

Since there is no guarantee that the microstructure is aligned similarly relative to the gradients in the ConFiG phantom and real data (indeed, it is highly unlikely that they are), the Bingham-NODDI model [10] is fit to the both the real and simulated data to give a basis for each, defined by three orthogonal vectors: the principal diffusion direction μ_1, the principal fanning direction μ_2 and a vector mutually orthogonal to both of these, μ_3, defined as $\mu_1 \times \mu_2$.

A rotation matrix is use to bring these two bases into alignment, defined as follows[1]: The rotation matrix R that brings a basis of unit vectors$(\vec{a}, \vec{b}, \vec{c})$ onto another basis $(\vec{d}, \vec{e}, \vec{f})$ must fulfil

$$R\vec{a} = \vec{d} \qquad R\vec{b} = \vec{e} \qquad R\vec{c} = \vec{f}\,. \tag{5.7}$$

This matrix can constructed through the addition of three dyads

$$R = \vec{d}\,\vec{a} + \vec{e}\,\vec{b} + \vec{f}\,\vec{c}\,, \tag{5.8}$$

[1]This solution is based on `https://math.stackexchange.com/questions/1125203/finding-rotation-axis-and-angle-to-align-two-3d-vector-bases`

which can be seen to work if we multiply through by one of the first basis vectors $R\vec{a} = \vec{d}\vec{a} \cdot \vec{a} = \vec{d}$, since $\vec{a}$ multiplied with itself gives unity and is orthogonal to $\vec{b}$ and $\vec{c}$. Each of these dyads can be written as a 3x3 matrix using the vector direct product, giving the rotation matrix

$$R = \vec{d} \otimes \vec{a}^T + \vec{e} \otimes \vec{b}^T + \vec{f} \otimes \vec{c}^T , \tag{5.9}$$

where $\otimes$ denotes the Kronecker product. This rotation matrix, R is used to bring the simulated and real data into the same space for visualisation.

This approach to signal visualisation enables us to get a better feel for how close the signals are and whether there is any pattern to the discrepancy (for instance, is the measured signal high along one axis and low along another). Besides its application here to see how close our simulated and real signals are, this could be a useful tool to see how fitting models are behaving. For instance plotting the data against a surface from a NODDI fit may enable easy identification of cases where the isotropic orientation dispersion assumed in Watson-NODDI is a poor fit to the data and where Bingham-NODDI [10] may be better. In this case of plotting data against a fit, the alignment approach described above does not need to be followed since the fitting parameters will already be in the same space as the real data.

5.4 Results

5.4.1 Impact of biological mechanisms

Each of the proposed biological mechanisms enabled ConFiG to generate phantoms with increased density over the minimal case of chemoattraction only, as is shown in Figure 5.4. Global optimisation resulted in the largest improvement, 17-24%, consistently giving a large improvement. Other improvements performed better for specific phantom configurations. For instance, fasciculation and the dynamic network produced only modest improvements in crossing fibre configurations (4-6%), but performed well in the single bundle cases (11-14%). Fibre collapse was particularly effective in the three perpendicular case, offering 10% improvement.

When combining all of the proposed mechanisms together, the achievable density is higher than any of the improvements individually. This improved performance is comparable to the state of the art, MEDUSA [111], with particularly good performance relative to MEDUSA in the crossing fibre configurations.

This improvement in density can be appreciated visually in Figure 5.5 which demonstrates virtual histology (in which a thin slice through a phantom is rendered) of a parallel fibre phantom for each of the mechanisms. Additionally, Figure 5.6 visually shows the difference in density of the phantoms in 3D between the minimal case of chemoattraction and all biological mechanism for each fibre configuration.

The improvement in the density of phantoms leads to a much more realistic simulated diffusion MRI signal as demonstrated in Figure 5.7. The root mean square error to the real data is reduced by 10 times when using improved ConFiG.

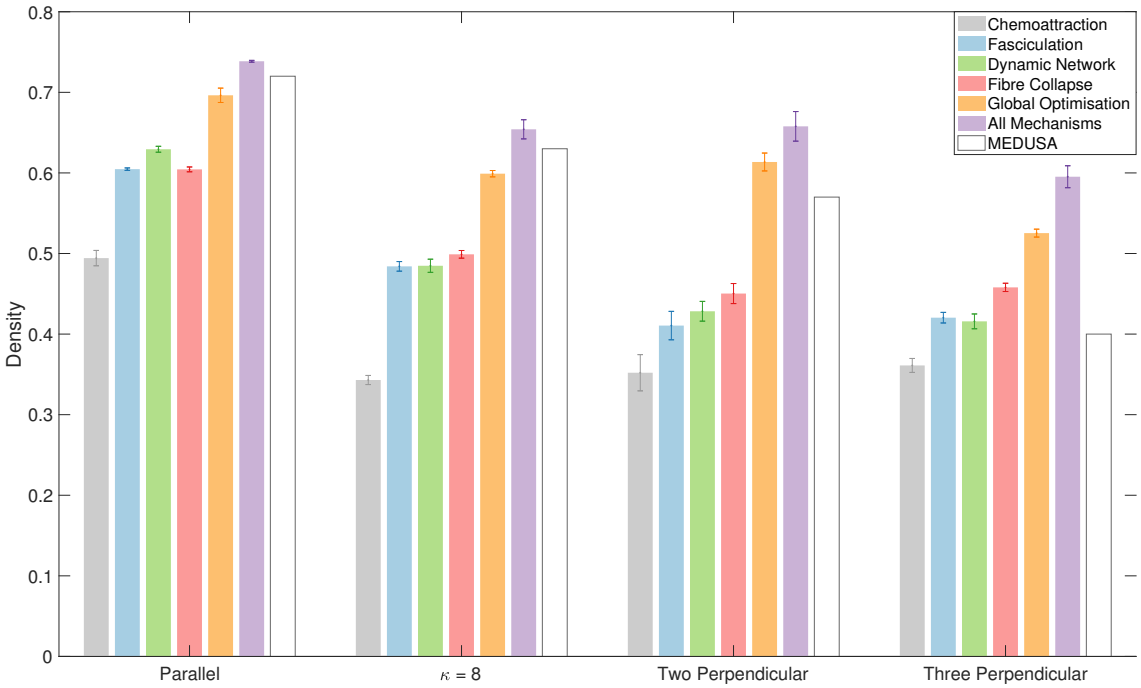

Figure 5.4: Demonstration of the impact of each biological growth mechanism on the density achievable with ConFiG. Each bar shows the mean density for each proposed mechanism, error bars show ± standard error on the mean. MEDUSA values are estimated from Fig. 14 in Ginsberger et al. (Ginsburger et al., 2019).

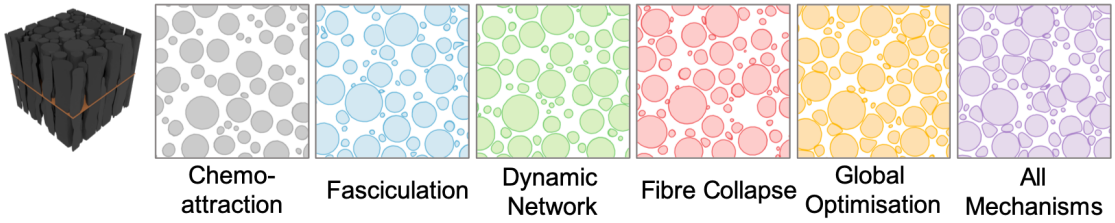

Figure 5.5: Virtual histology demonstrating the impact of biologically inspired mechanism on the final phantom created for one of the parallel phantoms tested. This visually demonstrates the improvement in density. Leftmost image shows the phantom generated with all mechanisms in 3D and the cutting plane used to produce the virtual histology.

5.4.2 Interaction of ConFiG parameters

High amounts of orientation dispersion limit the ability of ConFiG to achieve high densities as shown in Figure 5.8(a). Additionally, the number of ConFiG network nodes affects the outcome, with only the case of $\kappa = 50, N = 5 \times 10^6$ being able to achieve the target density of 75%. Generally, as the number of network nodes increases, the density achievable is higher at the cost of increased time taken for growth as demonstrated in Figure 5.8(b). The time taken for growth does not show a strong dependence on κ indicating that the primary driver of the time required for growth is the number of nodes in the network.

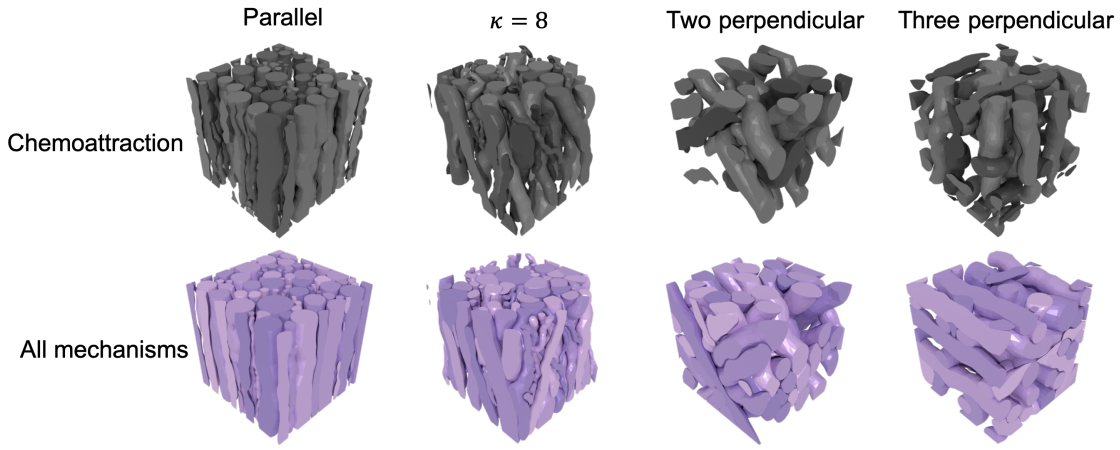

Figure 5.6: Demonstration of the improvement in density achieved when using all mechanisms in ConFiG compared to the minimal implementation using only chemoattraction. Colours chosen to match Figure 5.4.

The results shown are just for a single instance of each combination of κ, ρ and N (for the sake of time) and so there are fluctuations in the density achieved and time taken, however the results give an idea of the limitations and interaction of some the ConFiG parameters.

The result indicates that when attempting to grow very dense phantoms, the number of ConFiG network growth nodes should be carefully considered along with the time available.

5.4.3 Diffusion MRI simulation

Simulated data from ConFiG substrates match real dMRI data well, as shown in Figure 5.9. The direction averaged signal matches well in each case, in particular, for the corpus callosum and three crossing phantoms, the simulated signal matches the real signal closely. The b = $3\,\mathrm{ms}\,\mu\mathrm{m}^{-2}$ signal in the internal capsule and corpus callosum is lower in simulation than in real data. This is to be expected however because as $|n \cdot G|$ approaches 1, the signal reaches the noise floor and the noise-free simulations fall below the measured data.

Figure 5.10 shows the difference between the simulated and measured signal in 3D for the corpus callosum and internal capsule. In this visualisation it is possible to see some systematic differences that are not visible in the signal plotted against $|n \cdot G|$ in Figure 5.9 such as the anisotropic dispersion in the IC data. It can be seen that the real signal is generally lower than the ConFiG along the principle dispersion direction, μ_2, and higher in the orthogonal dispersion direction, μ_3, highlighting that the real data shows some anisotropic dispersion which isn't captured by the ConFiG phantom which was generated using the isotropic Watson distribution.

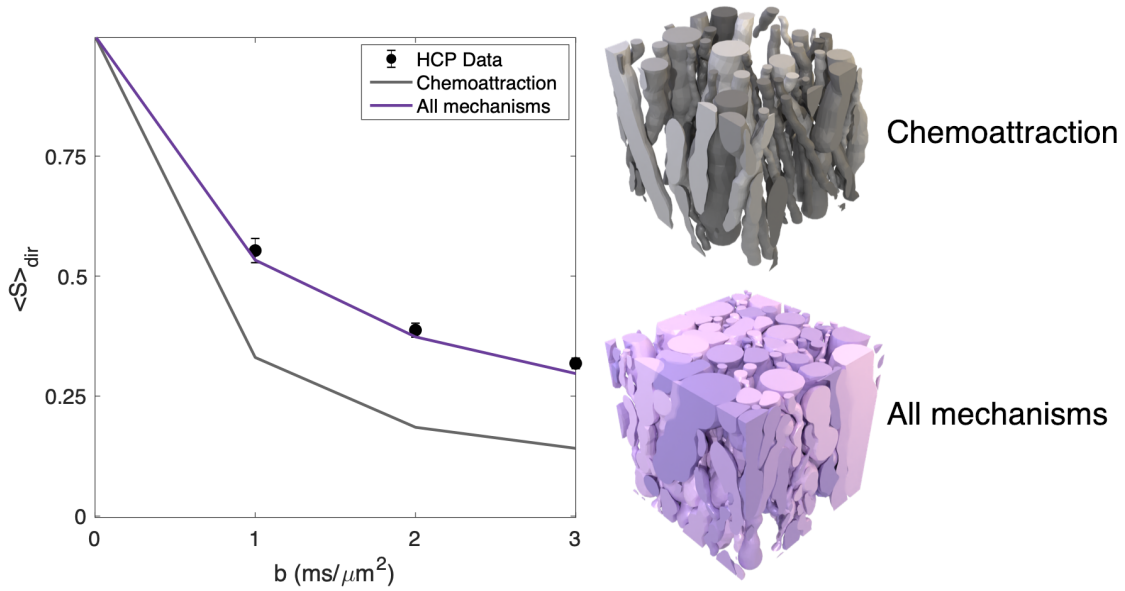

Figure 5.7: Left: Direction averaged signal attenuation for real HCP data (± standard deviation over ROI) and simulated data from the minimal ConFiG implementation using only chemoattraction and using all growth mechanisms ConFiG showing that ConFiG can produce realistic dMRI signals. Right: The original and improved ConFiG phantoms used to generate the signal on the left. Simulations performed with 10^5 spins, 2000 timesteps and HCP measurement scheme (Stamatios N. Sotiropoulos et al., 2013). Diffusivity set to $2.0\mu m^2\ ms^{-1}$, chosen to be consistent with previously reported values (Hall and Alexander, 2009; Nilsson et al., 2009; Rensonnet et al., 2017).

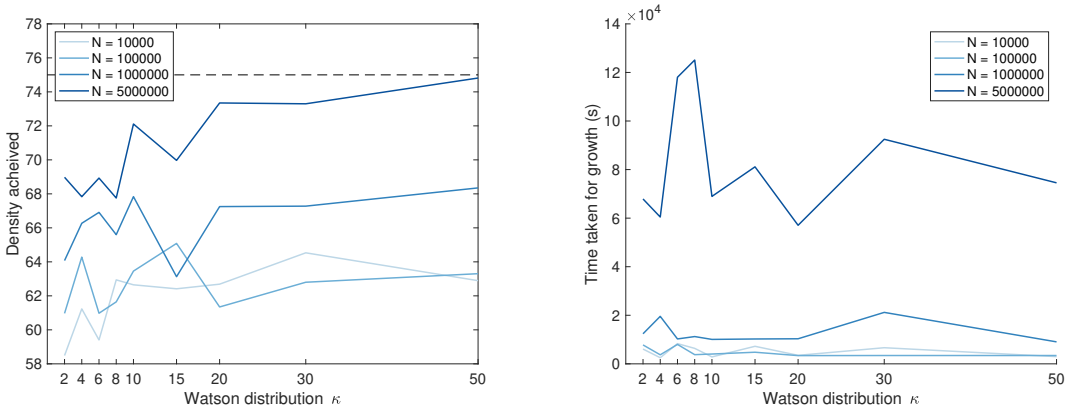

Figure 5.8: Interaction between target density and OD as a function of number of ConFiG network nodes.

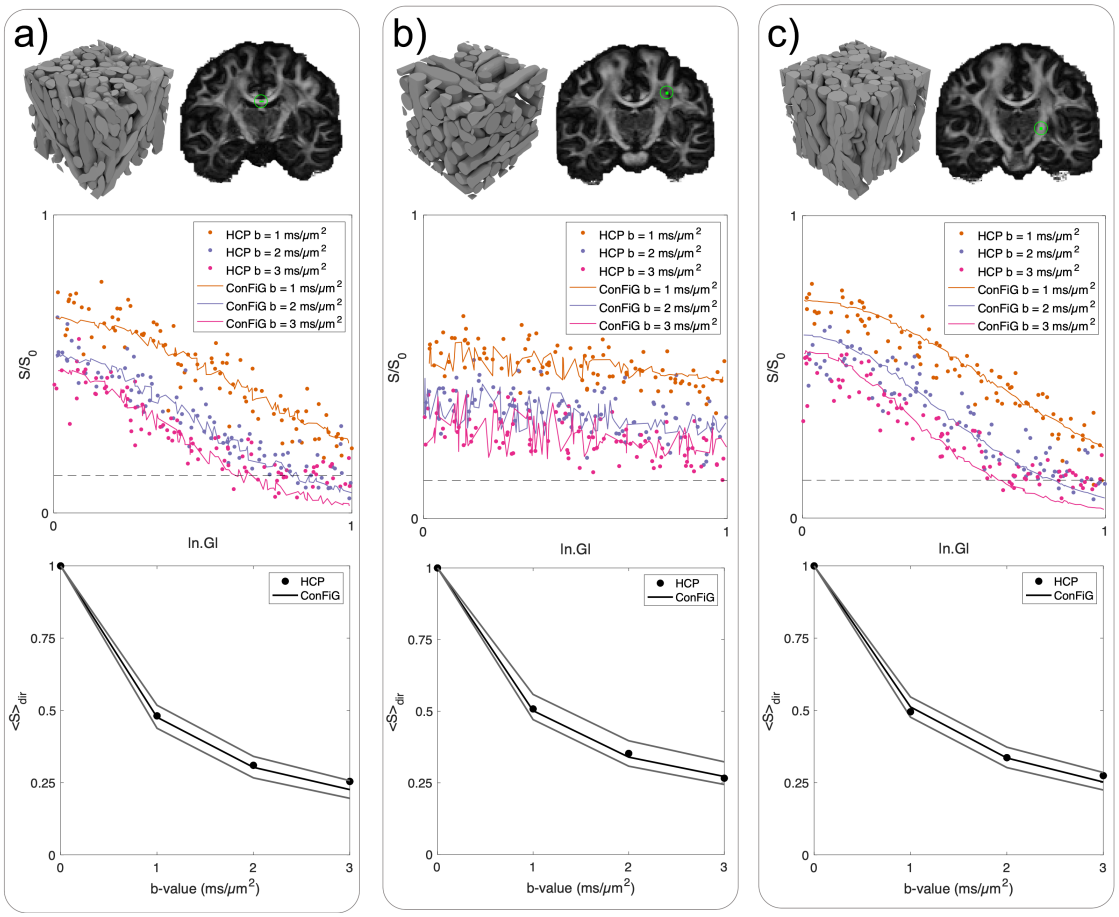

Figure 5.9: Comparison of diffusion MRI simulations and real data from three different brain regions: a) a voxel in the midbody of the corpus callosum, with phantom with volume fraction 55% and mean orientation from z 25°. b) a voxel in which there are three crossing bundles, with phantom of three crossing bundles with volume fraction 50% and c) a voxel in the internal capsule, with phantom with volume fraction 58% and mean orientation from z 22°. Top row shows the ConFiG phantom and corresponding WM voxel. Middle row shows the direction dependent signal for ConFiG (lines) and HCP data (dots). Grey dashed line demonstrates estimated noise floor. Bottom row shows the direction averaged signal. Black lines correspond to phantom in top row. Grey lines are signal from phantoms with the same orientation distribution as the black line in each plot but different densities to show that ConFiG has the flexibility to generate a wide range of realistic signals. Simulations performed with 10^5 spins, 2000 timesteps, diffusivity $2.0\,\mu m^2\,ms^{-1}$ and HCP measurement scheme [137]

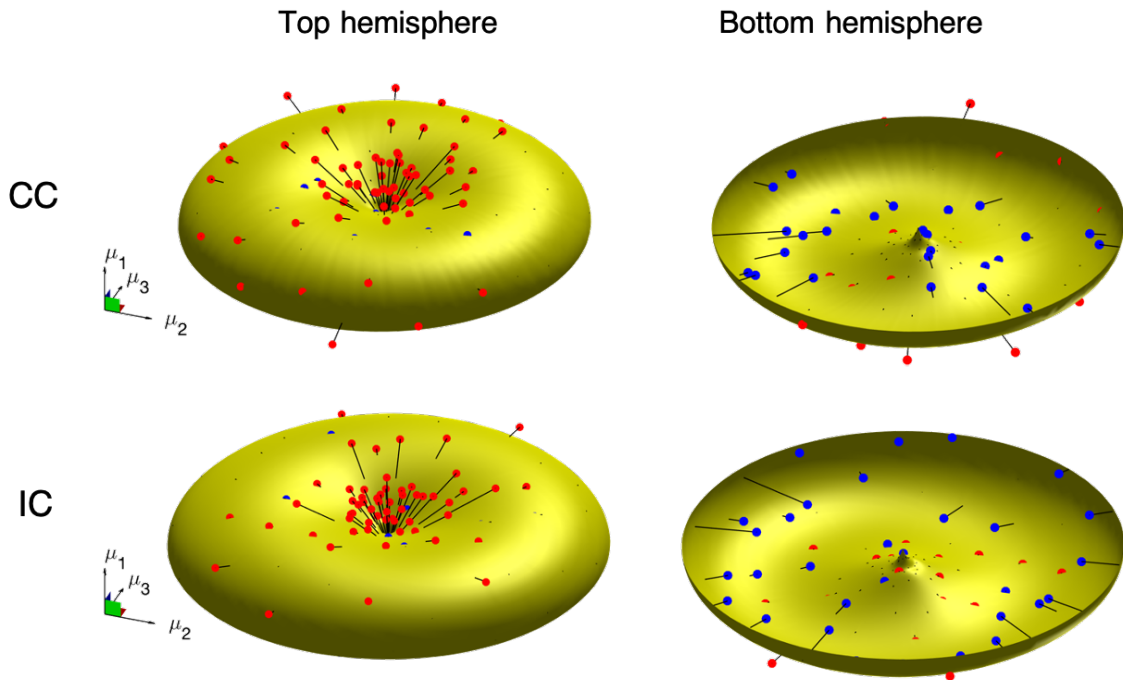

Figure 5.10: 3D visualisation of simulated and real dMRI signals in the corpus callosum and internal capsule for $b = 3\,\mathrm{ms}\,\mu\mathrm{m}^{-2}$. Surface is 6th order SH representation of the simulated signal. Points are measured data coloured such that red points have measured signal higher than corresponding simulated signal and blue points have signal lower than simulated signal. Lines connected to each point show distance to corresponding point on the simulated signal surface.

5.5 Discussion

ConFiG is shown to produce WM numerical phantoms with state-of-the-art performance, producing phantoms with higher fibre density than the MEDUSA approach, particularly in complex arrangements of fibres. These improvements lead to much more realistic simulated signals from ConFiG phantoms than the preliminary fibre growth algorithm as demonstrated in Figure 5.7. ConFiG is able to produce this realistic microstructure at high fibre density by following simple biologically inspired growth rules.

The ConFiG growth algorithm of course depends on the specific instance of the growth network, meaning that the resulting phantom for the same input fibre configuration will be different for different network choices. This is alleviated to an extent by using the dynamic network introduced here, however the phantom will still be dependent on the initialisation of the network. The dependence appears to be relatively minor as is demonstrated by the small standard errors on the mean density shown in Figure 5.4 across the five repetitions.

The diffusion MRI simulations shown in Figure 5.9 demonstrate the ability of ConFiG to generate phantoms which reproduce real diffusion MRI data well. These simulations, however, are just three examples of ConFiG phantoms and corresponding simulations. Using NODDI as input to ConFiG means that the resulting phantoms have sensible morphologies and are shown to generate signals that match the real tissue well, though there may be other configurations that can better reproduce the signal. As an example, the b = $3\,\mathrm{ms}\,\mu\mathrm{m}^{-2}$ signal from the internal capsule is higher at low $|n \cdot G|$ in the simulated versus the real data (Figure 5.9c). One explanation of this is that the phantom generated does not have microstructure accurately representing this region, for instance the phantom may have too little dispersion caused by ConFiG under representing the target orientation dispersion, as seen at low κ in the following chapter in Table 6.1.

Another explanation, as highlighted in Figure 5.10, is that the orientation dispersion distribution used is not right to fully characterise the signal. The fact that the real signal is lower in the principle dispersion direction μ_2 and higher in the orthogonal dispersion direction μ_3 suggests that that there is anisotropic dispersion in the real data which is not captured by the ConFiG. In this case, the ConFiG phantom was generated using a Watson distribution for the orientation dispersion, which gives the isotropic OD seen in Figure 5.10, however, it may be better to use something like the elliptically symmetric angular Gaussian (ESAG) or Bingham distribution to better represent the IC.

Eventually, it may be possible to find a better matching phantom using a computational modelling approach such as that proposed in [18], however the simulations presented are sufficient to demonstrate a proof-of-concept that ConFiG can be used to generate realistic simulated dMRI data.

Limitations and future work

One limitation of ConFiG is that the algorithm relies on the space being sufficiently densely sampled by the growth network. This can require a large number of nodes for a large phantom, becoming prohibitively memory and time expensive. This is investigated in Section 5.3.1.1 which demonstrates that in order to achieve high densities, a very large number of nodes are required ($> 5 \times 10^6$). In this case, the phantoms generated are quite small ($20 \times 20 \times 20\,\mu\mathrm{m}$), while larger phantoms will require a larger number of nodes. In fact, the more relevant value than the *number* of nodes is the *density* of nodes which may be remedied in future versions of ConFiG by changing the input parameter from the number of nodes to the density.

The dependence of the resulting phantom on the density of network nodes can be addressed by growing the fibres in small sub-regions local to the head of the fibres rather than the whole space at once. For instance, rather than filling the entire space of growth with nodes, it is possible to fill a small layer of the space with points and then grow layer by layer. In this way, it is possible to achieve a high density of nodes using fewer nodes than when covering the entire space.

One further potential limitation of ConFiG is that once a fibre has grown, it is static. The fibre will remain fixed in place and all other fibres will have to grow around it. One problem with this is that once the fibres are fixed, they may create pockets of inaccessible space which limits the space available for following fibres. Additionally, in real tissue, axons are flexible and non-rigid, meaning that it may be more realistic that growing fibres can push existing fibres out of the way to make more space for growth. A potential approach to ameliorate this would be to have an optimisation procedure during growth, similar to the global optimisation introduced in this work but optimising the shape of a fibre as it grows.

A limitation of the current study is that the simulations assume a single diffusivity for the intra and extracellular spaces and no permeability of the axonal membranes. Furthermore, effects such as T2 and magnetic susceptibility are ignored. These effects are a limitation of the simulator used rather than ConFiG, and work is planned to improve these aspects of the simulator for more realistic simulated signals.

While the dMRI simulations presented here provide a good sanity check that the ConFiG phantoms are realistic since the simulated signals match well with real dMRI signals, a more thorough validation of the microstructure generated is planned. Chapter 6 outlines a series of experiments performed to compare the microstructure generated with ConFiG to real axons segmented from EM [28].

We will work towards decreasing the difference between the input and output morphological measures, particularly in complex situations, such as high orientation dispersion and crossing bundles. This can be addressed through the improvements to ConFiG mentioned here and also by improving the strategy for the generation of starting and target points for each fibre. For instance, currently it is not intuitive how starting and target points should be arranged to achieve a desired density in crossing regions of fibres.

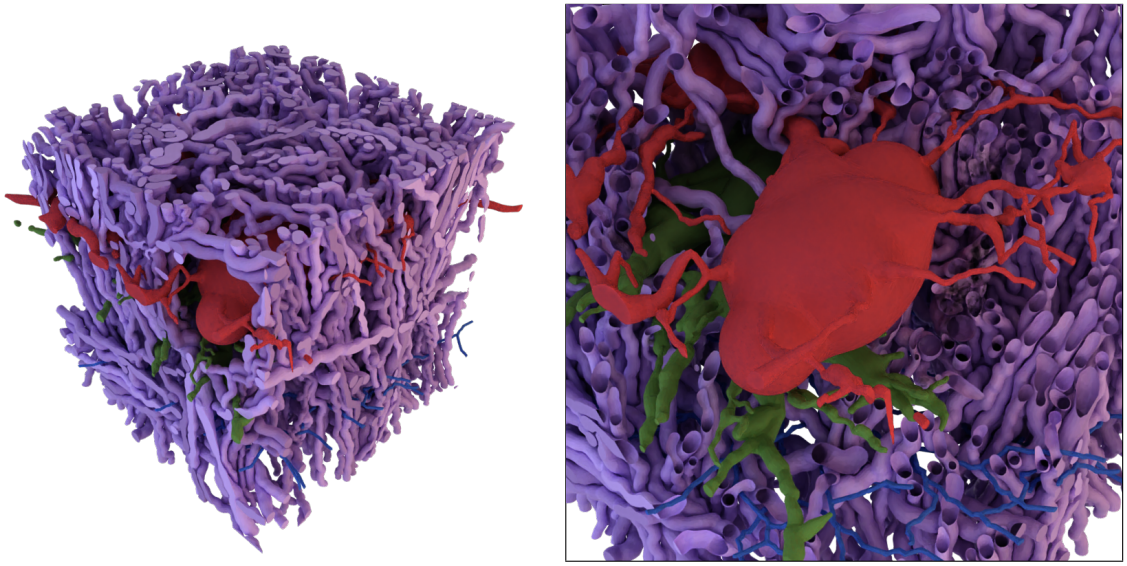

Figure 5.11: Proof-of-concept ConFiG phantom grown around real cells from `https://neuromorpho.org`. One astrocyte (green), oligodendrocyte (red) and microglial cell (blue) were added to the network before fibres grew and the network updated in the same as once a fibre has grown to prevent fibres intersecting the cells. Fibres then grew as normal.

One planned extension of ConFiG is to implement periodic boundary conditions in the growth network, enabling the generation of fully periodic phantoms. This would enable ConFiG phantoms to be generated in relatively small volumes and tiled for simulation, accelerating the process of generating a wide range of phantoms and the memory required to store each phantom.

The core growth algorithm for ConFiG relies on a set of starting and target points, a connected network of nodes and some rules defining the growth. As such, ConFiG is very flexible since the exact form of each of these components can be modified based on the application. One example of a simple modification that may be explored is the order of growth of the axons. Currently, in the absence of any clear biological precedent know to the authors, fibres grow in a random order, but it may be possible that there is a better order such as growing large diameter axons first, or central axons in a bundle growing first.

In this work, ConFiG is applied to the case of densely packed axons, without contributions from neuronal cell bodies or other processes. A planned future extension of ConFiG is to allow for the addition of glial cells such as astrocytes and oligodendrocytes [60] to the extracellular space to make the virtual WM tissue more realistic. A proof-of-concept of this is shown in Figure 5.11 where we generated fibres which grow around realistic cells from `https://neuromorpho.org`. An astrocyte, an oligodendrocyte and a microglial cell were manually placed into the space before the fibres grew and the network was updated in the same as once a fibre has grown meaning that the growing fibres grow around the cells.

Additionally, to further add to the realism of ConFiG phantoms, realistic myelin may be modelled, creating spiral layers wrapped around the axons [123]. Furthermore, intra-axonal structures such as mitochondria and microtubules may be added to investigate their contributions to the diffusion weighted signal.

A planned future application will be to use ConFiG to generate a wide range of phantoms with different microstructural features. These can then be used to create a computational model to estimate microstructural features directly from the diffusion MRI signal in an approach similar to previous works [15–18, 138].

Applications beyond diffusion MRI

As mentioned in the introduction, axonal configuration impacts MR signals beyond dMRI. One potential avenue of exploration would be to investigate the impact of realistic axonal configurations on magnetic susceptibility in a similar way to Xu et al. [21], extending their 2D simulations to use realistic 3D geometries generated in ConFiG.

The virtual histology presented in Figure 5.5 shows an approximation of a histology image generated using ConFiG substrates. In this work, the purpose of this is to show the impact the growth mechanisms have on ConFiG phantoms. For this reason, the virtual histology is simply produced by rendering images with false colours. It may be possible, however to generate more realistic electron microscopy images using a physically realistic electron microscopy simulator [139–141] which may be used to train and test axon segmentation routines. This may be of particular use for cases of fibres parallel to the electron microscopy plane or crossing bundles which are typically difficult for 3D reconstruction and segmentation algorithms.

The 3D meshes generated by ConFiG are saved in the PLY format, a widely used format for storing meshes for many purposes. This means that the ConFiG phantoms may be used in other types of simulations such as polarised light imaging [142, 143] or molecular dynamics simulations using software such as MCell (`https://mcell.org`) [90, 91, 144] or LAMMPS (`http://lammps.sandia.gov`) [145].

5.6 Conclusion

ConFiG enables the generation of realistic white matter numerical phantoms achieving state of the art fibre density whilst ensuring realistic microstructural morphology by following biologically motivated rules. This realistic microstructure is shown to generate realistic simulated diffusion MRI signals, opening up the possibility to use ConFiG to create a realistic computational model of WM microstructure. ConFiG outputs fibre meshes which can be used for realistic diffusion MRI simulations or can be processed to produce virtual histological slices, allowing for further potential applications outside of diffusion MRI.

6 Microstructural Evaluation

Chapter Contents

6.1 Introduction 88
6.2 Microstructural measurements 88
6.3 Experiments 91
6.4 Results 93
6.5 Discussion 99
6.6 Conclusion 100

Chapter Summary

This chapter presents a series of experiments performed using ConFiG phantoms to test their microstructural realism, something essential to prove the value of ConFiG phantoms.

Firstly, the methodology used to calculate microstructural features from ConFiG phantom meshes are presented. The dependency of some of these features on the input morphology is demonstrated, followed by a series of experiments which compare ConFiG phantoms to real white matter reconstructed from electron microscopy.

6.1 Introduction

As demonstrated in Chapter 5, ConFiG is able to generate white matter (WM) numerical phantoms which can be used to generate realistic simulated diffusion MRI (dMRI) signals. Whilst this a valuable check of the realism of the phantoms, this does not guarantee that the underlying microstructure is realistic since simplistic representations such as straight parallel cylinders can be used to generate realistic signals in certain cases.

To investigate how well axons generated using ConFiG represent real WM axons, a series of experiments were conducted to compare ConFiG axons to real axons reconstructed from electron microscopy (EM). Typical EM techniques such as scanning electron microscopy (SEM) and transmission electron microscopy (TEM) provide only 2D images of tissue cross-sections, giving only limited information on axonal morphology. Recent technological improvements - primarily improvements to and increased accessibility of serial block-face scanning electron microscopy (SBEM) [146, 147], a technique to take SEM images of many sequential tissue slices - have enabled new possibilities for the study of axonal microstructure in 3D.

Although high-resolution 3D imaging of tissue is possible using these SBEM techniques, segmentation and quantification of the resulting images has remained challenging, meaning that there have only been a few studies investigating the 3D microstructure of segmented axons in WM [27, 28, 106]. Of these, only Lee et al. [28] have made their axonal segmentation publicly available so it is this data set that we use to compare with ConFiG.

The rest of the chapter is arranged as follows, Section 6.2 describes the procedures used to extract microstructural features from the 3D meshes generated by ConFiG. Section 6.3 outlines a set of experiments that were performed using these microstructural measurement techniques to assess how the complexity of fibre arrangements affect the axonal morphology and to compare ConFiG axons to real axons segmented from EM. Section 6.4 presents the results of these experiments and Section 6.5 summarises the contributions and discusses future work.

6.2 Microstructural measurements

In order to test how realistic the microstructure generated using ConFiG is, microstructural measurements of diameter distribution and orientation distribution were calculated using methods designed to be comparable with previous studies on ex-vivo tissue [27, 28].

A centre line is generated from each of the fibre meshes by aligning the ends of each fibre with the z-axis and connecting the centre of mass of 100 equidistant slices through each fibre, following the approach taken by Lee et al., [28]. This is illustrated in Figure 6.1.

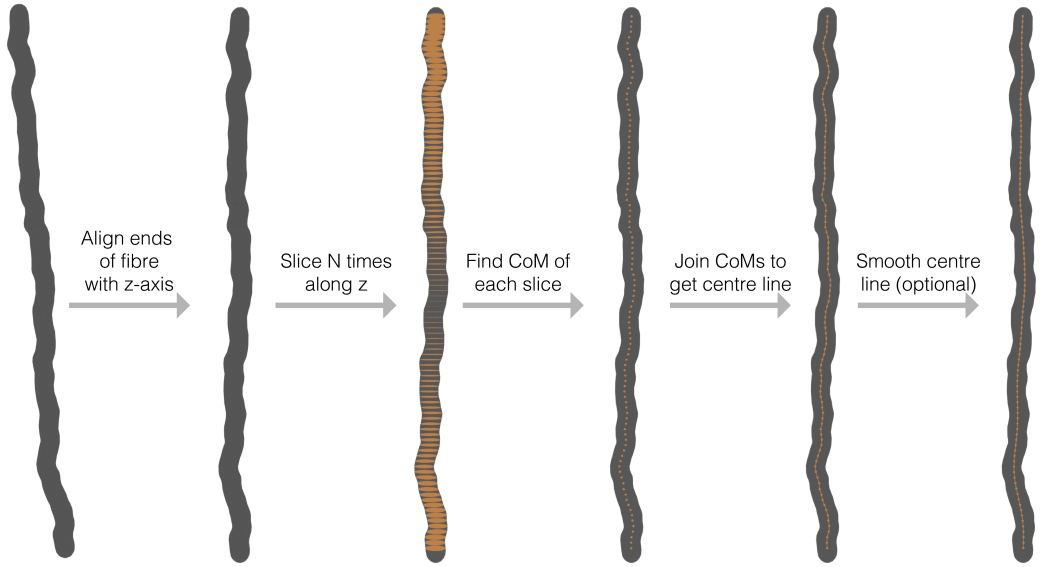

Figure 6.1: Centre line extraction of fibres. Each fibre was sliced N times along the z-axis, connecting the centre of mass of each slice to create the points in the centre line. This line could then be optionally smoothed according the diffusion time coarse-graining effect, as in [28]

Each segment in this centre line could then be used to assess the microstructure of the phantom. The direction of each segment was used to assess the orientation distribution of the phantom, illustrated in Figure 6.2. Following the approach of Lee et al., [28], the direction of each segment was projected onto the surface of a triangulated unit sphere [148]. For each triangle, the number of segments pointing in that direction was used to colour the triangle to visualise the orientation distribution.

A second approach was devised to better visualise orientation distributions in 3D to aid differentiation of crossing bundles and antipodal symmetry. In this approach each vertex was raised above the surface of the sphere proportionally to the number of segments pointing in its direction as illustrated in Figure 6.2.

To measure the diameter profile along fibres, the direction of each segment gave the normal to a plane used to cut the fibre using Boolean intersection to give a cross section of each fibre at each segment. The diameter profile along the axon was generated by calculating the equivalent diameter of a circle with the same area as the fibre cross section. This process is illustrated in Figure 6.3.

6.2.1 Virtual histology and 2D morphological measures

Virtual histological slices were generated to compare ConFiG substrates to real white matter analysed using histology. Histological slices were found by calculating the Boolean intersection of a cutting plane with the generated fibre meshes using Blender (`https://blender.org`). A myelin sheath was added to the fibres when generating virtual histology for visualisation purposes. Virtual histological images were rendered

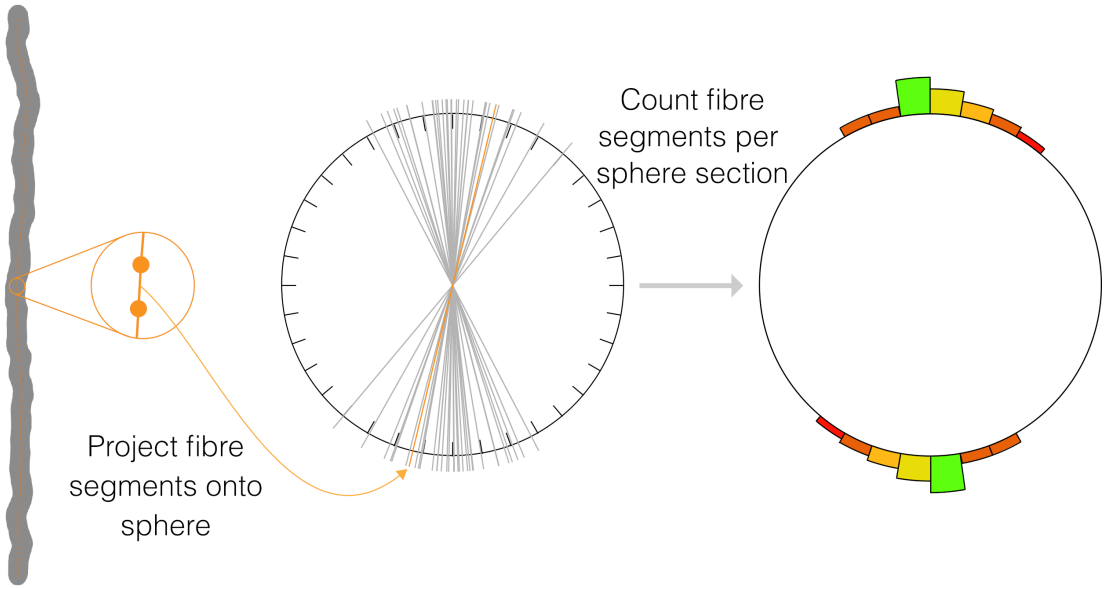

Figure 6.2: Orientation distribution calculation. Each segment of a fibre was projected onto the surface of a triangulated sphere, here illustrated with a sectioned circle. For each section in the sphere, the number of fibre segments going through that section was used to colour and/or raise the surface to visualise the orientation distribution. Since the diffusion process is symmetric about the origin, each fibre segment was projected onto the sphere forwards and backwards.

with a resolution of $5\,\mathrm{nm} \times 5\,\mathrm{nm} \times 100\,\mathrm{nm}$, chosen to be comparable to real histological white matter measurements [27, 28].

In order to compare ConFiG virtual histology to real histology, virtual histological slices were rendered in binary black and white to compare against intra-axonal segmentations from [28]. Slices from real histology, ConFiG phantoms and a parallel cylinder phantom were processed using the MorphoLibJ plugin for ImageJ [149–152] to extract morphological features: circularity (4π × Area/Perimeter^2), convexity (Area of shape/Area of convex hull), eccentricity of fitted ellipse and Area/πr_{max}^2 for each axon. Axons touching the edge of the image were removed since truncation from the image edge would skew these microstructural metrics.

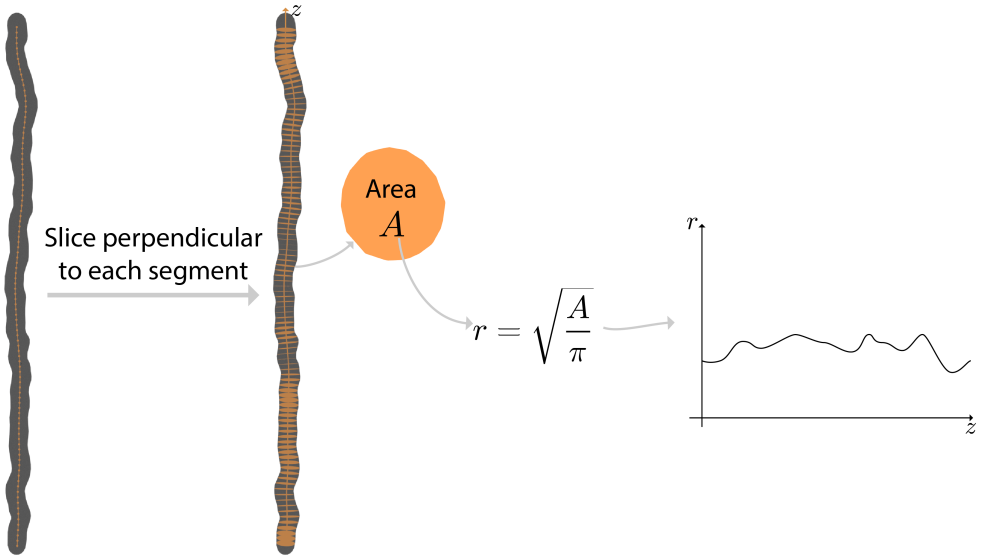

Figure 6.3: Calculation of the diameter distribution. A slice is taken through each fibre perpendicular to every segment in the centre line. The area of each of these slices is used to find a circle equivalent radius or diameter using $A = \pi r^2$.

6.3 Experiments

6.3.1 Relationship between input and output morphology

As mentioned in Section 5.2, the nature of the ConFiG growth algorithm means that the microstructural morphology of the phantoms may not match the user input. Some fibres may become stuck and fibres cannot typically grow in a straight line, affecting the density and orientation distribution.

To investigate this, we generated a series of ConFiG phantoms with Watson distributed [115] orientation dispersion with $\kappa = [8, 10, 15, 20, 30, 50, 100]$ and target density, $\rho = 75\%$. The target density of 75% is chosen as this is the upper limit of what is achievable empirically and towards the higher end of expected axonal volume fraction. Additionally, with 75% achieved, lower densities can be generated easily, either by running ConFiG in full, or simply by removing or shrinking fibres.

The mean and standard deviation of the angle from z, μ_θ and σ_θ respectively, for each κ was calculated by taking 10000 samples from the Watson distribution and this was compared to μ_θ and σ_θ of the ConFiG fibres. Additionally, the density of the ConFiG phantoms was compared to the target density of 75%. Each phantom was generated in a $20\,\mu\text{m} \times 20\,\mu\text{m} \times 20\,\mu\text{m}$ region, using 2.5×10^6 nodes in the growth network.

6.3.2 Packing induced microstructural complexity

When axons grow, they must bend and bulge to grow around one another without colliding. It is a reasonable assumption to make that the level of divergence from straight cylinders is higher in more complex arrangements of fibres (i.e. in complex fibre arrangements such as high orientation dispersion, fibres are more likely to cross paths and so are more likely to need to deform around one another).

In order to test whether this effect is present in ConFiG phantoms, we used the phantoms from Section 6.3.1, with the addition of two phantoms generated with target κ of 2 and 4 to attempt to generate even more complex arrangements. Each of these phantoms were generated with fibres that attempted to go straight from start to target point with a constant radius, so any deviations arise from the growth procedure itself rather than input settings.

For each fibre, the centre line was extracted as described in Section 6.2 and the end points of each fibre were used to calculate σ_θ for the fibres if they were joined by straight lines. This is used as a proxy for the complexity of the phantom since a higher σ_θ will lead to more fibres crossing one another. Additionally, for each phantom, the coefficient of variation of the diameter along each fibre (CV) was calculated as a measure of the amount of fibre beading and the tortuosity for each fibre was calculated as $\tau = L_{\text{fibre}}/L_{\text{ends}}$, the ratio of the total length of the fibre and the length of the straight line between endpoints.

6.3.3 Variability of repeated instances of ConFiG

As discussed in Chapter 4, ConFiG uses a network of pseudorandom points to facilitate the growth of WM phantoms. This means that different instances of ConFiG with the same inputs (other than a different random seed) will produce slightly different phantoms.

In an effort to investigate this affect, repeated phantoms were generated using exactly the same input parameters for three conditions (a subset of the conditions used in the previous experiment):

- Watson distributed $\kappa = 2$ - highly dispersed
- Watson distributed $\kappa = 6$ - Typical WM dispersion
- Watson distributed $\kappa = 100$ - Highly coherent

Phantoms were generated for each condition 40 times, each time keeping all other parameters the same but using a different random seed. The mean coefficient of variation of diameter, μ_{CV}, and tortuosity τ for each fibre in each phantom was calculated to assess how much microstructural variation can arise due to the random instantiation of the ConFiG algorithm.

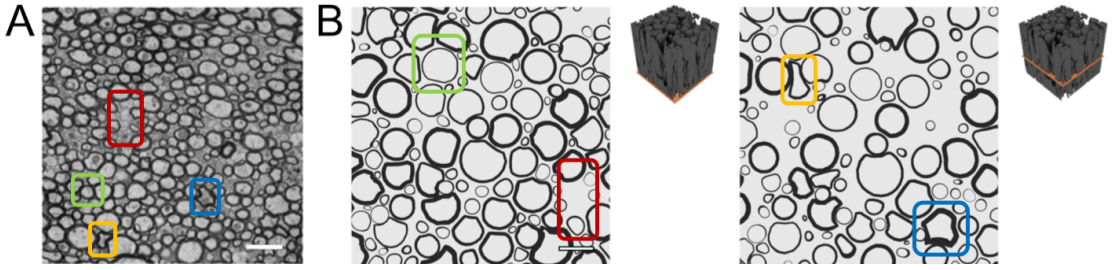

Figure 6.4: Comparison of real and virtual histology. A) Light microscopy of rat ventromedial WM in thoracic spinal cord. Reproduced from Baxi et al. 2015 (Baxi et al., 2015), scale bar 2 µm. B) Two virtual histological slices from a ConFiG generated phantom. Phantoms are rendered to have similar colours to electron microscopy studies. The exact contrast and fibre bundle configurations are different between the real and virtual tissues, but the general morphology of the myelinated axons are captured well using ConFiG as highlighted by corresponding boxes. Yellow and Blue: axons severely deformed between other axons. Red: Pockets of empty space forming. Green: Largely circular axon surrounded by other axons deforming around it. Scale bar 2 µm.

6.4 Results

6.4.1 Microstructural measures and virtual histology

The microstructural morphology generated using ConFiG is comparable to results from real data as demonstrated in Figures 6.4, 6.5, 6.7 and 6.8. Figure 6.4 demonstrates virtual histology of a ConFiG phantom alongside a real EM image from mouse corpus callosum [153]. The exact microstructural features, such as diameter distribution, as well as the EM contrast do not exactly match between ConFiG and the real data. However, ConFiG is able to capture the general morphology of real axons as highlighted in Figures 6.5 and 6.7. In particular, ConFiG is able to capture complex fibre cross-sections such as in the case of fibres squashed into small spaces. This is the first model of white matter able to handle complex fibre cross-sections such as this to our knowledge.

ConFiG morphological metrics calculated slice-wise on the virtual histology correspond much more closely to real axons than the same metrics calculated for parallel cylinders, as shown in Figure 6.5. While cylinders produce a delta function at one extreme of each metric, ConFiG phantoms produce much closer distributions to the real data. Figure 6.6 shows each of these slices coloured by their morphological metrics, demonstrating which fibres are contributing to the distributions of each metric shown in Figure 6.5.

The diameter distribution of a ConFiG substrate is compared to a reconstruction from real EM data [28] in Figure 6.7. ConFiG is able to capture the general profile of axonal variations well, with the overall shape of the diameter distribution matching well. The distribution of the coefficient of variation along ConFiG axons is slightly narrower with a smaller mean than real axons, though these discrepancies may be alleviated with a different choice of input parameters to ConFiG.

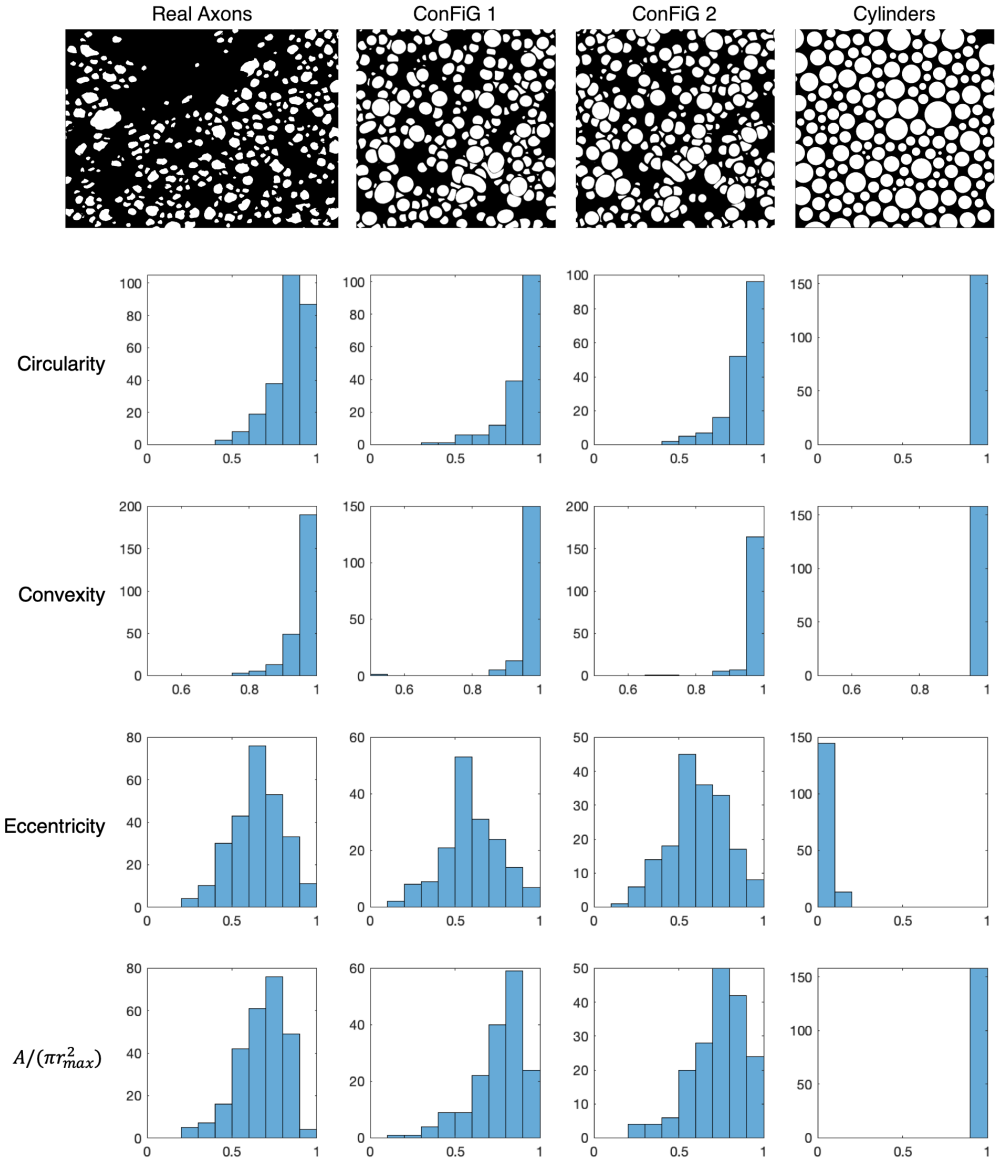

Figure 6.5: Slice wise morphological metrics calculated for real axons, ConFiG phantoms and parallel cylinders. Across each of the metrics, ConFiG produces much more realistic distributions than the cylinder phantom. Some of the cylinders have a non-zero eccentricity, but this arises since the metrics are calculated from binary images where the pixelated circles may appear to not be perfectly circular.

ConFiG is also able to generate fODFs comparable to real tissue as shown in Figure 6.8. Here the fibre paths are smoothed with a Gaussian kernel equivalent to having a diffusion time of $1\,\mathrm{ms}$ and diffusivity of $2\,\mu\mathrm{m}^2\,\mathrm{ms}^{-1}$ as in Lee et al. [28]. The orientation dispersion is introduced to ConFiG phantoms using the elliptically symmetric angular Gaussian (ESAG) [154] to best approximate the EM data and also using isotropic Watson distributed directions to demonstrate the flexibility of ConFiG.

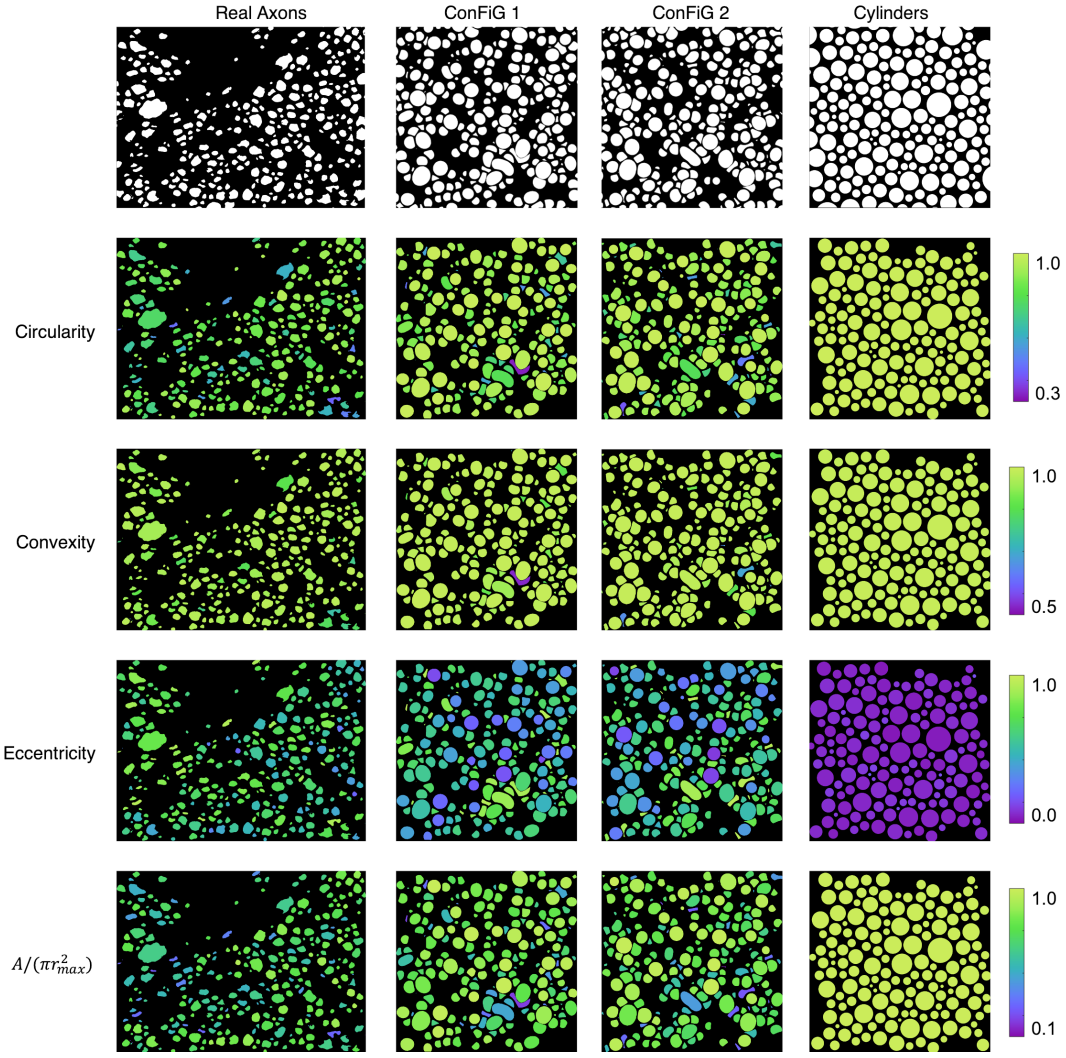

Figure 6.6: Slice-wise microstructural measurements with colourmap demonstrating which fibres are being picked out by which metrics. Again, this shows that the standard cylinder approach generates uniform cross-sections while nature produces much more variation which ConFiG is much closer to capturing.

6.4.2 Relationship between input and output morphology

The morphology of ConFiG phantoms matches the input morphology well, as shown in Table 6.1. Whilst the input and output μ_θ and σ_θ, do not match exactly, the values are close and increasing the input μ_θ and σ_θ also increases the output μ_θ and σ_θ. Additionally, the output density generally matches the input target density well, achieving higher densities than MEDUSA for the same angular dispersion.

These phantoms took an average of 6 hours to grow plus an average of 20 minutes for the meshing and microstructural measurement procedure, using 9.4GB of RAM on average. These values give an estimate of the time taken to generate a typical ConFiG phantom, though it is strongly dependent on user inputs (number of nodes in the network etc.).

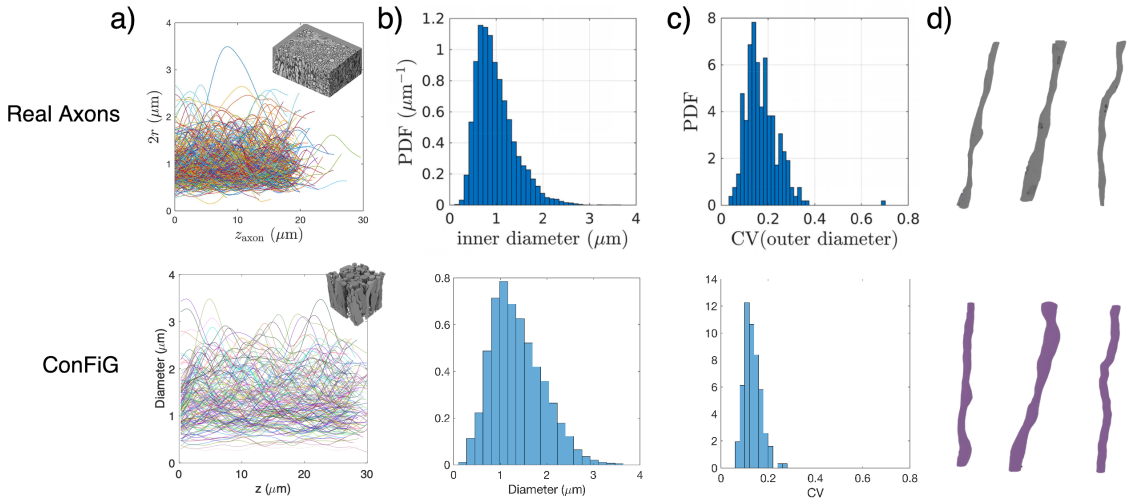

Figure 6.7: a) Along fibre diameter variation in ex vivo mouse corpus callosum, reproduced from Lee et al. [28] compared to along axis diameter variation in the phantom inset demonstrating the ability of ConFiG to generate realistic microstructure. b) Histograms of the inner diameter of axons from Lee et al. [28] and diameter of ConFiG axons. c) Coefficient of variation along axons for real and ConFiG axons and d) Three example fibres reconstructed from the EM data used by Lee et al. to make a). d) Three example ConFiG fibres selected for similarity to the EM examples

6.4.3 Packing induced microstructural complexity

The more complex the fibre orientation in a phantom, the more complex the microstructure that is generated will be as demonstrated in Figure 6.9. Both the mean tortuosity and mean coefficient of variation grow as more orientation dispersion is introduced to the phantom (Figure 6.9a&b), with the two showing a positive correlation (Figure 6.9).

6.4.4 Variability of repeated instances of ConFiG

Different random seeds in the ConFiG phantom generation algorithm lead to different microstructural arrangements as shown in Figure 6.10. In particular, the amount of beading introduced can vary quite a lot relative to the mean amount of beading, with certain instantiations of the more dispersed phantoms ($\kappa = 2$) showing a CV more typical of the highly coherent phantoms ($\kappa = 100$).

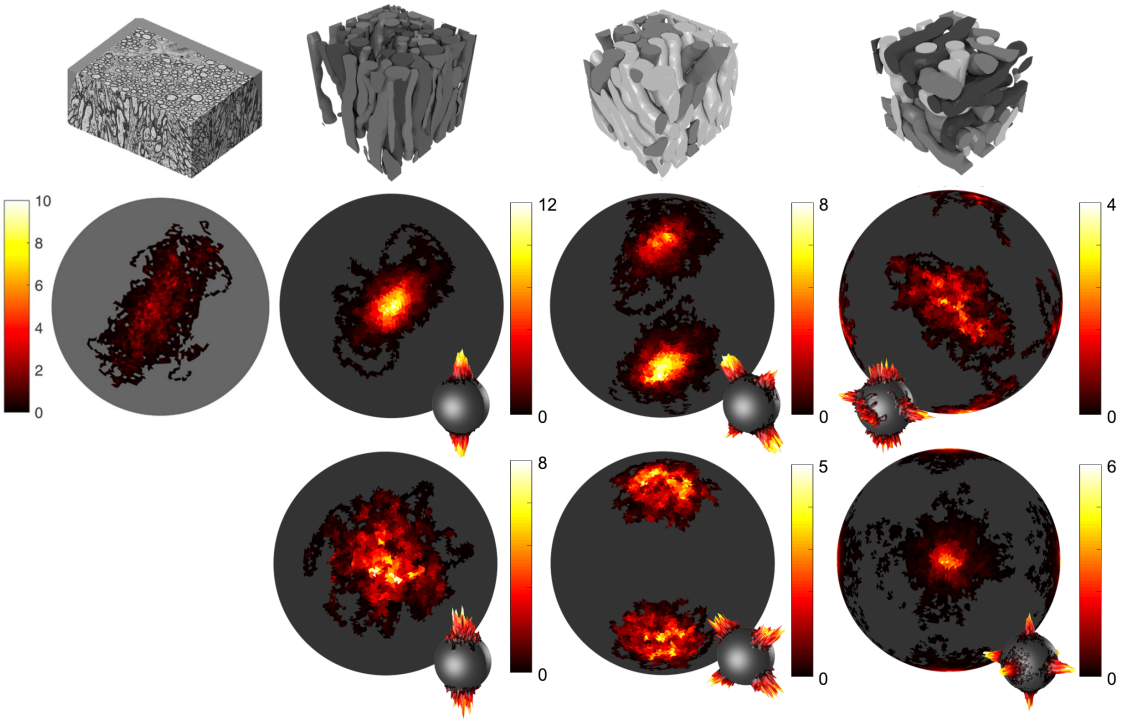

Figure 6.8: Fibre orientation distributions for EM data and a series of numerical phantoms. Top row: EM data used to generate fODF, reproduced from Lee et al. [28] and three ConFiG phantoms with one, two and three crossing bundles (each crossing bundle coloured a different shade of grey). Middle row: fODF profile for real EM data and fODF profiles corresponding to ConFiG phantoms above, generated using an elliptically symmetric dispersion. Bottom row: Three fODF profiles generated from ConFiG phantoms generated using isotropic orientation dispersion. Colourmap has units of sr^{-1}.

Table 6.1: Comparison between input microstructural parameters and the microstructure measured in the resulting ConFiG phantoms. For each phantom, an input target density, ρ, of 75% was used with each phantom having a different value of κ used in the Watson distribution. Each κ is associated with a target μ_θ and σ_θ, the mean and standard deviation of the angle away from the main bundle direction. Angles reported in degrees.

Input κ	Input ρ	Output ρ	Target μ_θ	Output μ_θ	Target σ_θ	Output σ_θ
8	75%	70.6%	19.60	17.46	11.32	9.92
10	75%	73.4%	17.11	16.47	9.62	9.43
15	75%	73.4%	13.60	13.93	7.37	8.75
20	75%	70.7%	11.68	12.69	6.23	7.88
30	75%	72.0%	9.45	11.60	5.02	6.60
50	75%	73.6%	7.26	9.36	3.83	5.51
100	75%	74.9%	5.10	7.75	2.68	4.23

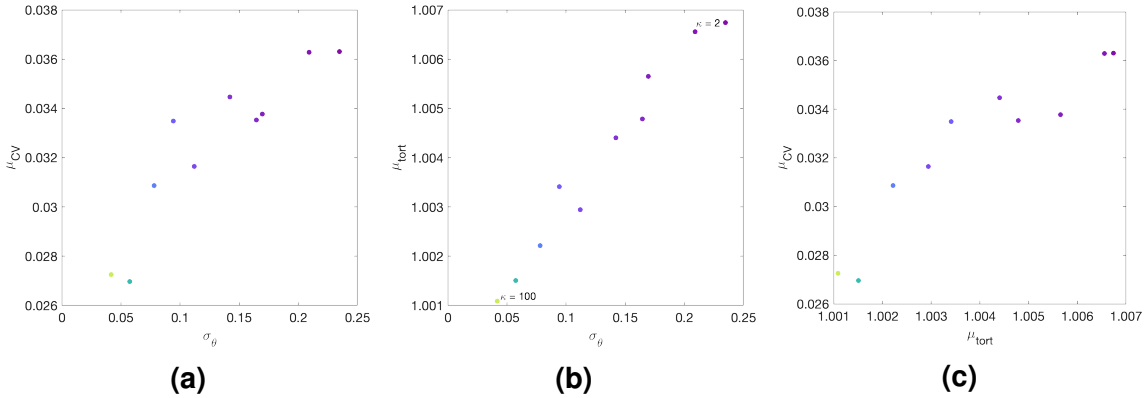

Figure 6.9: Packing induced microstructural complexity. (a) the mean coefficient of variation in a phantom (μ_{CV}) grows with increased orientation dispersion in the phantom, as does (b) the mean tortuosity (μ_{tort}). (c) μ_{tort} and μ_{CV} are positively correlated showing that for ConFiG phantoms, complex arrangements of fibres are achieved through both undulation and beading of fibres.

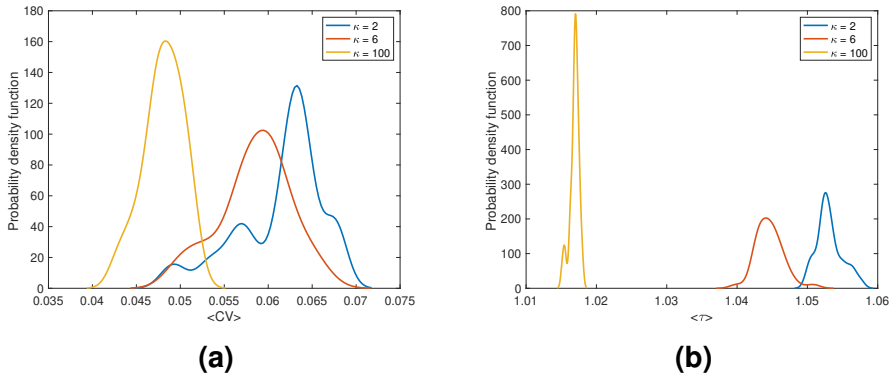

Figure 6.10: Variability of ConFiG phantoms under repetition. (a) Variability of the coefficient of variability of diameter of the axons which is a measure of beading. (b) Variability of the tortuosity of axons which is a measure of undulation.

6.5 Discussion

ConFiG is shown to produce realistic WM numerical phantoms, capturing microscopic structural features such as diameter and orientation distributions. The amount of real data containing 3D microstructural morphology information available to compare to is limited, so we have only compared to one sample in this study. Whilst limited, this shows that ConFiG is able to produce realistic microstructure by following simple biologically inspired growth rules.

Figures 6.5 and 6.6 demonstrate that ConFiG is able to create fibre morphologies that match real axons much more closely than previous methods based on cylinders. Whilst some of the features such as eccentricity may be achievable with cylinders oriented obliquely to the cutting plane, ConFiG phantoms capture morphological features that are otherwise impossible with cylinders such as convexity less than one.

Whilst the input morphological priors do not necessarily correspond to the morphology of the resulting ConFiG phantom, Table 6.1 shows that even for relatively high orientation dispersion and density, this effect is small.

Related to this effect, Figure 6.9 shows that as the input fibre orientation becomes more complex, with increased OD, the microstructural complexity increases. Both the tortuosity in the fibre paths and the coefficient of variation in diameter along fibres grows with increased OD in the phantom. Furthermore, the fact that ConFiG relies on a pseudorandom network of points leads to variation in the generated microstructure, even when using the same input to the algorithm as shown in Figure 6.10.

These factors mean that ConFiG should not be seen as capable of producing exactly the desired inputs but instead as a tool which can generate realistic phantoms as close to the target morphology as currently possible. Even so, for use in further analyses microstructural measures such as orientation dispersion and density should be calculated based on the resultant phantom, rather than taking the input microstructural parameters.

6.5.1 Limitations and future Work

As mentioned above, this study only compares ConFiG to one EM sample of real tissue. Future work will also aim at more extensive validation of the digital phantoms generated using ConFiG, making comparison with larger EM datasets, including different WM configurations from different brain regions.

One limitation of this and future work is ambiguity in the definition of the fODFs which are displayed in Figure 6.8. In this work, since we're concerned with the actual microstructure of the phantoms we calculate the fODF with a small amount of smoothing to compare with Lee et. al [28]. One consequence of this 'microscopic' definition of the fODF is that individual fibres contribute to more than one direction in the distribution (this can be seen in Figure 6.8 with the loops on the edges of the distributions which come from individual fibres). This is contrary to the fODF more commonly used in dMRI techniques such as tractography [70, 72, 75, 76] which assumes each fibre contributes to a single direction only. It remains to be seen how these two definitions of the fODF can be brought together, thought it would certainly be a valuable contribution to be able to test the performance of fODF estimation techniques to the 'true' microscopic fODF.

We will work towards decreasing the difference between the input and output morphological measures, particularly in complex situations, such as high orientation dispersion and crossing bundles. This can be addressed through the improvements to ConFiG mentioned in Section 5.6 and also by improving the strategy for the generation of starting and target points for each fibre. For instance, currently it is not intuitive how starting and target points should be arranged to achieve a desired density in crossing regions of fibres.

6.6 Conclusion

The experiments presented here demonstrate the microscopic realism of ConFiG phantoms. Microstructural measurement methods have been developed which enable the quantification of axonal microstructure from ConFiG meshes and are used to demonstrate that ConFiG captures axonal diameter and orientation distributions that agree well with real axonal microstructure. Not only that, but ConFiG captures subtle, complex features such as bulges and non-circular cross-sections bringing WM phantoms much closer to real tissue than the previous standard of cylinders. Whilst only comparing to a single tissue sample, these experiments give us confidence that the microstructure generated using ConFiG is a big step forwards towards realistic WM numerical phantoms.

Part III

Applications

7 Impact of realistic axonal geometry on intracellular diffusion

Chapter Contents

7.1 Introduction 103
7.2 Method 105
7.3 Results 109
7.4 Discussion 114
7.5 Conclusion 117

Chapter Summary

In this chapter we explore how realistic axonal morphology impacts the intracellular diffusion within axons, from the pure diffusion dynamics to models used to estimate the diffusivity from dMRI data.

7.1 Introduction

The diffusivity is an intrinsic property of a fluid describing how quickly the diffusion process happens, and as such is very important for diffusion MRI (dMRI) techniques since our measurements are intrinsically dependent on the diffusivity. As such, virtually all dMRI models have some parameter which is used to capture the diffusivity. The diffusion tensor, for instance, is effectively a 3D descriptor of the diffusivity in different directions with the apparent diffusion coefficient (ADC) being the apparent overall diffusivity that we can measure from our data.

Other compartment models assign a diffusivity to various tissue compartments and either assume that they have some fixed value, or try to estimate them directly from the data. An example of this is neurite orientation dispersion and density imaging (NODDI) [13], in which the axons are treated as zero-radius 'sticks' meaning they only have diffusivity along the axons which is assumed to be $1.7\,\mu\mathrm{m}^2\,\mathrm{ms}^{-1}$. The extracellular space has a diffusivity based on the estimated axonal volume fraction under the assumption that more densely packed fibres leads to a more hindered diffusion perpendicular to the fibres.

This representation of axons as straight sticks or cylinders is a common one in the white matter (WM) imaging community [13, 62, 63, 155–158] since it leads to a simple relationship between the microstructure (the sticks) and the dMRI via the diffusivity. In the simplest case of diffusion in straight, parallel sticks, the dMRI signal is

$$S = S_0 \exp(-b d_\parallel (\vec{g} \cdot \vec{n})^2)\,, \tag{7.1}$$

where b is the b-value, S_0 the signal without diffusion weighting, $\vec{g}$ the gradient direction, $\vec{n}$ the stick orientation direction and $d_\parallel$ is the axial diffusivity of the sticks (that is, diffusion along the sticks). In this model, diffusion is Gaussian along the axon and zero radially.

In these models the axons are represented as straight sticks or cylinders, however *ex vivo* studies have revealed that real axonal morphology is much more complex than that, including undulations and variable diameters along axons [27, 28, 106, 159]. One thing that *ex vivo* histology studies such as these inherently cannot measure however, is the diffusivity, and how it interacts with the complex microstructure.

It is important that we are able to accurately estimate (or assign an accurate value to) the diffusivity since misestimation (or poor assignment) will bias the other parameters that our model is trying to estimate (for instance axon diameter, or orientation dispersion (OD)). A few studies have attempted to estimate the axonal axial diffusivity from dMRI data, typically getting values around $1.9\text{-}2.5\,\mu\mathrm{m}^2\,\mathrm{ms}^{-1}$ [155, 160–162]. A recent study, attempting to estimate $d_\parallel$ while accounting for OD, however, estimates higher values, up to $3\,\mu\mathrm{m}^2\,\mathrm{ms}^{-1}$ [163].

In this chapter, we investigate how complex microstructural features interact with the intracellular diffusivity, initially in simple geometric environments of undulation and beading separately. Previous works have investigated the impact of undulation [54, 122] and beading [58] but here we look into whether such simple representations can explain the intracellular diffusivity in complex environments introduced by the dense packing of axons in realistic contextual fibre growth (ConFiG) phantoms and axons reconstructed from electron microscopy (EM) [28]. We further investigate how the complex microstructure in realistic axons affects the estimation of the intracellular diffusivity using a model based on Equation (7.1) with OD.

7.2 Method

7.2.1 Impact of beading and undulation

To investigate how microstructural morphology affects the intracellular diffusivity, we began by investigating the simplistic cases of beading (variable diameter along axons) and harmonic undulation. A set of fibres were generated with varying amounts of undulation and beading. In each case, we created 300 fibres, each $1000\,\mu\mathrm{m}$ long, broken down into $1\,\mu\mathrm{m}$ sections.

For the undulating fibres, each fibre had a diameter drawn from a gamma distribution with mean $D_0 = 1\,\mu\mathrm{m}$ and standard deviation $\sigma = 0.2\,\mu\mathrm{m}$ and variable undulation amplitude, $a = 2, 4, 8\,\mu\mathrm{m}$ and wavelength $\lambda = 10, 50, 100\,\mu\mathrm{m}$.

For the beading fibres, each fibre had a base diameter, D_0, gamma distributed with mean $D_0 = 1\,\mu\mathrm{m}$ and $\sigma_0 = 0.2\,\mu\mathrm{m}$. Each $1\,\mu\mathrm{m}$ long fibre section was then given a diameter D drawn from a gamma distribution with mean D_0 and $\sigma = 0.5D_0, 1D_0, 2D_0$. Each fibre's diameter profile was smoothed with a $15\,\mu\mathrm{m}$ wide moving mean to ensure a smooth diameter profile along each fibre. Examples of undulating and beading fibres can be seen in Figure 7.1.

As a summary metric of the beading of each set of fibres the mean coefficient of variation of diameter of each fibre was calculated. That is

$$\langle \mathrm{CV} \rangle = \left\langle \frac{\sigma_D}{\mu_D} \right\rangle , \tag{7.2}$$

where σ_D is the standard deviation of the diameter and μ_D the mean diameter for each fibre and $\langle \cdot \rangle$ represents averaging across fibres in a set. As a summary metric of undulation, the mean microscopic orientation dispersion, $\mu\mathrm{OD}$, across fibres was calculated as in Brabec et al. [122]. For a single fibre, the $\mu\mathrm{OD}$ is

$$\mu\mathrm{OD} = \left\langle \sin^2 \theta \right\rangle , \tag{7.3}$$

where $\theta = \tan^{-1}\left(l_\perp / l_\parallel\right)$, with $l_\parallel$ being the length of a fibre segment along the main axis and $l_\perp$ the length in the radial direction. This means a fibre segment perfectly aligned with the main axis has $\theta = 0 \rightarrow \mu\mathrm{OD} = 0$ while a segment perfectly perpendicular to the axis has $\theta = \pi/2 \rightarrow \mu\mathrm{OD} = 1$. The mean $\mu\mathrm{OD}$ across all the fibres in the set is used as a summary metric of the undulation for that set.

For comparison, a set of straight cylinders were generated with diameters drawn from the same gamma distribution with mean $D_0 = 1\,\mu\mathrm{m}$ and standard deviation $\sigma = 0.2\,\mu\mathrm{m}$.

The diffusion of $500,000$ spins distributed proportionally to the volume of each fibre was simulated using Camino [25, 89]. From the spins' trajectories, the axial diffusivity $d_\parallel$ was computed for $t = [0 : 0.1 : 500]\mathrm{ms}$ from the mean squared displacement (MSD) along the axon as

$$d_\parallel = \frac{\mathrm{MSD}}{2t} . \tag{7.4}$$

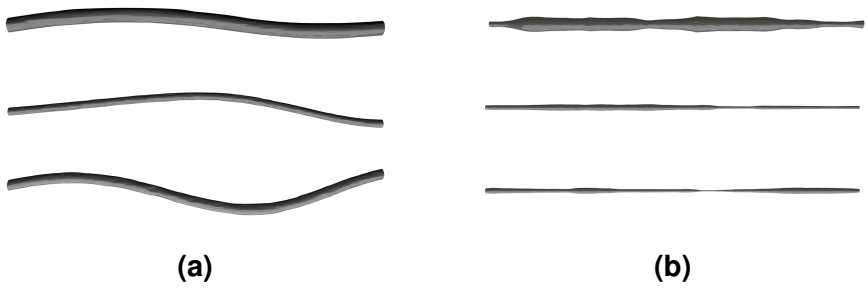

Figure 7.1: Example fibres used to investigate impact of undulation (a) and beading (b).

7.2.2 A phenomenological model of undulation and beading

In an attempt to explore the impact of beading and undulation on the axial diffusivity simulated in Section 7.2.1, we created a phenomenological model to describe the time-dependent diffusivity. As discussed in Sections 7.3 and 7.4, the results of Experiment 7.2.4 show that undulation leads to a long diffusion time dependence in the axial diffusivity while undulation leads to largely time independent shift in the axial diffusivity (Figure 7.3).

Therefore, to describe the impact of beading and undulation on the diffusivity, we split the diffusion time into two regimes, a long diffusion time ($t \geq 25\,\mathrm{ms}$) regime which informs us about beading and a short diffusion time ($4 < t < 20\,\mathrm{ms}$) regime which informs us about undulation.

A relationship was estimated from the data between the slope, c, of a linear fit to the long diffusion time $d_{\parallel}$ and the amount of beading of axons (measured by $\langle \mathrm{CV} \rangle$ as in Equation (7.2)) and between the mean $d_{\parallel}$ at short diffusion time, $d_{\parallel}^{\mathrm{short}}$ (Equation (7.5)), and the undulation (measured by mean $\mathrm{microOD}$).

$$d_{\parallel}^{\mathrm{short}} = \langle d_{\parallel} \rangle \quad \text{for } \; 4 < t < 20\,\mathrm{ms} \tag{7.5}$$

The effectiveness of this model was tested on the $d_{\parallel}$ measured from realistic ConFiG phantoms and EM fibres by comparing the $\langle \mu\mathrm{OD} \rangle$ and $\langle \mathrm{CV} \rangle$ estimated from $d_{\parallel}$ in the phantoms to that measured from the microstructure using Equations (7.2) and (7.3).

7.2.3 Realistic axonal morphology

In order to test how the complex morphology introduced by packing fibres together around one another impacts on the intracellular diffusivity, a set of six ConFiG were generated. Three phantoms were generated with target OD drawn from a Watson distribution [115] with $\kappa = 2, 6, 100$ and constant target radius to test the impact of OD driven fibre microstructure. One phantom was generated with $\kappa = 5$ and variable target diameter with standard deviation $\sigma_D = 1D_0$ in the same way as in Section 7.2.1 to test how packed fibres with beading affect the diffusivity. Phantoms were generated with two crossing and three crossing perpendicular bundles of fibres to assess the impact of complex morphology introduced by interleaved crossing bundles of fibres.

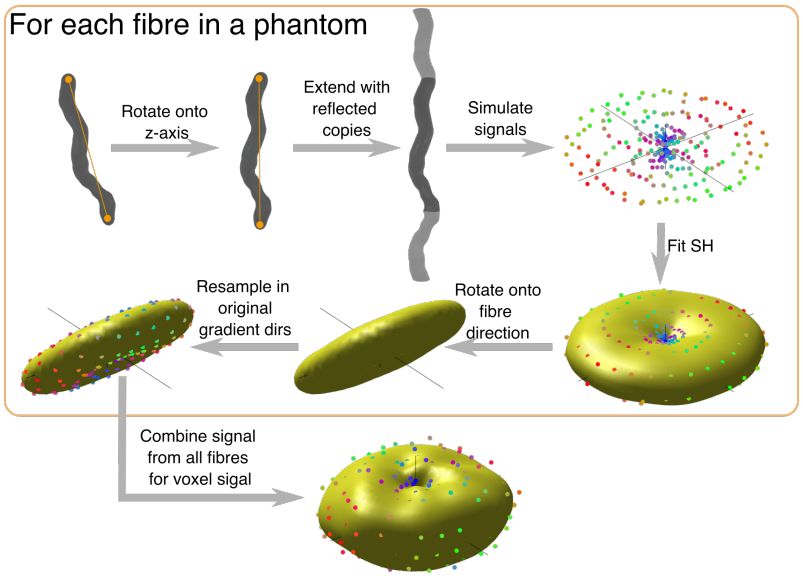

Figure 7.2: The method used to generate simulated intracellular signals. Firstly, the fibres are aligned with the z-axis and then extended with replicated copies before the signals are simulated using Camino. Each fibre's signal is then rotated back onto its original direction using the SH representation and all fibre signals combined to give the overall voxel signal.

Additionally, a set of fibres meshed from the segmentations of 3D EM of mouse corpus callosum (CC) in Lee et al. [28] were used to assess the impact of realistic geometries. The procedure used to mesh these axons used the `isosurface` function in MATLAB to create a surface from a binary volume segmentation of axons and is described in more detail in Section 8.2.1.2.

Since this experiment was only interested in the intracellular diffusivity and the b-values used were very high, the extracellular signal was assumed to be negligible. This meant that the intracellular signal from each fibre could be simulated individually and so to avoid issues with capping at the end of each fibre mesh, each fibre was aligned with the z-axis and extended with a reflected copy as in Lee et al. [134]. The rotation matrix used to align the fibre with z was stored so that signals could be rotated back into the dispersed directions to generate an overall voxel signal. The process used to simulate the signals used in this experiment is shown in Figure 7.2

To test whether the diffusivity time-dependence seen in simulation is present in real data, dMRI measurements were performed in ex-vivo mouse CC. Pulsed-Gradients-Stimulated-Echo data were measured at 16.4 T (scanner Bruker-BioSpin, Germany) with: $\mathrm{TE/TR} = 46/2500\,\mathrm{ms}$ $\delta = 2.5\,\mathrm{ms}$, $b = 1.5, 3\,\mathrm{ms\,\mu m^{-2}}$, $\Delta = 6.85, 15.85, 25.85, 50.85, 75.85, 100.85\,\mathrm{ms}$, 15 gradient directions, resolution $100\,\mu\mathrm{m} \times 100\,\mu\mathrm{m} \times 250\,\mu\mathrm{m}$, 4 averages. The principle eigenvalues of the diffusion tensor, $\mathrm{ADC}_{\|}$, at each diffusion time, $\Delta - \delta/3$, were computed using FSL (`https://fsl.fmrib.ox.ac.uk/fsl`) and these principle eigenvalues taken to be an estimate of the axial diffusivity.

7.2.4 Impact on estimated $d_{\parallel}$

In order to test how microscopic variations in fibre morphology affect the estimation of $d_{\parallel}$ from the dMRI, we used a modified NODDI model based on recent work from Howard et al. [163] to estimate $d_{\parallel}$ in the presence of OD. Here the assumption is that at a high b-value the extracellular signal has decayed to a negligible level, and the signal can be described by the intracellular term in the NODDI model alone. That is, the dMRI signal in a gradient direction $\vec{g}$ at a given b-value is

$$S \sim S_0 f \cdot \frac{1}{C_W} \int_{S^2} \exp\left(\kappa \left(\vec{\mu} \cdot \vec{x}\right)^2\right) \exp\left(-b d_{\parallel} \left(\vec{g} \cdot \vec{x}\right)^2\right) dS \,, \tag{7.6}$$

where f is the intracellular volume fraction, C_W is the Watson distribution normalisation constant, κ is the Watson concentration parameter, $\vec{\mu}$ is the mean fibre direction and $\vec{x}$ is a unit vector on the sphere which is integrated over. This model is equivalent to the standard NODDI model with ν_{ic} set to one, ν_{iso} set to zero and $d_{\parallel}$ set free.

To test how this model performs and is affected by microscopic fibre morphology, simulations were performed in the intracellular space of the same Watson distributed phantoms as in Section 7.2.3 with each fibre aligned with z. The per-fibre signals were rotated to align with the original fibre direction and the volume-weighted mean signal calculated for each phantom as illustrated in Figure 7.2.

The simulated scheme had $b = 7.75, 11.6, 15.5\,\mathrm{ms\,\mu m^{-2}}$, $\delta = 7\,\mathrm{ms}$, $\Delta = 55\,\mathrm{ms}$ with 256 directions per shell. These b-values and timings were chosen to be in agreement with Howard et al. [163] who use such high b-values to ensure the extracellular contribution is minimal. Here we use 256 directions per shell to ensure a dense sampling of directions so that signals can be rotated into alignment with the original fibre direction effectively.

Additionally, to investigate the effectiveness of the model in a simple, theoretically ideal scenario, the model was tested against simulations in simple cylinders. Cylinders with the same diameter distribution as the ConFiG phantoms had the same dMRI sequence as above simulated in them, with varying levels of dispersion introduced into the overall signal by rotating each signal onto a direction sampled from the Watson distribution. The set of κ values used was $\kappa = 2, 6, 10, 20, 30, 50, 100, \infty$, where ∞ leads to no dispersion (i.e. parallel fibres).

The model in Equation (7.6) was fit to the simulated dMRI using the NODDI toolbox (`https://www.nitrc.org/projects/noddi_toolbox`), with the modifications described above - that is, setting ν_{ic} to one and ν_{iso} to zero. To see how axonal morphology affects estimation of $d_{\parallel}$, firstly the model was fit with $d_{\parallel}$ set as a free parameter and estimated. Additionally, the model was fit with a range a of fixed $d_{\parallel}$ values from $d_{\parallel} = 0.5 : 0.1 : 3\,\mathrm{ms\,\mu m^{-2}}$ to see how the fitting performance varied with diffusivity.

7.3 Results

7.3.1 Impact of beading and undulation

Beading and undulation have different effects on the time-dependence of the intracellular diffusivity as seen in Figure 7.3. Figure 7.3a&c shows that undulation effects the intracellular diffusivity at very short diffusion times ($t < 4\,\mathrm{ms}$, $t^{-0.5} > 0.5$), leading to an almost plateau of $d_\parallel/d_0$ at moderate to long t ($t > 4\,\mathrm{ms}$, $t^{-0.5} < 0.5$).

Beading leads to time-dependence at long diffusion times as seen in Figure 7.3b&d with increased beading leading to a stronger fall in $d_\parallel/d_0$ at long diffusion times ($t > 25\,\mathrm{ms}$, $t^{-0.5} < 0.2$).

7.3.2 Phenomenological model of undulation and beading

A phenomenological model of undulation was created by finding the relationship between the mean level of the $d_\parallel/d_0$ plateau ($4 < t < 25\,\mathrm{ms}$), which we label $d_\parallel^{\mathrm{short}}$, and the microscopic OD which we label $\mu\mathrm{OD}$. This relationship is linear, as seen in Figure 7.3e.

Beading, meanwhile, gives a quadratic relationship between $\langle CV \rangle$ and the slope, c, of the linear fit to the long diffusion time diffusivity as seen in Figure 7.3f.

7.3.3 Realistic axonal morphology

Generally, the axial diffusivity measured in realistic fibres shows a combination of the plateau shift from undulation and time dependence from beading as seen in Figure 7.4.

When using the phenomenological model created in Section 7.3.2, the undulation is not well captured as seen in Figure 7.4b, with all of the $\langle d_\parallel^{\mathrm{short}} \rangle$ falling below the fit line. For beading, the phenomenological model performs better as seen Figure 7.4c, with most of the points lying close to the fit.

The real data from mouse CC shown in Figure 7.4a, also show the time dependence in $\mathrm{ADC}_\parallel$ that is typical of beading. It is difficult however, to quantitatively compare this data with the simulations since here the raw measured $\mathrm{ADC}_\parallel$ is used without being normalised by d_0 since this is unknown in the real data.

7.3.4 Impact on estimated $d_\parallel$

The microscopic variations in fibre morphology introduced when fibres are packed densely around one another lead to a lower estimated $d_\parallel$ than when using cylinders as shown in Figure 7.5.

In Figure 7.5, we see that the higher the κ, the higher the estimated $d_{\parallel}$. This is investigated in cylinders in Figure 7.6, showing that at low κ, $d_{\parallel}$ is estimated well, but as κ increases, the estimated $d_{\parallel}$ increases, deviating from the true value. As seen in Figure 7.6b-e, this could be explained by the fact that at high κ, the objective function becomes flat at $d_{\parallel} > 1.7$, meaning that the model cannot accurately estimate $d_{\parallel}$.

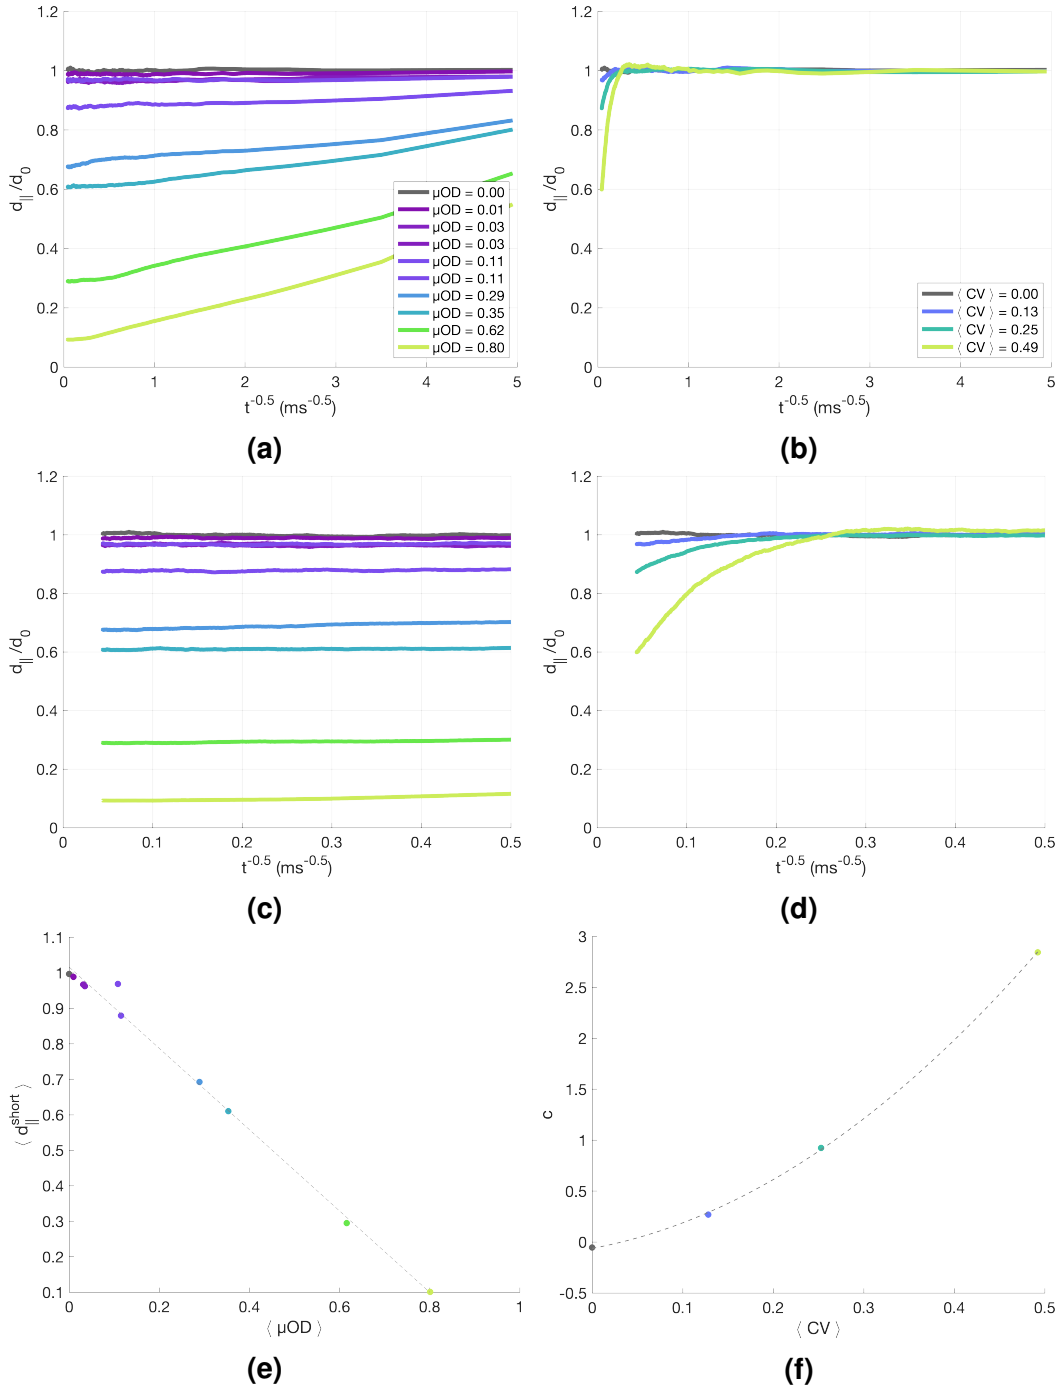

Figure 7.3: Intracellular axial diffusivity in beading and undulation. a) and e) undulation leads to a time-dependence in diffusivity at very short diffusion times and a μOD dependent plateau at long diffusion time. b) and d) beading leads to a long diffusion time time-dependence of $d_{\parallel}/d_0$. e) The phenomenological model of undulation finds a linear relationship between μOD and the diffusivity plateau and f) the phenomenological model of beading finds a quadratic relationship between $\langle CV \rangle$ and the slope, c, of the diffusivity at long diffusion time. Grey dashed lines show phenomenological model.

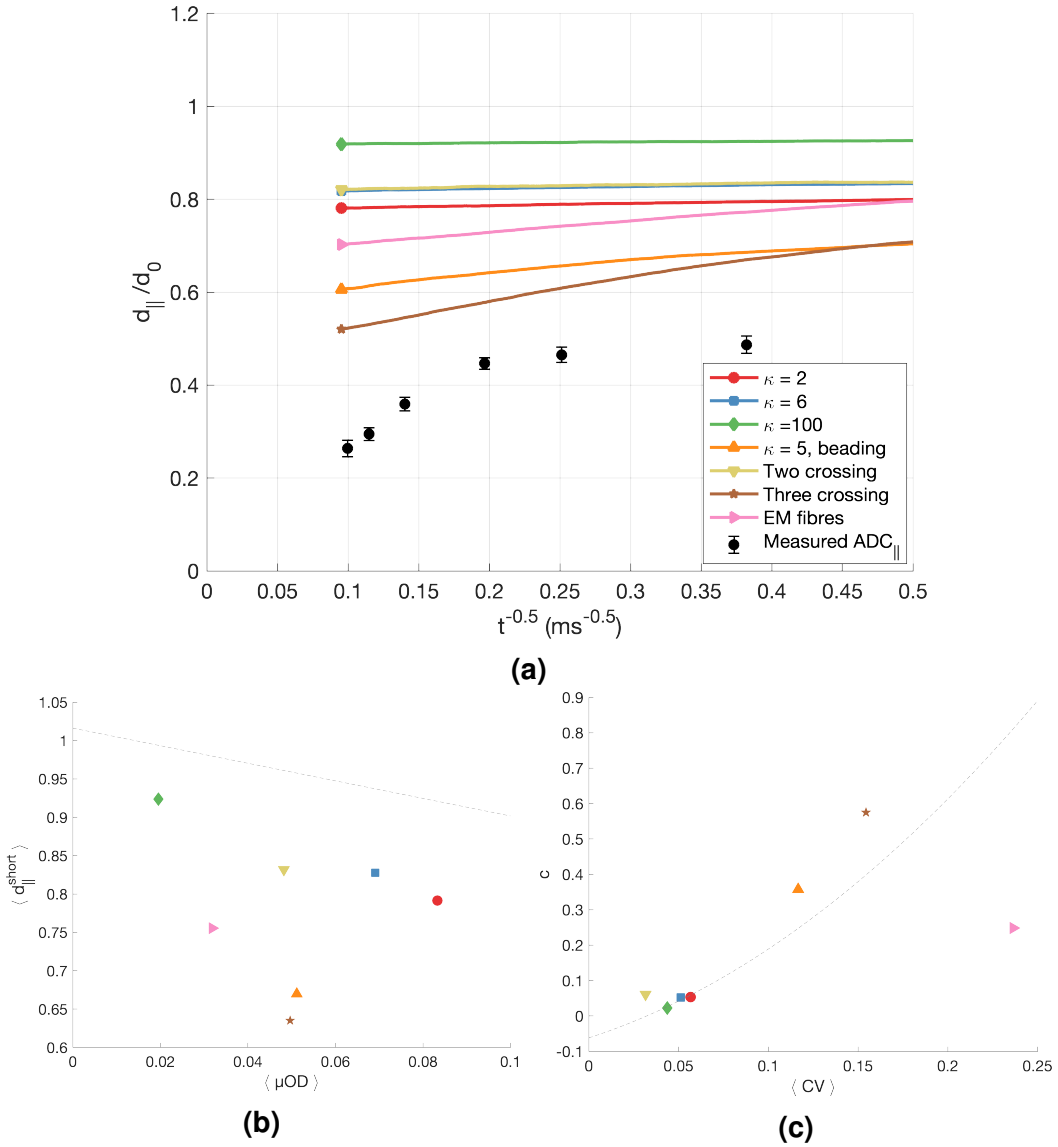

Figure 7.4: Intracellular axial diffusivity in realistic phantoms generally shows a combination of the plateau shift from undulation and long time-dependence from beading. a) $d_{\parallel}/d_0$ for a variety of phantoms (lines) as well as measured $\mathrm{ADC}_{\parallel}$ from mouse corpus callosum. It should be noted that $\mathrm{ADC}_{\parallel}$ is purely the measured valued and not normalised by d_0 (since this is unknown). The phenomenological model for undulation, b), does not match well with realistic phantoms, however the phenomenological model for beading, c), does a better job. Shapes and colours of points correspond to a).

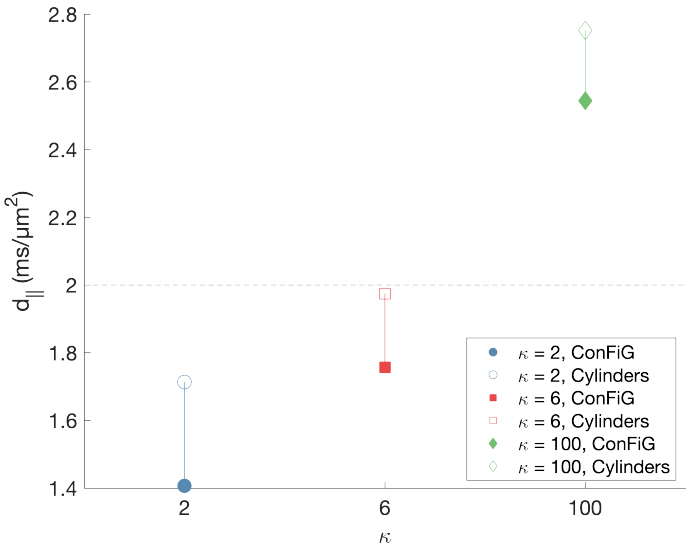

Figure 7.5: Impact of microstructure on estimated $d_{||}$ in realistic phantoms. Estimated $d_{||}$ is consistently lower in ConFiG phantoms than cylinders with the same OD.

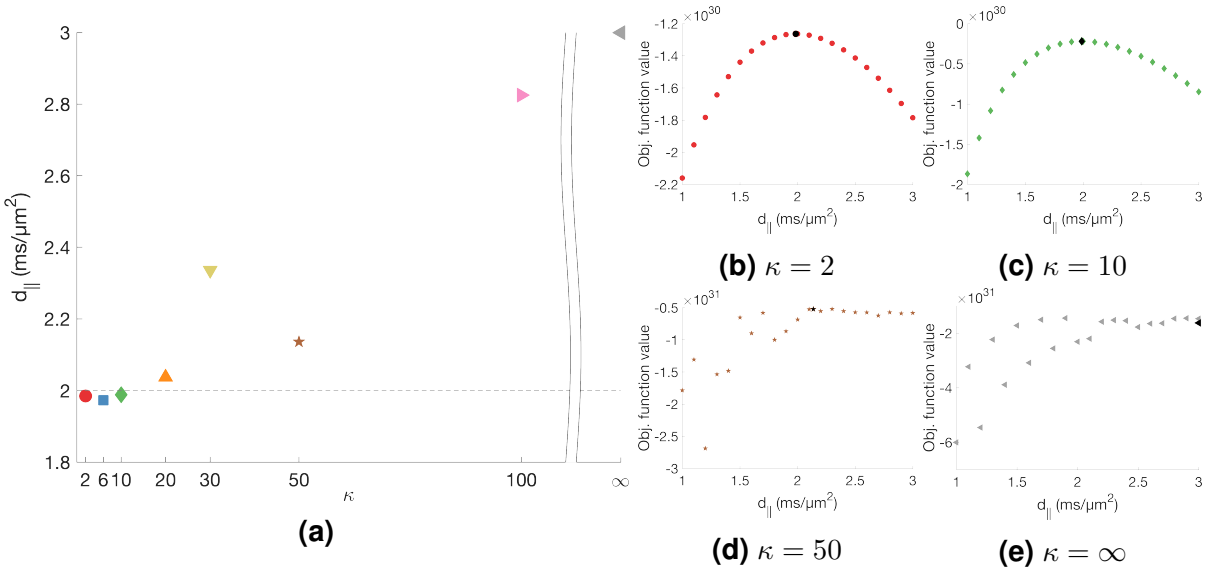

Figure 7.6: Estimated $d_{||}$ in cylinders with varying OD. a) as κ increases (meaning decreasing OD), the estimated $d_{||}$ goes up, straying from the true d_0 set in simulation at $2\,\mu m^2\,ms^{-1}$ (grey dashed line). b) to e) show fitting performance when varying the fixed value of $d_{||}$. At low κ, fitting performance is as expected, being optimal close to the true value of d_0, but at high κ, fitting performance flattens for a range of $d_{||} \geq 2\,\mu m^2\,ms^{-1}$, leading to the misestimation of $d_{||}$ when it is set free. Different symbols represent phantoms with different $kappa$. Black points represent fit with $d_{||}$ set free and estimated.

7.4 Discussion

Impact of beading and undulation Undulation and beading affect the intracellular axial diffusivity in different diffusion time regions as demonstrated in Figure 7.3. This is in agreement with recent results from Lee at al. [134], who perform a similar experiment in axons segmented from EM (in fact, the same real EM fibres used in Section 7.2.3). In their work, they are limited to the EM sample they have, however here we perform a more thorough exploration of the impacts of undulation. In particular, here we more clearly show the effect of undulation which is only present to a small degree in the Lee et al. fibres.

The time-dependent behaviour of $d_{\parallel}$ under undulation can easily be reasoned, since undulation acts to increase the path length that spins must diffuse in the axial direction which leads to a reduced apparent diffusivity. The behaviour under beading is less intuitively understood, however it could arise due to beading behaving as short-term disorder on long diffusion timescales, as suggested by Lee et al. [134, 164].

Phenomenological model of undulation and beading The phenomenological model developed in this work is merely meant as an exploration of the observed pattern in $d_{\parallel}$ to see whether anything meaningful could be potentially gleaned from the data. Were it to be used in practice, a more thorough theoretical model should be developed for linking diffusivity to the microstructural features such as Brabec et al. [122] do for radial diffusivity with undulation and Lee et al. [134] do with short-term disorder to describe beading.

It is, nonetheless, interesting to understand the relationship between undulation and beading and the axial diffusivity, particularly in comparison to realistic axons. Indeed, as shown in Figure 7.4c, although the beading phenomenological model is based on simple beading phantoms, it can be used to describe the axial diffusivity in realistic phantoms fairly well. The undulation model however, fails to represent the realistic phantom diffusivity well (Figure 7.4b), though one possible explanation of this is the relatively low $\langle \mu\mathrm{OD} \rangle$ in the real phantoms when compared to the undulating phantoms. Another factor is that the undulation model is based on harmonic undulation while real axons display more complex undulation patterns.

Realistic axonal morphology Diffusion in realistic numerical phantoms generally shows a combination of undulation-like and beading-like diffusion time dependence as seen in Figure 7.4 which is sensible since each phantom shows some degree of both. Indeed, the long diffusion time behaviour in ConFiG phantoms generally matches well with the phenomenological beading model, however the real EM fibres show a weaker diffusion time dependence given their $\langle \mathrm{CV} \rangle$ which could suggest that the beading model used in building the phenomenological model does not match real beading very well. This also points towards the beading that is produced by ConFiG not matching realistic beading very closely which can be used to inform future development of ConFiG.

While the realistic phantoms do not match the undulation phenomenological model very well, there is a downward trend in $\langle d_{\parallel}^{\text{short}} \rangle$ with increasing $\langle \mu\text{OD} \rangle$ which is what would be expected. From the realistic phantom $d_{\parallel}$ data, one possible explanation for the poor agreement can be seen in that, unlike in Figure 7.3, $d_{\parallel}$ does not generally completely plateau, particularly with beading, in Figure 7.4a. This suggests that $\langle d_{\parallel}^{\text{short}} \rangle$ is not a good value to use since $d_{\parallel}$ is not constant in the time regime in which it is calculated ($4 < t < 25\,\text{ms}$).

The real data from *ex vivo* mouse CC also demonstrate the beading-like time-dependence at long diffusion time, however it is difficult to make a direct comparison with the simulation results since the true d_0 for this case is unknown. It is, however interesting that we see the same kind of time-dependence in this real data that we see in the simple beading phantoms. A confounding factor here is that this is simply the ADC in the axial diffusivity and as such contains contains contributions from non-axonal tissue compartments such as the extracellular space and, potentially, cell bodies. These could be a significant contribution since the b-value used in this experiment is relatively low ($b = 1.5, 3\,\text{ms}\,\mu\text{m}^{-2}$).

Impact on estimated $d_{\parallel}$ Realistic microstructural phantoms lead to a lower estimated $d_{\parallel}$ than when using equivalent cylinders as demonstrated by Figure 7.5. This behaviour should be expected since as shown in Figure 7.4, the apparent axial diffusivity is lower in realistic phantoms than the free diffusivity d_0.

Indeed, the relative drop in estimated axial diffusivity from cylinders to ConFiG, $d_{\parallel}^{\text{ConFiG}}/d_{\parallel}^{\text{cyl}}$, for $\kappa = 2, 6, 100$ is $d_{\parallel}^{\text{ConFiG}}/d_{\parallel}^{\text{cyl}} = 0.82, 0.89, 0.92$ which matches well with the $d_{\parallel}/d_0$ at the diffusion time of the dMRI measurements ($\Delta - \delta/3 = 55 - 7/3 = 52.67\,\text{ms}$) which are $d_{\parallel}/d_0|_{t=52.67} = 0.78, 0.82, 0.92$. This shows that although the actual value of $d_{\parallel}$ estimated in Figure 7.5 is not accurate to the true d_0, the change due to microstructural complexity is captured well.

As mentioned, the estimated $d_{\parallel}$ is not accurate and is dependent on the level of orientation dispersion which was investigated in further detail in Figure 7.6. At low κ, i.e. high dispersion, the estimated $d_{\parallel}$ in dispersed cylinders is accurate and the fitting performance while changing the fixed value of $d_{\parallel}$ behaves as expected, with optimal fitting performance around the true value. At high κ, meaning low OD, however, the $d_{\parallel}$ estimated is inaccurate overestimating $d_{\parallel}$ up to the extreme of $3\,\mu\text{m}^2\,\text{ms}^{-1}$ (which is the upper limit constraint set in the fitting procedure) for parallel fibres. This is also explained by the fitting performance at varying $d_{\parallel}$, which becomes flat after $d_{\parallel} \sim 1.7\,\mu\text{m}^2\,\text{ms}^{-1}$ meaning that the fitting procedure cannot differentiate different $d_{\parallel}$ values very well.

One possible reason for this may be that at the high b-values used here, the signal parallel to fibres has decayed almost entirely to zero for a low level of dispersion in these simulations with no added noise, meaning that many values of $d_{\parallel}$ explain the data equally well since they are all almost zero. To test whether this effect may be present in real data, Rician noise was added to the simulations for a signal-to-noise ratio (SNR) of 30 and the same analysis performed, with the results shown in Figure 7.7.

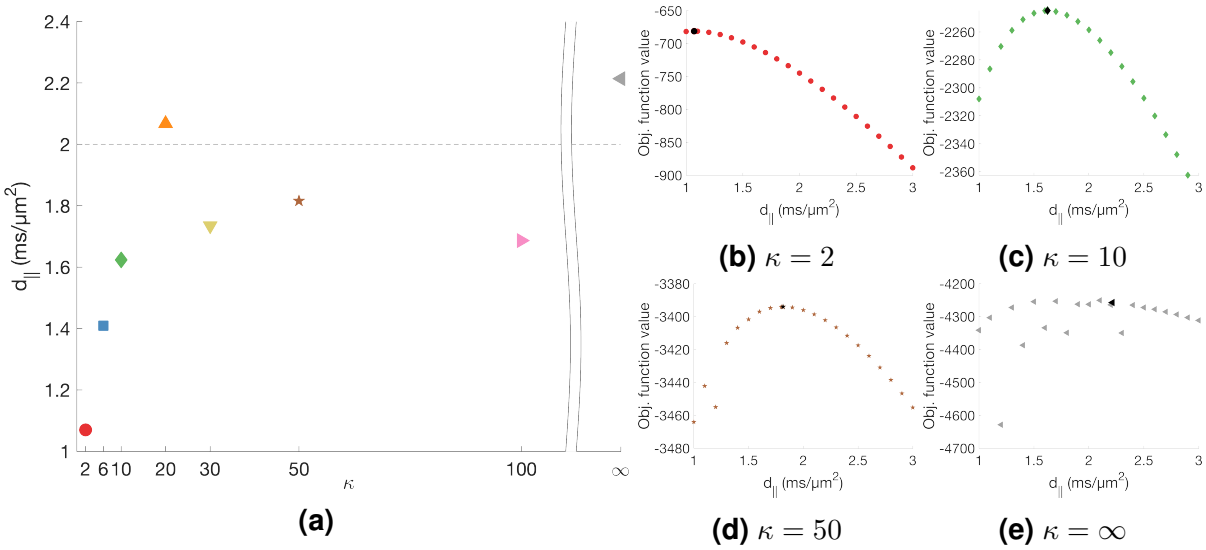

Figure 7.7: Estimated $d_{||}$ in cylinders with varying OD at SNR=30, showing a different behaviour than at SNR=∞ in Figure 7.6.

The estimation of $d_{||}$ behaves differently in the presence of noise, with few phantoms yielding a good estimation of $d_{||}$. For most values of κ, $d_{||}$ is underestimated, though the fitting curves shown in Figure 7.7b-e are generally quite stable, peaking at the estimated free value of $d_{||}$ except for with parallel fibres where the fitting curve is again quite flat. One explanation for this behaviour, linked to the overestimation in the $\mathrm{SNR} = \infty$ case is that here the Rician noise floor masks the full decay of the signal and since the b-value is already so high, all the data points along the axial direction have hit the noise floor and the diffusivity can't be accurately estimated.

This suggests that perhaps the model proposed here cannot be used to accurately estimate $d_{||}$ owing to the vary high b-values used and perhaps can help to understand why the $d_{||}$ estimated by Howard et al. [163] differs from typically reported values.

7.4.1 Limitations and future work

In this work we have identified signatures of undulation and beading in the time-dependence of the axial diffusivity. One potential limitation of this study is that the patterns of undulation and beading are generated using simple parametric forms which may not represent the patterns of undulation and beading that are seen in realistic axons. A potential way to address this is to investigate these patterns in real data to understand what realistic undulation and beading look like, however the amount of high-resolution 3D imaging of WM is limited, both in terms of the number of samples available and the size of the samples.

Another feature which has not been investigate here is the effect of orientation dispersion. It has been shown that OD will affect estimates of $d_{||}$ [107, 134]so future work will look to explore this effect.

The simple phenomenological model presented in Section 7.2.2 is used to explore the effects of undulation and beading, though it is limited. Firstly the model is just a simple way to explain the data we see in Figure 7.3 and as such has no theoretical basis. Future work will look towards building a more complete, theoretically described model. Secondly, the model presented here is based on the measured axial diffusivity normalised by the true free diffusivity used in simulation, $d_{\parallel}/d_0$. This means that the model can't be applied to the real data from mouse CC since the true d_0 is not known in that case. One potential way to address this is to simply use $d_{\parallel}$ to build the model, however the effect of d_0 on this value is not known. For instance, if the same level of beading gives the same high t slope in $d_{\parallel}/d_0$ independently of d_0, then the slope without normalisation will also depend on d_0 and the two cannot be simply disentangled. Future work to attempt to understand this will repeat the simulations in Section 7.2.1 at different values of d_0 to explore its effect.

The effect of microscopic morphological variations on estimated $d_{\parallel}$ using a the model based on Equation (7.6) which attempts to estimate $d_{\parallel}$ under the assumption of diffusion in Watson distributed sticks. Without noise, this model can estimate $d_{\parallel}$ at low κ values (high OD), however it fails at high κ (low OD). When this same model is applied to data with noise at $\mathrm{SNR} = 30$, the model cannot estimate $d_{\parallel}$ accurately, typically underestimating its value. This suggests that this model is not effective for estimating $d_{\parallel}$, at least at these very high b-values and may need to be reassessed.

One potential approach is to attempt the same experiment at lower b-values where the attenuation along fibres will be less strong. The problem with this approach, however is linked to another limitation of this study which is the lack of extracellular space in the simulations. The high b-values used by Howard et al. [163] and in this study are chosen to attenuate the extracellular space strongly, however, so the same model is likely to be invalid at lower b-values where something like the standard NODDI model will be more appropriate. Future work will investigate this effect, simulating a realistic extracellular space along with the intracellular simulations presented here.

Another avenue for investigation is to use a measurement scheme designed for axial diffusivity measurements such as that presented by Dhital et al. [160]. They use planar diffusion encoding to suppress signal orthogonal to the axial direction and then standard linear diffusion encoding in the axial direction to estimate $d_{\parallel}$. It would be an interesting experiment to see how this approach works in the realistic environments created by ConFiG or EM reconstructed fibres.

7.5 Conclusion

Undulation and beading affect the axial diffusivity in fibres in separate diffusion time regimes, offering the potential to disentangle the two effects. We have demonstrated here that realistic axonal morphology, including microscopic variations in fibre shapes, leads to a diffusion time dependent axial diffusivity which shows both undulation-like and beading-like behaviour and demonstrated that a simple phenomenological model can capture some of these features.

Additionally, we have demonstrated that this microscopic complexity of realistic axons impacts on $d_{\parallel}$ estimated using a model which attempts to estimate the axial diffusivity when there is orientation dispersion. This demonstrates that the models typically used to represent axons as sticks do not represent the true geometry fully, which may lead to a bias when estimating $d_{\parallel}$ and, as such, other microstructural parameters.

8 Impact of within-voxel heterogeneity in fibre geometry on spherical deconvolution

Chapter Contents

8.1 Introduction 120

8.2 Method 121

8.3 Results 131

8.4 Discussion 136

8.5 Conclusion 139

Chapter Summary

This chapter presents a series of experiments performed to probe to what extent microscopic variations in fibre geometry impact on dMRI techniques based on spherical deconvolution (SD), using ConFiG phantoms and real fibres reconstructed from EM.

Firstly, the theory behind SD techniques is described, outlining some of the assumptions that could be affected or violated by microscopic fibre heterogeneity. dMRI simulations are performed in ConFiG phantoms and EM fibres to test how realistic, complex axonal morphology impacts SD techniques for fibre orientation distribution function (fODF) estimation.

8.1 Introduction

Diffusion magnetic resonance imaging (dMRI) has been widely used to probe the structure and organisation of brain tissue, with one particular area of focus being the estimation of the orientational distribution of neuronal fibres in a voxel. This fibre orientation distribution function (fODF) is particularly interesting since it is used in tractography techniques to probe the structural connectivity of the brain which is important in many clinical and basic neuroscience studies [72, 165, 166]. Whilst tractography has found many uses, there remain a number of challenges to the technique, including typically generating a large number of false positive connections [73, 77, 79, 167]. One potential source of these issues could be due to difficulties in reliably estimating the fODF, where minor differences in fODF can lead to large differences in the tractograms created [73,76]. Accurate and reliable estimation of the fODF is therefore important to improve the accuracy of tractography techniques.

Many techniques have been developed for estimating the fODF, of which perhaps the most prominent are based on spherical deconvolution (SD). While there are a variety of spherical deconvolution methods, the central principle is the same - the diffusion weighted signal as a function of the azimuthal (ϕ) and elevation (θ) angles is modelled as a spherical convolution of the fODF, $F(\theta, \phi)$, with a kernel (called the fibre response function (FRF)), $R(\theta)$, the typical diffusion weighted signal from a single fibre population estimated *a priori*. By estimating an FRF from voxels where the signals are deemed typical of a single coherent fibre population, the fODF is determined by deconvolving this FRF from the signal. Implicit in this formulation is an assumption that one common FRF is shared across all fibre populations in the white matter (WM). Recently, some works have challenged this assumption, for instance Schilling et al. [76] use known fODFs from histology to estimate the FRF in different WM regions, showing that the FRF does indeed vary across the WM and that this variation does affect the estimated fODF and tractography results.

Indeed, while the FRF is typically taken to represent the typical response from a fibre population on a voxel level, the way that the convolution is defined mathematically requires that it be identical across individual fibre populations within a voxel. In fact, recent works using electron microscopy (EM) and high-resolution x-ray imaging to investigate WM axonal morphology show that axons within a voxel have different shapes [27, 107, 134], with varying diameters along their length and non-straight trajectories. It is reasonable, therefore, to propose that this heterogeneity in fibre geometry could lead to different fibres having different responses. This may lead to misestimation of the fODF when assuming a single FRF for all fibres, which can have downstream consequences for techniques including tractography.

A related factor arising from the convolution is that an assumption is made that there is no exchange between fibres or, equivalently, diffusion in different directions using this representation. In essence, this means that the fibres are implicitly assumed to be perfectly straight and pointing a given direction since any deviation from straight (i.e. curved or undulating fibres) would introduce directions that are connected, violating the non-exchange assumption. Additionally, this means that it is hard to reconcile the effect

of tissue compartments which contribute to multiple directions, such as isotropically restricted compartments or the extracellular space, with this formulation. Under some experimental conditions, such as those used in current clinical applications, these effects may not be negligible and may affect subsequent techniques such as tractography.

In this work we investigate what effect, if any, violation of these assumptions introduced by within-voxel heterogeneity in axonal morphologies has on SD techniques. We use contextual fibre growth (ConFiG) [168], our recently developed white matter numerical phantom generator capable of generating realistic WM morphology to investigate this in controllable environments, as well as real digital tissues reconstructed from EM [28] to test a limited sample of real tissue. Firstly, we investigate how microscopic variations in fibre geometry affect the diffusion within each fibre and whether the dMRI signal from each fibre is the same. We further evaluate what effect this has on fODF estimates by calculating them using FRFs representing the variable responses present in a voxel.

The rest of this paper is organised as follows: Section 8.2 describes the experiments performed to probe the assumptions outlined above, Section 8.3 presents the results and Sections 8.4 and 8.5 summarise the contributions and discuss future work.

8.2 Method

In order to test the impact of fibre geometry heterogeneity on the dMRI signal per-fibre and how any variability in response may affect fODF estimation. Experiments were performed with a range of numerical phantoms generated using ConFiG and reconstructed from EM [28], using an acquisition typical of SD applications.

Two primary experiments were conducted:

1. **Per-fibre response heterogeneity** - To investigate the impact of fibre geometry heterogeneity on the dMRI signal per-fibre
2. **Impact on fODF estimation** – To investigate what impact variation in the FRF can have on fODF estimation

In this section, we begin by describing how the phantoms were generated and gold-standard fODFs extracted and quantified, before describing the dMRI simulation experiments that were performed to investigate the impact of microscopic structural variability.

8.2.1 Phantom Generation

In order to test SD techniques in realistic geometries, a set of digital-tissue phantoms were generated to represent a range of WM tissue configurations:

- A single bundle of fibres generated by ConFiG with varying amounts of orientation dispersion:

  - Watson distributed [115], $\kappa = 2$ - highly dispersed
  - Watson distributed, $\kappa = 6$ - typical WM dispersion
  - Watson distributed, $\kappa = 100$ - highly coherent
- Crossing bundles of fibres generated by ConFiG
  - Two perpendicular bundles
  - Three perpendicular bundles
- Real fibres from mouse corpus callosum (CC) reconstructed from EM [28]

In the case of the single bundle phantoms, a low κ means high orientation dispersion, so phantoms with a lower κ were expected to have more complex morphology since higher OD means that they must grow around one another more to avoid intersections. A typical κ, estimated using NODDI [13], for the corpus callosum of a healthy Human Connectome Project (HCP) [132,137] subject is $\kappa \sim 6$ [168]. Since the CC is expected to contain some of the most coherent fibre bundles in the brain, $\kappa \sim 6$ will be towards the lower end of OD (higher end of κ) *in vivo*. Crossing bundle phantoms were generated by using starting and target points arranged into two- or three-crossing bundles and grown using ConFiG to generate complex phantoms with interleaved fibres.

8.2.1.1 Real WM fibres from EM

To simulate diffusion in real axons, 3D meshes were generated from WM axon segmentations from EM of mouse corpus callosum presented by Lee et al. [28].

The axonal segmentations are provided in the NIfTI format, a volumetric format. In order to convert these into surface meshes for dMRI simulation, the `isosurface` function in MATLAB was used, however this produces meshes with some artifacts such as loose surfaces inside the fibres. In order to account for this, a further mesh refinement procedure was developed using the shrinkwrap feature in Blender to create a smooth, closed surface mesh around each fibre.

8.2.1.2 Gold standard fODF extraction from microstructure

In order to generate a ground truth to evaluate fODFs estimated from the simulated dMRI signals, a gold standard fODF was estimated directly from the WM numerical phantom meshes. As an attempt to generate a microstructural fODF comparable to that estimated from the simulated dMRI based fODFs, the microstructural fODF was calculated using the assumption of one direction per fibre, namely the direction connecting the endpoints of the fibre that would subsequently be used to align each fibre to the z-axis (see Figure 1).

A triangulated unit sphere was used to store this fODF, with each triangle in the sphere storing the number of fibres whose direction went through that triangle scaled by the volume of each fibre, as illustrated in Figure 1. In order to compare this fODF to those calculated using SD from dMRI, the microstructural fODF was expanded in spherical harmonics (SHs). A spherical function $f(\theta, \phi)$, can be expressed in terms of spherical harmonics as:

$$f(\theta,\phi) = \sum_{l=0}^{l_{\max}} \sum_{m=-l}^{l} c_l^m Y_l^m(\theta,\phi)\,, \tag{8.1}$$

where

$$c_l^m = \int_0^{2\pi} \int_0^{\pi} f(\theta,\phi) Y_l^{m*}(\theta,\phi) \sin(\theta) d\theta d\phi\,. \tag{8.2}$$

Y_l^m are the so-called spherical harmonics of degree l and order m up to a maximum degree $l_{\max}$ and $*$ denotes complex conjugation. In our case, θ and ϕ are discrete samples in the centre of each triangle in our unit sphere meaning one approach to finding c_l^m is to turn the integral in into a summation. We adopt a more robust approach based on least squares [71, 169] in which the spherical harmonics are re-indexed to have single index $j(l,m) = l^2 + l + m$. The discrete fODF values stored in each triangle are turned into a vector of length n_{tri}, $[f] = \{f(\theta_i, \phi_i), i = 1, \ldots, n_{\text{tri}}\}$ and an $n_{\text{tri}} \times j(l_{\max}, l_{\max})$ matrix, X, constructed with elements $X_{i,j(l,m)} = Y_l^m(\theta_i, \phi_i)$. Essentially, X maps the SH coefficients for each l, m into amplitudes along each θ_i, ϕ_i. The $j(l_{\max}, l_{\max})$ vector of SH coefficients, $[c]$ can then be found as

$$[c] = (X^{*T} X)^{-1} X^{*T} [f]\,. \tag{8.3}$$

The number of coefficients in $[c]$ can be reduced since the fODF is real-valued and antipodally symmetric. Being real valued means that the SH coefficients exhibit conjugate symmetry (that is, $c_l^m = (-1)^m c_l^{-m*}$) and the antipodal symmetry means that all odd m terms are 0 [70,71]. In the end, this means that $[c]$ has $(l_{\max}+1)(l_{\max}+2)/2$ elements.

Each fODF was normalised such that

$$\int_0^{2\pi} \int_0^{\pi} f(\theta,\phi) \sin(\theta) d\theta d\phi = 1\,, \tag{8.4}$$

to ensure that the fODF is a probability density function (PDF), describing the probability of fibre pointing in a given unit of solid angle.

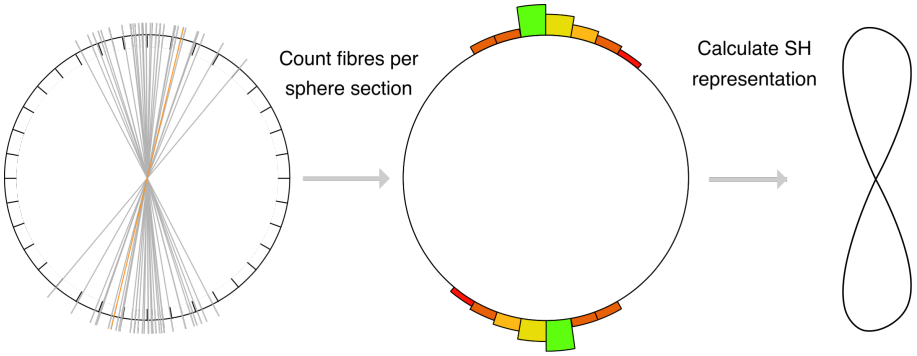

Figure 8.1: Gold standard fODF estimation from microstructure of WM numerical phantoms. Each fibre's main direction (line in left image, extracted as shown in Figure 2) is projected onto a sphere and a spherical harmonic representation calculated

8.2.1.3 fODF metrics

In order to quantify differences in fODF, a series of metrics were extracted by fitting a Watson distribution to each single bundle fODF. The Watson distribution probability density function is defined as a function of a unit vector $\overrightarrow{x}$, parameterised by $\overrightarrow{\mu}$ and κ, the mean direction and concentration parameter respectively.

$$f\left(\overrightarrow{x}|\overrightarrow{\mu},\kappa\right) = C(\kappa)\text{exp}\left(\kappa\left(\overrightarrow{\mu}\cdot\overrightarrow{x}\right)^2\right) , \tag{8.5}$$

where

$$C(\kappa) = M\left(\frac{1}{2},\frac{3}{2},\kappa\right)^{-1} \tag{8.6}$$

is Kummer's confluent hypergeometric function which is a constant for any given κ. If we have a series of measured directions $\{x\} = \{\overrightarrow{x}_1, \overrightarrow{x}_2, \ldots, \overrightarrow{x}_n\}$, we can determine the Watson best-fit parameters with maximum-likelihood estimation. The likelihood function can be computed as

$$\begin{aligned} L\left(\overrightarrow{\mu},\kappa|\left\{x\right\}\right) &= \prod_{i=1}^{n} f\left(\overrightarrow{x}_i|\overrightarrow{\mu},\kappa\right) , \\ &= \prod_{i=1}^{n} C(\kappa)\text{exp}\left(\kappa\left(\overrightarrow{\mu}\cdot\overrightarrow{x}_i\right)^2\right) . \end{aligned} \tag{8.7}$$

In this experiment, $\{x\}$ is our set of measured fibre directions from the phantom, however since each fibre has a different volume, we need to create a weighted likelihood which weights larger fibres more heavily than small fibres. Each fibre should contribute to the overall distribution proportionally to its volume, so we assign each fibre a weight which is its volume, v_i, normalised by the overall fibre volume $V = \sum_i v_i$, i.e. the weight for each fibre is $w_i = {}^{v_i}/_{V}$.

In order to incorporate this into the likelihood, we treat each fibre contributes as if it contributes w_i times so the overall likelihood, so Equation (8.7) becomes

$$L\left(\overrightarrow{\mu}, \kappa | \{x\}\right) = \prod_{i=0}^{n} \left[C(\kappa) \text{exp}\left(\kappa \left(\overrightarrow{\mu} \cdot \overrightarrow{x}_i\right)^2\right)\right]^{w_i}. \tag{8.8}$$

Taking the logarithm and exploiting logarithmic properties, this can be simplified to give the log-likelihood

$$\begin{aligned} l\left(\overrightarrow{\mu},\ \kappa | \{x\}\right) &= \sum_{i=0}^{n} w_i \left[\log\left(C(\kappa)\exp\left(\kappa\left(\overrightarrow{\mu}\cdot\overrightarrow{x_i}\right)^2\right)\right)\right], \\ &= logC(\kappa) + \kappa \sum_{i=0}^{n} w_i \left(\overrightarrow{\mu}\cdot\overrightarrow{x_i}\right)^2. \end{aligned} \tag{8.9}$$

In order maximise the log-likelihood, we follow the approach of Mardia and Jupp [115], to directly estimate $\overrightarrow{\mu}$ using the scatter matrix, $\mathbf{T}(\overrightarrow{w}, \mathbf{x}) = \sum_{i=0}^{n} w_i \overrightarrow{x_i} \bigotimes \overrightarrow{x_i}$, where $\bigotimes$ denotes the outer product. The maximum likelihood estimate of $\overrightarrow{\mu}$ is then found by taking the eigenvector of $\mathbf{T}$ with the largest eigenvalue [115]. κ is then estimated with an iterative grid search evaluating Equation (8.9) for a range of possible $\kappa \in [0, 100]$.

Additionally, this approach is used to quantify dMRI derived fODFs by sampling uniform directions on the sphere and taking the weight, w_i, for each direction to be the normalised fODF amplitude in that direction.

8.2.2 Experiments

8.2.2.1 Per-fibre response heterogeneity (Experiment 1)

We test to what extent variable fibre geometry results in variable fibre responses by simulating the intracellular dMRI signal from the phantoms described in Section 8.2.1 following the procedure outlined in Figure 8.1.

As mentioned in Section 8.1, SD assumes that the overall voxel signal is the sum of signals from the constituent fibre directions. This requires no exchange between different fibre directions which makes the extracellular space difficult to account for. Indeed, in practice the signal is often assumed to come from the intra-axonal space alone to aid in interpretation of the resulting fODF [74,170]. Therefore, to test SD under the assumptions inherent in the model, only the intra-axonal signal was simulated.

To create a periodic intra-axonal space for simulation, each fibre was rotated to be aligned with the z-axis and then extended with a reflected copy as in [171]. The rotation matrix used to align the fibre with z was stored so that signals could be rotated back into the dispersed directions to generate an overall voxel signal. In the case of EM fibres, only axons that were longer than $18\,\mu\text{m}$ (chosen to be slightly shorter than the height of the EM volume in the principle fibre direction) before extension were used in simulation to remove very short fibres that had been segmented at the edge of the volume.

Diffusion MRI signals were simulated from ConFiG phantoms using the Camino dMRI simulator [25, 89] to perform the experiments described below. For all experiments a bulk diffusivity $D = 2.0\,\mu\mathrm{m}^2\,\mathrm{ms}^{-1}$ was used in accordance with similar Monte Carlo experiments [5, 25, 160, 171]. Standard Camino periodic boundaries were used in which copies of the phantom are tiled in the x-, y- and z-directions to create an effectively infinite, but periodic, simulation domain [59].

For each fibre 10,000 spins were initialised uniformly within the intra-axonal space and the simulations were performed using 5000 timesteps. Each phantom had ~ 300 fibres giving $\sim 3 \times 10^6$ spins in total per phantom. These settings were confirmed to be adequate by comparing to a set of test simulations performed using 10^5 spins per fibre and 10^4 timesteps. The measurement parameters were $\Delta = 28\,\mathrm{ms}$, $\delta = 24\,\mathrm{ms}$, $b = 1000, 2000, 3000\,\mathrm{s\,ms}^{-1}$ and 256 equidistributed gradient directions [172] at each shell. This gives a diffusion time $d_t = 20\,\mathrm{ms}$, chosen so that the diffusion length scale ($\sqrt{2Dd_t} \approx 8\,\mu\mathrm{m}$ at $D = 2.0\,\mu\mathrm{m}^2\,\mathrm{ms}^{-1}$ is small relative to the axon length ($\geq 18\,\mu\mathrm{m}$). Additionally, these settings give $G = 60\,\mathrm{mT\,m}^{-1}$ at $b = 3000\,\mathrm{s\,mm}^{-2}$, a feasible gradient strength on a high-end clinical system.

To compare to the collection-of-straight-fibres assumption implicit in SD techniques, an infinite cylinder representing each fibre was generated using the endpoints of each fibre to give the direction and the mean radius of the fibre as the cylinder radius. For simulation, each cylinder was aligned with the z-axis similarly to the ConFiG fibres so that everything was in the same space to compare the signals. The same measurement scheme was simulated in each cylinder in order to compare to the ConFiG fibres.

Since the individual axons have been aligned with the z-axis, the signals from each fibre can be directly compared with one another as the gradient directions are aligned with respect to each fibre. To demonstrate the variability in dMRI response, the median, 10th and 90th percentile responses were found (in terms of mean squared difference between fibre and cylinder responses). Additionally, to relate the type and size of morphological variation to the signal changes, the microscopic orientation dispersion [122] (μOD), a measure of undulation, and the coefficient of variation of diameter (CV), a measure of beading [171], was calculated for each fibre.

Throughout this work, SH representations of signals are used. The MATLAB implementation of constrained spherical deconvolution (CSD) [75] available from (`https://github.com/jdtournier/csd`) is used to calculate SH decompositions of signals. This is the same technique as used in popular dMRI tractography tool MRtrix3 [136] and follows the procedure outlined in Section 8.2.1.2 for SH decomposition of the dMRI signal.

8.2.2.2 Impact on fODF estimation (Experiment 2)

To investigate the impact on the fODF of assuming a single fibre response per voxel, we compared fODF estimates from CSD using six different FRFs per phantom derived from Experiment 1: the 10th, 50th and 90th percentile FRF (representing the spread of responses), the volume-weighted mean signal (representing a best case CSD scenario), a cylinder (representing an ideal stick-like response) and the $\kappa = 100$ voxel signal (representing a typical CSD-like approach of estimating the FRF in a region of coherent fibres). These FRFs are illustrated in Figure 8.2.

For each FRF, the fODF was estimated in two scenarios per phantom, illustrated in Figure 8.3:

1. Full voxel signal: The voxel signal contains each fibre's signal as simulated. This represents the realistic scenario where each fibre may have a different signal.
2. Single FRF signal: The voxel signal comes from using the same FRF signal per fibre. This represents the ideal CSD scenario in which the signal truly comes from a single FRF.

The approach to generating the voxel signals is illustrated in Figure 8.4, achieved by rotating each fibre's signal onto the original fibre direction and weighting by fibre volume. To investigate the impact of FRF variation relative to noise, the experiment was repeated for 500 Rician noise instances at 30 SNR. The fODF estimated using CSD is not a true PDF as it does not integrate to one, so throughout this work the fODF from CSD is normalised as outlined in Section 8.2.1.2.

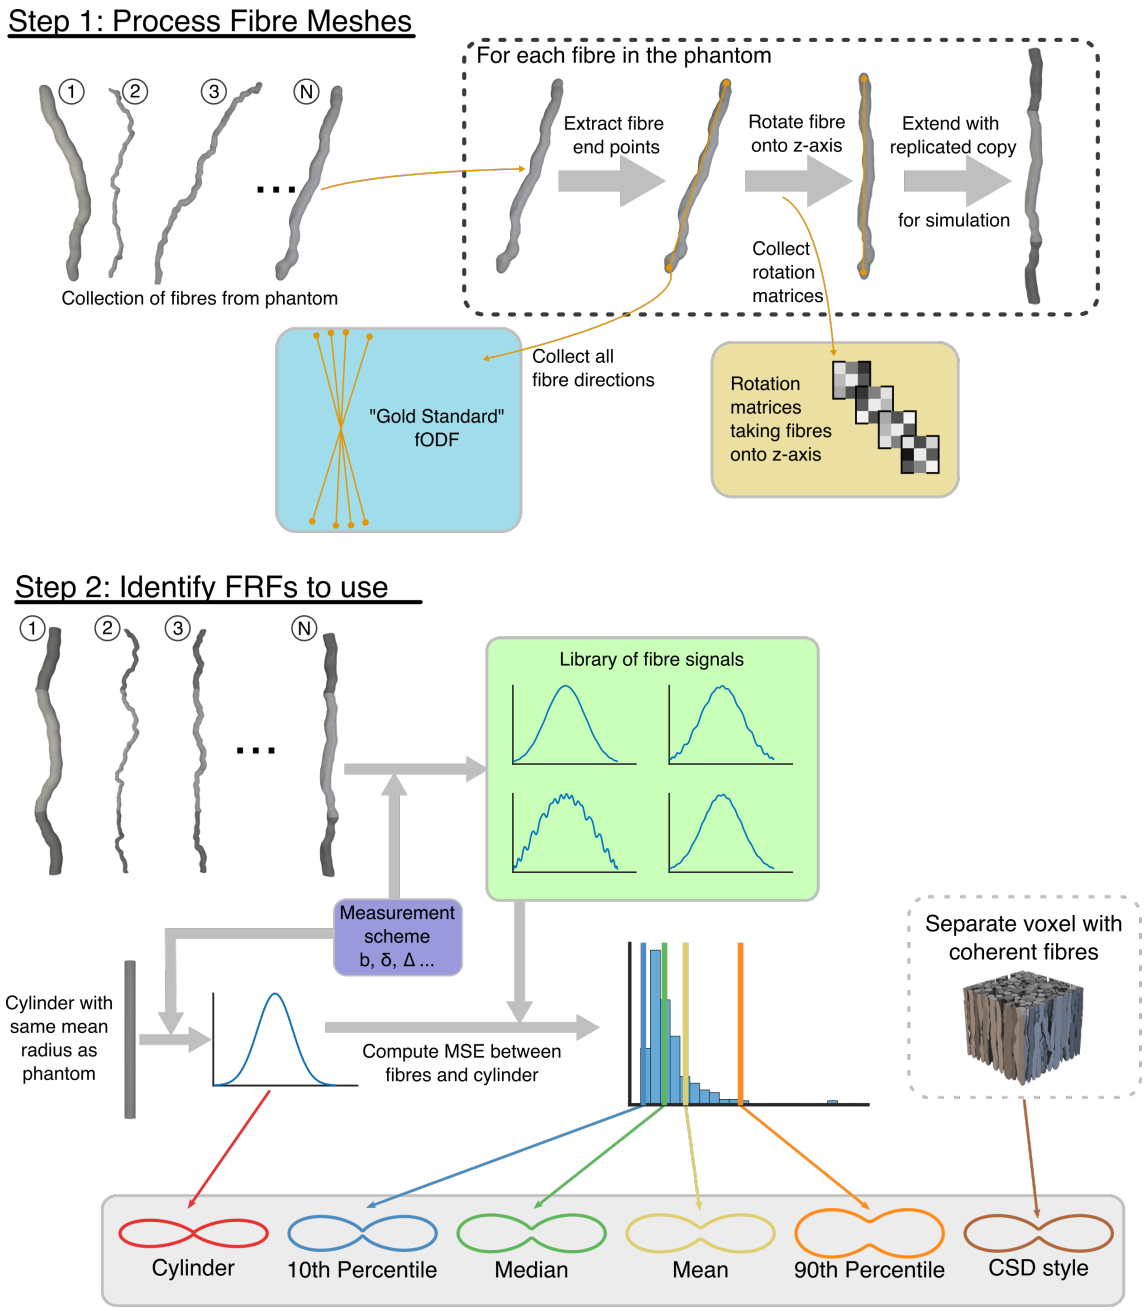

Figure 8.2: Steps 1 and 2 in the experiment pipeline. Step 1 processes fibre meshes to prepare them for simulation as well as to extract gold standard fODF metrics. Each fibre is rotated to align with the z-axis and extended with a reflected copy for simulation as in [168, 171]. The main direction of each fibre as well as the rotation matrix used to align it to z are stored for later use. Step 2 simulates the dMRI signal in each fibre, taking the signal only from the original (not replicated) copy of the fibre and computes the distance between each fibre's signal and a cylinder signal to extract a set of FRFs. This gives a measure of variability in fibre responses in a voxel.

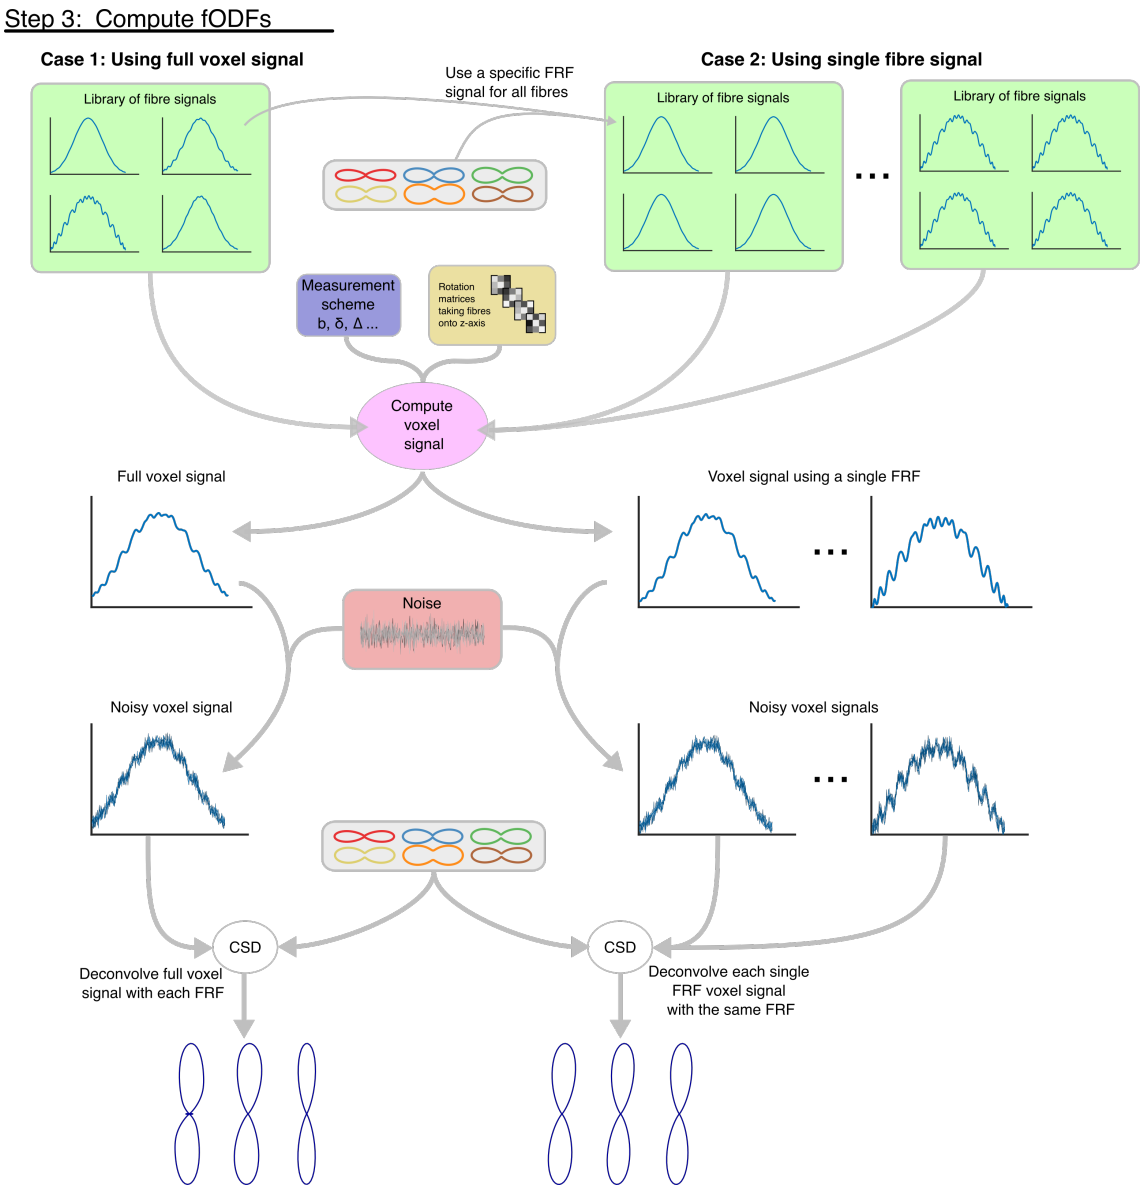

Figure 8.3: The third step in the experiment is to compute voxel signals and perform CSD. There are two strands to the experiment, one using the full voxel signal from each fibre's individual signal and one using a single signal for all fibres, representing the ideal case in which there is a single true FRF. In each case the voxel signal is computed by combining all fibre signals rotated into their original directions and then noise is added. The various FRFs used in the experiments are then deconvolved with the signal using CSD to estimate the fODF.

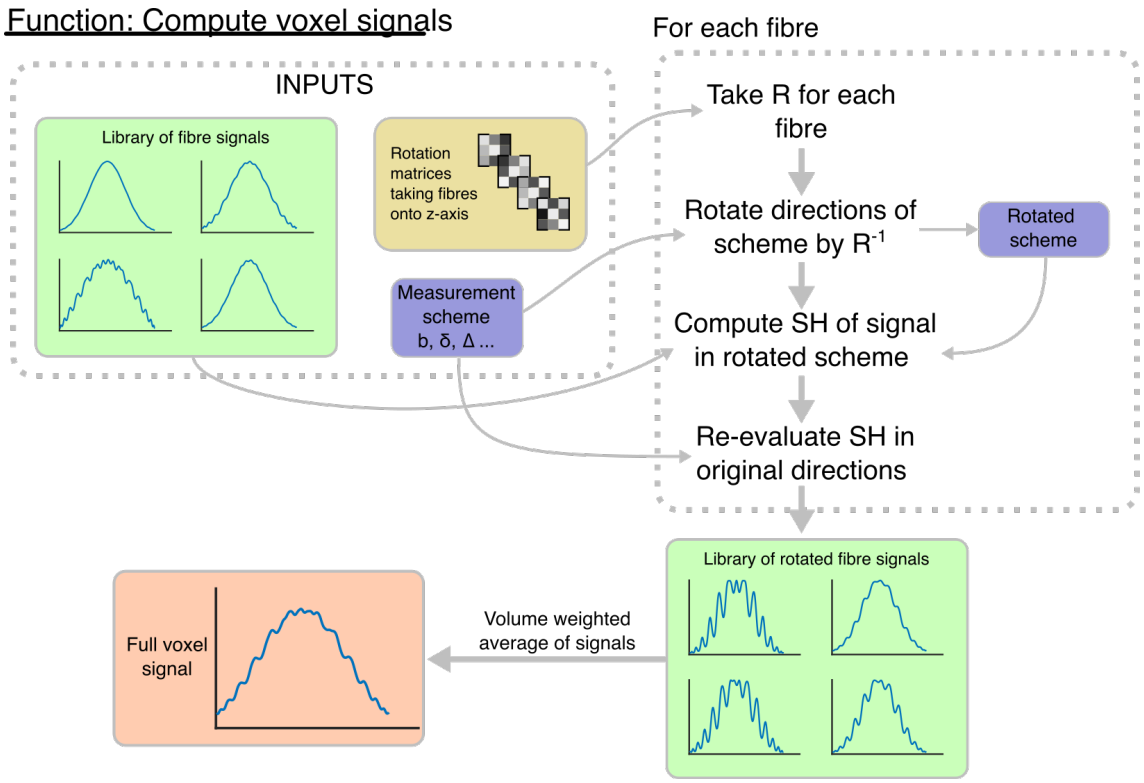

Figure 8.4: Illustration of the method to produce voxel signals. Given a library of fibre signals aligned with z, the rotation matrices taking each fibre onto z and the measurement scheme, each fibre's signal is evaluated in SHs in a rotated scheme and then these SH coefficients are re-evaluated in the original directions. Rotated signals are then combined, weighted by the volume of each fibre, to give the total voxel signal.

8.3 Results

8.3.1 Per-fibre response heterogeneity

Variations in intra-voxel fibre geometries are present in real fibres and ConFiG phantoms as demonstrated in Figure 8.5 which shows each digital phantom alongside the fibres which give the median, 10th and 90th percentile response.

Under the experimental conditions investigated, this morphological variation in the fibres causes the dMRI signal response per-fibre to vary as can be seen in Figure 8.5, which shows the mean, median, 10th and 90th percentile signals across all fibres in each phantom at $b = 3000\,\mathrm{s\,mm^{-2}}$ and 30 SNR. The variation in the response function depends on the complexity of the fibre arrangement, with the most complex three-crossing bundle arrangement leading to the largest variation in response functions.

This variation in the FRF is seen across each b-value from 1000 to $3000\,\mathrm{s\,mm^{-2}}$ as demonstrated in Figure 8.6 for the $\kappa = 2$, three-crossing and EM fibre phantoms, chosen since these display the most variation for each phantom category. Here $\mathrm{SNR} = \infty$ to isolate the effect of b from noise.

8.3.2 Impact on fODF estimation

The variation in the FRF for each fibre leads to a variation in the estimated fODF as seen in Figures 8.7 to 8.9. Again, the magnitude of differences in fODF tends to depend on the complexity of the fibre arrangements since the more complex arrangements have more variation in the FRF. Generally, the fODF calculated from SD picks out the correct main peak direction that is seen in the gold standard fODF from the microstructure, with differences in the overall peak amplitude and shape. Notable exceptions to this occur for some phantoms when using the cylinder FRF, such as EM fibres and three-crossing bundles in which fODF peaks are estimated which are not present in the gold standard fODF.

These qualitative differences in fODF produced with differing FRFs lead to quantitative differences as seen in Figure 8 for the single bundle phantoms. The more stick-like FRFs lead to an underestimation of κ (too much orientation dispersion in the fODF) while FRFs which are more representative of the overall fibre population (mean and median) perform better. In the ideal single FRF voxel signal case, in which the signal truly comes from a single FRF, the fODF is generally estimated well with the main exception being a misestimation of peak direction for $\kappa = 2$ and 6 with the 90th percentile FRF.

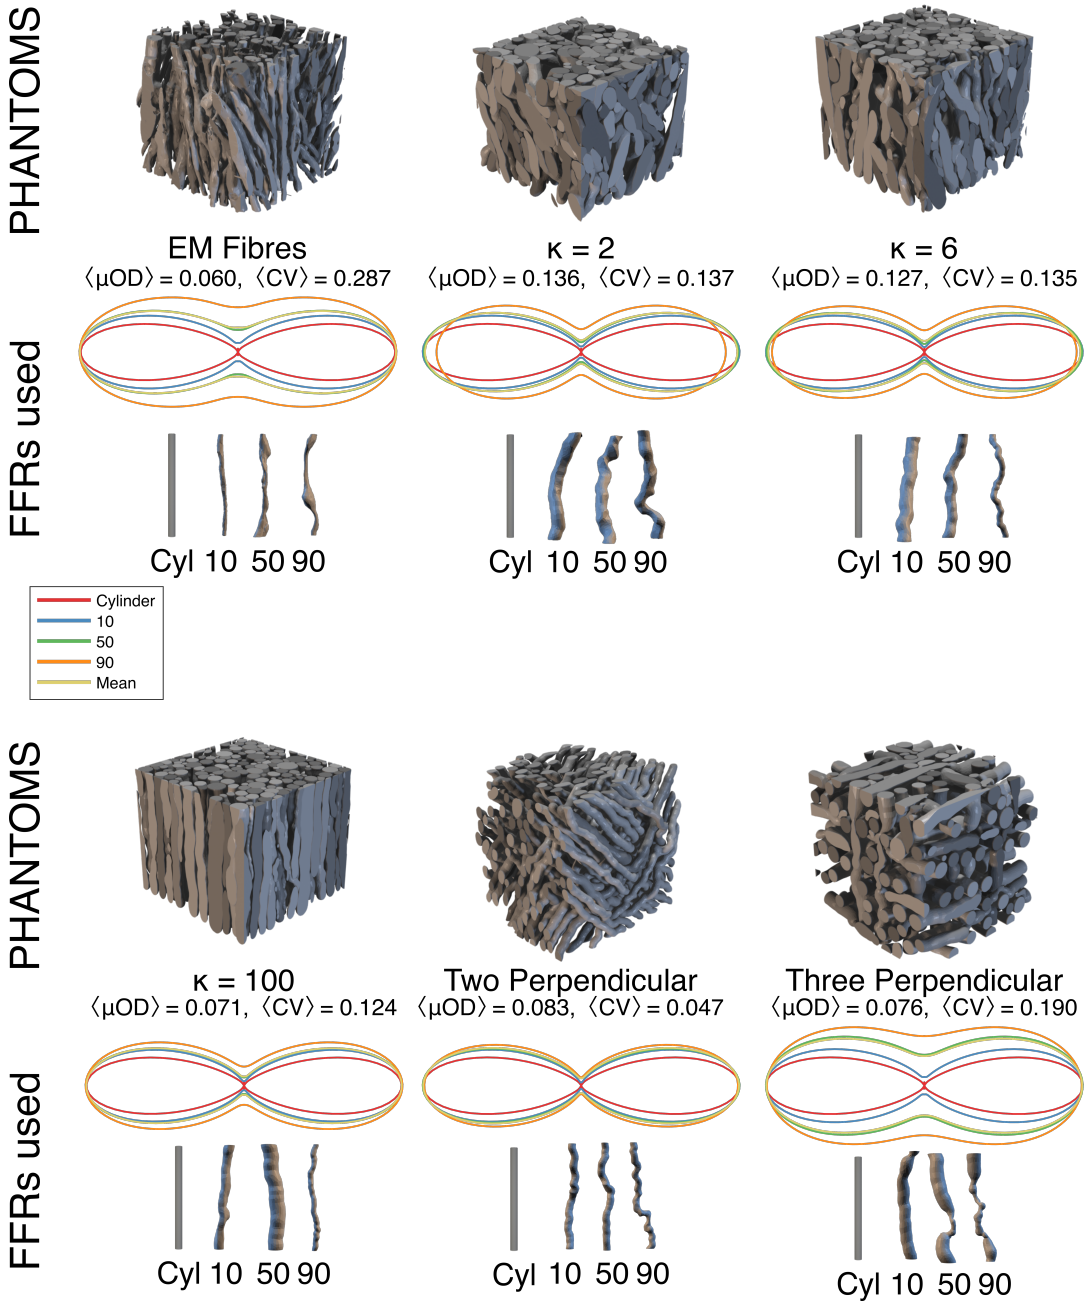

Figure 8.5: Variability in fibre responses within a voxel at $b = 3000\,\mathrm{s\,mm^{-2}}$ at SNR=30 along with geometrical variation in fibres responsible for median, 10th and 90th percentile response.

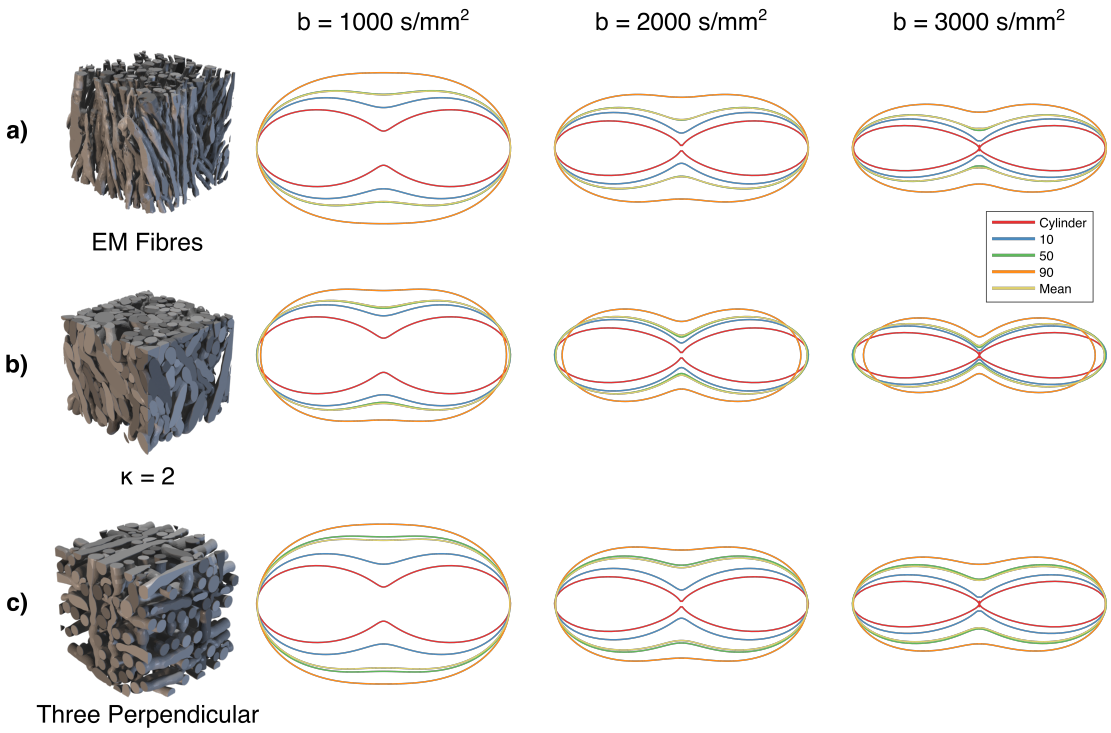

Figure 8.6: Per-fibre response function at $b = 1000,\ 2000,\ 3000\ s/mm^2$ (left-to-right) for (a) the EM fibres, (b) the $\kappa = 2$ phantom and (c) the three-crossing phantom.

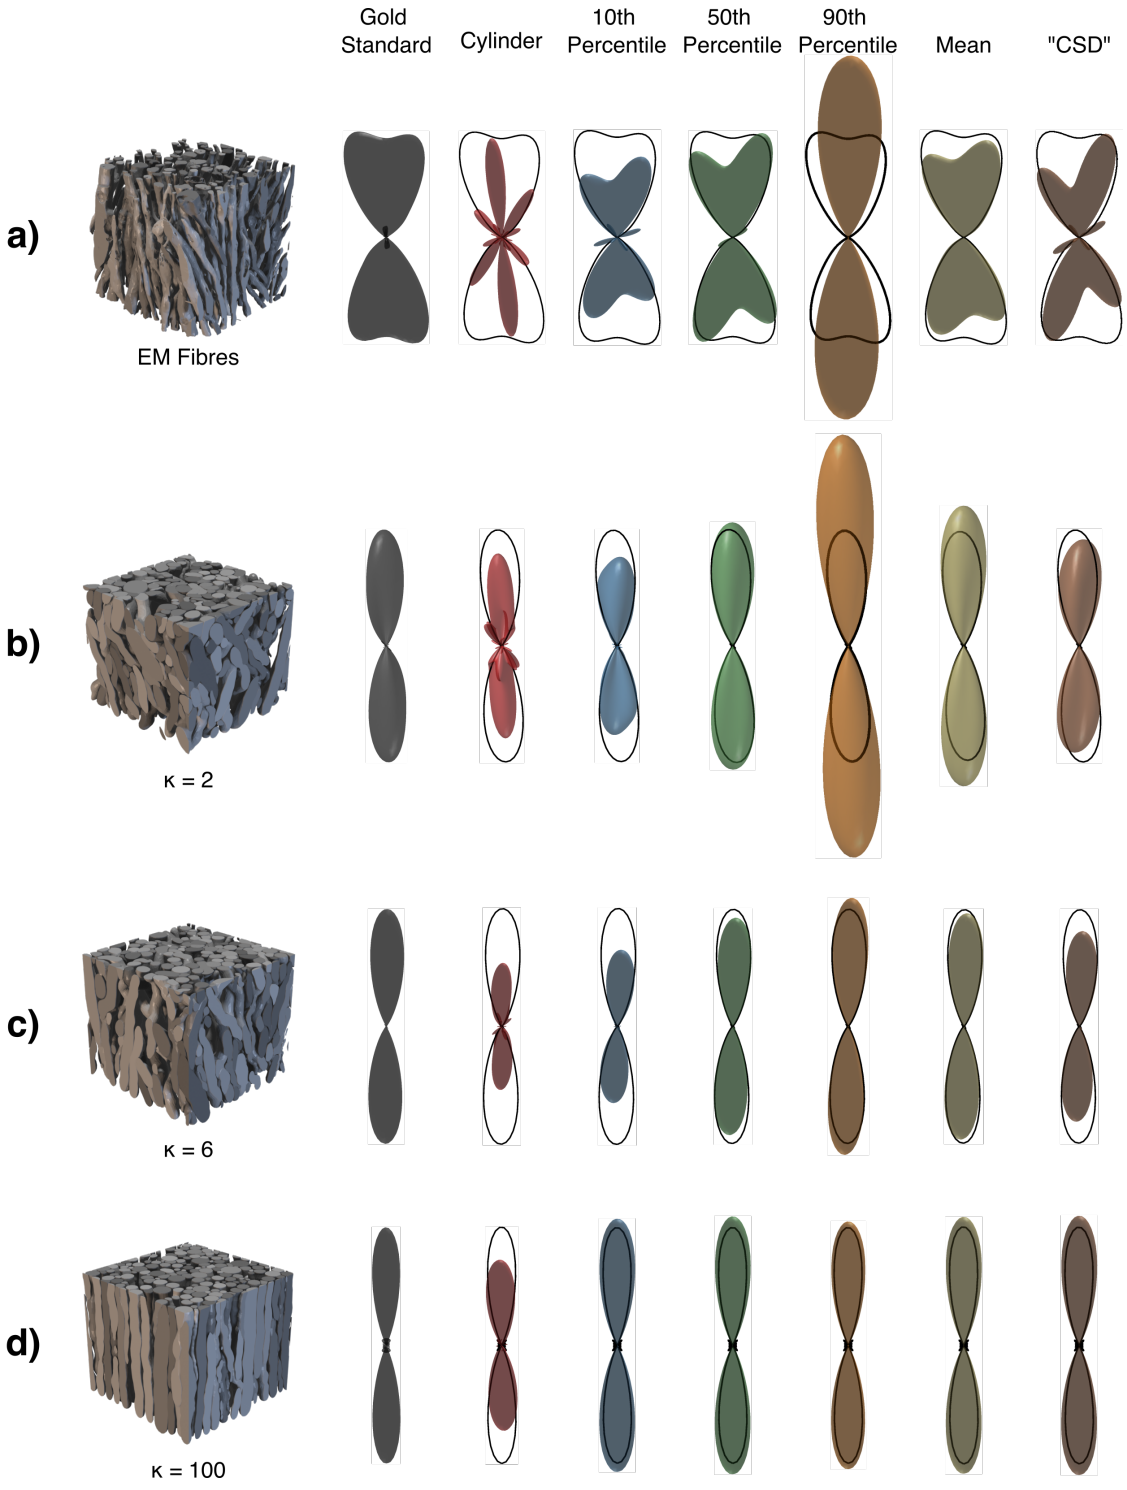

Figure 8.7: Variation in fODF estimated using the range of FRFs in Experiment 2 at b = $3000\,\mathrm{s\,mm^{-2}}$ for the single fibre bundle voxels

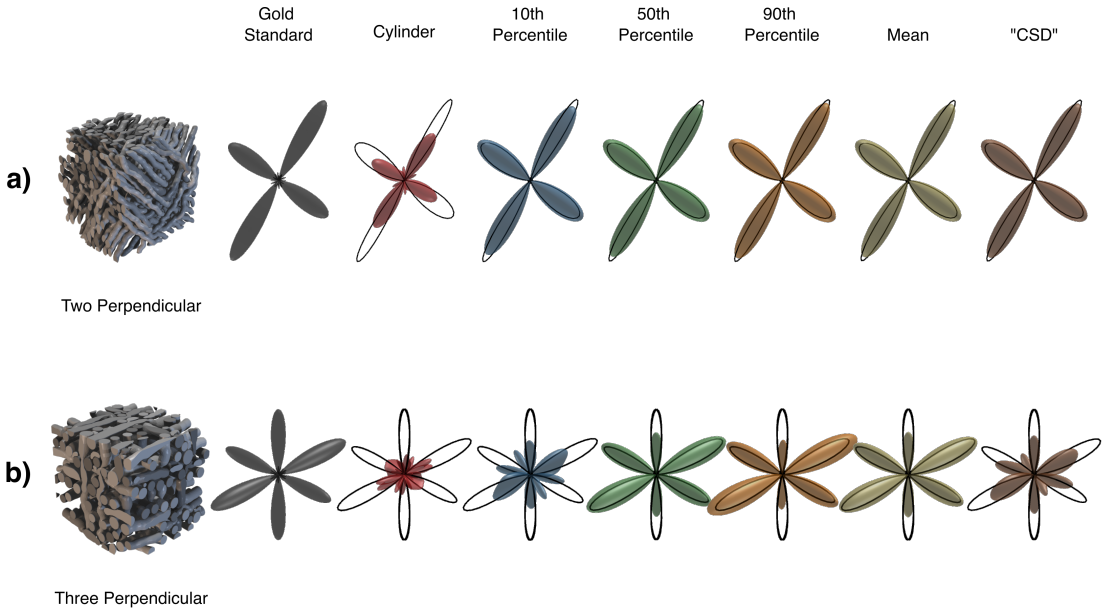

Figure 8.8: Variation in fODF estimated using the range of FRFs in Experiment 2 at b = $3000\,\mathrm{s\,mm^{-2}}$ for the crossing fibre bundle voxels.

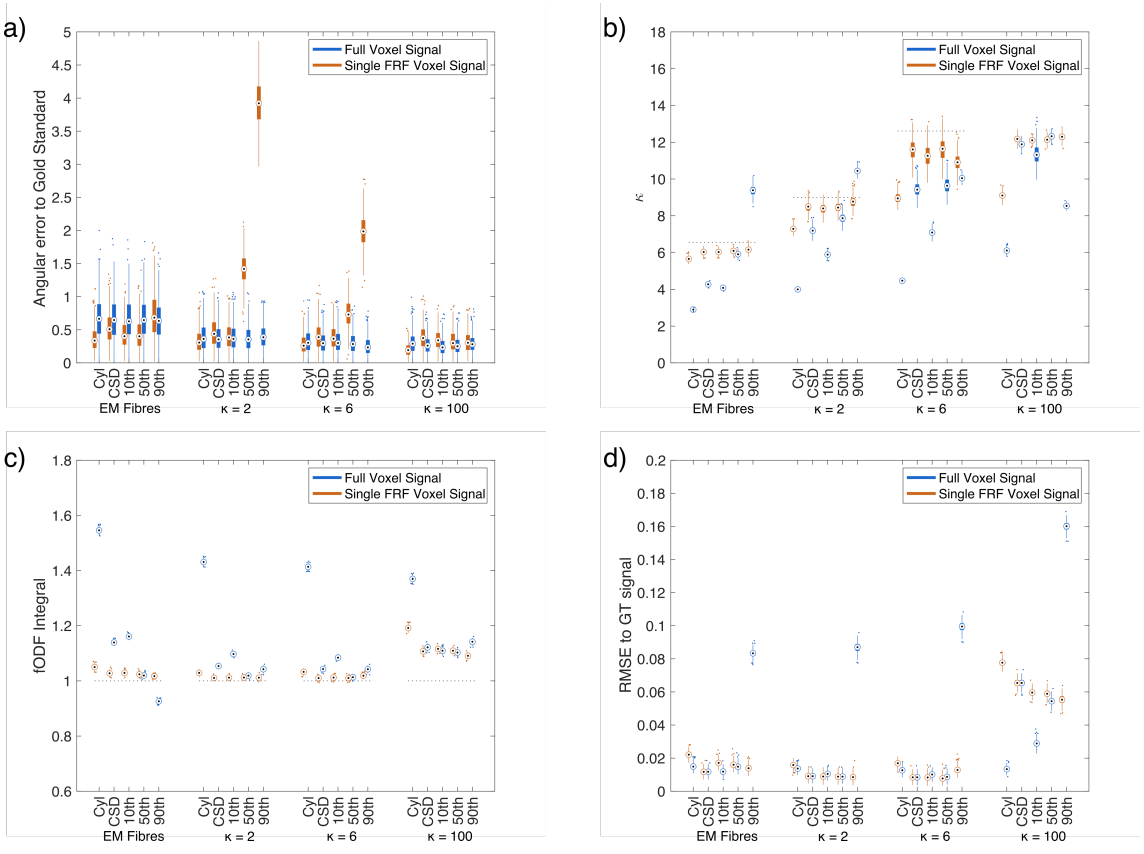

Figure 8.9: Variation in fODF properties when estimated using different FRFs. a) the angular error between the main peak in gold standard fODF and that of the fODFs estimated using CSD, b) κ estimated from the fODF (gold standard value is dashed line for each fODF), c) fODF integral for each fODF estimated and d) the RMSE to the ground truth signal.

8.4 Discussion

The microscopic variations in fibre morphology challenge the assumption in SD techniques that there exists a unique and shared FRF even within a single voxel. Here we have used dMRI simulations to demonstrate that variations in individual axonal morphology do indeed lead to different response functions per-fibre which in turn can have a knock-on effect on SD techniques to estimate the fODF.

As demonstrated in Figures 8.5 and 8.6, the response function can vary substantially across different fibres, particularly in complex fibre arrangements such as the three-crossing bundles scenario. Indeed, the variation in responses in fibres reconstructed from real EM images of WM is large, similar to that of the ConFiG three-crossing simulations, and generally larger than ConFiG phantoms containing a single fibre bundle. This suggests that there is in fact more variation in real axons than in the axons generated by ConFiG, even in simple arrangements of a single bundle of fibres. This is something that may be used to inform future versions of ConFiG to generate more realistic phantoms.

In the main, the largest variation in the per-fibre response, as well as the largest difference between cylinders and realistic axons happens in the axial direction which is to be expected. Even with diameter variations and undulation, the radial diffusion is still strongly restricted under the assumption of no axonal permeability, while recent studies have shown that real axonal morphology causes time-dependent deviations from Gaussian diffusivity along the axial direction [171]. Phantoms which show large amounts of beading (high CV in EM fibres and three-crossing bundles) show the largest variability in response, while μOD affects the response less, suggesting that fibre beading drives fibre response variability more than undulation. One exception to this in the $\kappa = 2$ case in which the 90th percentile fibre contains a large amount of undulation, leading to a reduction in radial signal at higher b-values (Figure 8.6).

Variations in the FRF have an impact on the estimated fODF as demonstrated in Figures 8.7 to 8.9. In the simplest fibre arrangements with low dispersion ($\kappa = 6,\ 100$) the FRF variation is relatively small and the resulting fODF variation is small. In phantoms with more variable FRFs (EM Fibres, three-crossing), the resulting fODFs show more variability, sometimes identifying fewer (e.g. EM fibres 90th percentile) or more peaks (e.g. EM fibres cylinder, three-crossing CSD) than are present in the gold standard fODF and often producing fODFs with differently shaped peaks. This is significant because as shown by Schilling et al. [76], even changes in fODF peak amplitudes without peak direction changes can have an impact on tractography results.

When constructing a voxel signal that truly comes from a single FRF and using that FRF to deconvolve, the fODF is generally recovered well as demonstrated in Figure 8.9. The main exception is the peak direction for $\kappa = 2$ and 6 for the 90th percentile FRF, though this may be caused by the FRF itself containing some orientation dispersion, which makes the voxel signal more difficult to deconvolve as the assumption that each fibre contributes to a single direction is not valid.

Conversely, when using the full voxel signal, the fODF reconstructed depends on the FRF used as shown in Figure 8.9. This demonstrates that when the voxel signal truly comes from fibres with a mix of different responses, assuming a single FRF can lead to misestimation of the fODF. This effect is particularly pronounced when using an FRF that does not represent the fibre population well (Cylinder, 10th percentile, 90th percentile) and would be expected in these cases.

When using an FRF that more broadly represents the fibre population (mean and median), the fODF is generally well estimated, however it is worth noting that the typical CSD approach of taking the signal from a voxel of coherent fibres does not always perform well, producing wrongly shaped peaks for the EM fibre and three-crossing phantoms.

The main takeaway from these investigations is that within-voxel heterogeneity in fibre geometry leads to heterogeneity in the per-fibre response to the extent that using a single FRF in CSD cannot always accurately recover the underlying fODF. Some techniques account for some FRF variation voxel-to-voxel [156, 173], however the investigations presented here suggest that the FRF may vary even within a voxel. In spite of this, when using a FRF that broadly represents the population, such as the median fibre response, CSD is able to recover the fODF well, however the fact remains that currently there is no way to determine this representative FRF which means that SD techniques may suffer from this effect when prescribing a single FRF.

Additionally, these simulations lend further support to challenge the assumption that FRF is constant across the brain as differences in the mean FRF per phantom demonstrate that the overall FRF from different fibre arrangements will be different as a result of the different axonal morphology in each environment. Further, as demonstrated by the CSD-style fODF experiments, using a single FRF across different voxels can lead to misestimation of the estimated fODF, meaning that using the wrong FRF in different brain regions could have large impacts on fODF-based techniques such as tractography. For instance, this could potentially explain fODFs with many peaks in the gyral blade such as those seen in Figure 8.10, which appear similar to those seen in this study in Figure 8.7a when using an incorrect FRF and can lead to spurious fibres crossing the gyral blade in tractography. In this region, it is reasonable to expect complex fibre geometry (and therefore variable FRF) as axons spread out to reach their grey matter termini, however other factors may also influence this such as partial volume. In the absence of in vivo ground truth we cannot say for certain the cause of these fODFs, however this study suggests axonal morphology variations could be a contributing factor.

8.4.1 Limitations and future work

In this work we investigated how complex fibre morphology affects SD techniques using a dMRI scheme chosen to use clinically feasible gradient strength and duration, though it may be possible that other measurement schemes exhibit greater or lesser sensitivity to these effects. For instance Yeh et al. [176] have shown that the gradient pulse duration can impact fibre orientation estimation. Another factor to investigate is the impact of the diffusion time on the observed effects, since longer diffusion times can ‘smooth out’ these microscopic morphological variations as spins are able to diffuse further. At present, the

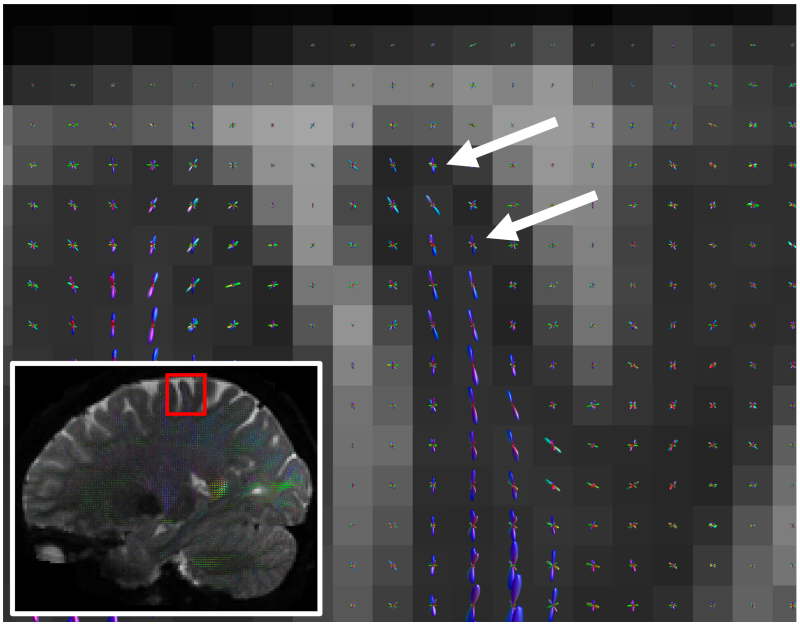

Figure 8.10: An example of spurious peaks in the fODF estimated in the gyral blade of an HCP subject [133, 137] using MRTrix3 [136]. Fibres are expected to project towards the gyral crown but may end up with erroneous peaks in the fODF which can lead to errors in tractography [174, 175].

diffusion time we can simulate is limited by the size of our phantoms, however work is ongoing to produce larger phantoms with ConFiG to allow us to study these effects at longer diffusion time.

Throughout this work we have referred to the fODF generated from the phantom microstructure as the 'gold standard' fODF in inverted commas. This is deliberate since it is not straightforward to define an fODF from microstructure that exactly corresponds to that from dMRI, in part due to the assumptions made in modelling the fODF from dMRI, which have been discussed here. Efforts have been made to make the two fODFs comparable in this work by defining the gold standard fODF using a single direction per fibre and normalising all fODFs to one.

The variability that is demonstrated in the per-fibre response in this experiment is suggested to arise due to the complexity of fibre morphology introduced due to complex fibre arrangements. This seems to be case, as shown in Figure 8.5, though the exact nature of the link is not known for certain since there are many sources of morphological complexity (undulation, beading, non-circular cross-sections, etc.) which could all contribute to these variations. Currently, ConFiG does not allow us to directly control such morphological features but instead the precise morphology generated is a product of the target fibre arrangement (orientation dispersion, density etc.) and the growth algorithm. Future work will aim to incorporate control over microscopic morphological features into ConFiG in order to isolate each of these effects to probe which morphological features have the largest impact on the FRF/fODF. Should these effects be understood, it may be possible to estimate them from the dMRI signal to improve the accuracy of SD techniques.

Another important consideration is that in this work we use CSD as in MRtrix3, however there are wide range of SD techniques for fODF estimation, each with slightly different derivations and assumptions. While this will affect the fODFs presented in this work, the per-fibre signals and compartmental signals presented do not rely on any SD model, so the FRF variations will impact any models which use the FRF.

It is worth pointing out that here we merely demonstrate that the response varies on a per-fibre basis, meaning that the concept of a fibre response function needs to be treated carefully. Most SD techniques estimate their FRF from averaging across a number of voxels, meaning that the FRF is an average single bundle response (including many fibres, extracellular space, other cells etc.) rather than purely a *fibre* response function. The variability in the per-fibre response may contribute less to the variable overall FRF seen across the brain in previous studies [76, 177] than other factors such as the extracellular space but it should be a consideration.

It is also worth noting that this effect will impact other dMRI modelling techniques which model the signal as a combination of a diffusion response with an orientation distribution. For instance, NODDI models the intracellular signal with a Watson distribution as the orientation component and diffusion in sticks as the MR response, assuming that all the fibres can be treated as sticks. As shown here, microscopic variations in fibre morphology mean that the signals from each fibre are not identical and this could affect results of NODDI and other similar models [63]. Similarly, previous studies have looked into the effect of morphological features such as undulation on axon diameter estimation [54, 122, 171]. Future work will aim to shed more light on these effects and investigate whether it is possible for it to be accounted for in our models.

8.5 Conclusion

The complex axonal morphology introduced by axons packing together in complex arrangements leads to differences in the dMRI response across different fibres. These variations in per-fibre response functions can lead to differences in the fODF estimated using CSD when using a single FRF across all fibres and all voxels. Indeed, when the FRF represents the voxel well, CSD can recover the fODF well, however when the FRF is slightly off (as can be the case for assuming a single FRF across all voxels), the fODF can be misestimated.

All of this means that the interpretation of the FRF and fODF in SD needs to be carefully considered and future models may seek to disentangle some of these effects for more accurate FRF and fODF estimation.

Part IV

Conclusions

9 Conclusions

Chapter Contents

9.1 Contributions 143
9.2 Future directions 144

Chapter Summary

The purpose of this chapter is to summarise the contributions made by the work presented in this thesis, including the new white matter numerical phantom generator, ConFiG, and discuss future directions for the project.

9.1 Contributions

In this thesis, we have presented a new method for generating realistic numerical phantoms of white matter (WM) for diffusion MRI (dMRI) simulations. Our primary contribution is the development of contextual fibre growth (ConFiG), a novel approach to phantom generation in which we mimic the processes governing the growth of real axons in an effort to produce realistic WM phantoms. We demonstrated, though a preliminary implementation and subsequent improvements, that ConFiG is able to produce phantoms with state-of-the-art performance in terms of the axonal density accomplished, by following simple rules inspired by the biological processes governing axonal guidance.

Diffusion MRI simulations show that the microstructure generated by ConFiG can be used to generate realistic synthetic dMRI data, acting as a proof-of-concept of the utility of ConFiG phantoms. Comparisons to real axons segmented from EM show that ConFiG produces axons with similar morphologies to real axons, capturing not just axon diameter and orientation distributions, but also subtle features such as variable diameters along axons and non-circular cross-sections. Taking all of this into account, ConFiG phantoms represent a new generation of WM numerical phantoms, greatly increasing phantom realism when compared to the current standard of WM numerical phantoms which are based on cylinders.

As example applications of ConFiG, we investigated the intra-axonal diffusion in realistic geometries and probed assumptions that are inherent in spherical deconvolution (SD) based dMRI modelling techniques. We have demonstrated that undulation and beading have distinct effects on the diffusion time dependence of axial diffusivity and that a simple phenomenological model can capture some of these effects. Realistic axons however, are not particularly well described by the undulation term in this phenomenological model. Further, we showed that models which attempt to model the intracellular diffusivity by modelling the axons as sticks may end up with biased estimates by ignoring the true microscopic complexity of axonal morphology.

Additionally, we were able to demonstrate that microscopic variations in the shapes of individual axons within a bundle mean that the intra-axonal signal varies per-fibre, violating one of the assumptions in SD techniques that each fibre has an identical dMRI response. The variability in the FRF also leads to a variability in the fODF estimated which could have a large impact fODF based techniques such as tractography.

9.2 Future directions

ConFiG opens up many possibilities for future investigation by giving us the ability to generate more realistic WM phantoms than ever before. One such possibility is to study the dynamics of the diffusion process itself. For instance, we could study the diffusion time dependence in realistic WM phantoms, similarly to recent work investigating intracellular diffusion [107, 134], however ConFiG offers us the potential to study the realistic extracellular space too. This could help to disentangle different sources of diffusion time dependence (axonal beading, undulation, extracellular space etc.) and identify which source is the most significant.

Further, we can investigate how best to sensitise the dMRI signal to these complex microstructural features. Recent advances in dMRI acquisition have given us novel diffusion encoding schemes such as b-tensor encoding [178] which provide additional information to standard pulsed gradient spin echo (PGSE) techniques. While this gives us a new tool to study microstructure, the potential space of possible imaging schemes becomes very large, making it difficult to know which schemes may be best to identify which features. This is something that ConFiG gives us an opportunity to empirically test, by generating a wide range of phantoms covering different microstructural environments and simulating various dMRI sequences to test which sequences or combination of sequences give us the best sensitivity.

On a related note, ConFiG enables us to study how realistic microstructure affects dMRI signals and subsequent modelling techniques using these new b-tensor encoding schemes. For instance, Kerkelä et al. [179] in collaboration with authors of this thesis, have used ConFiG phantoms to investigate how well different models are able to estimate microscopic fractional anisotropy, a measure of the anisotropy of tissue compartments regardless of their orientation dispersion. This is just one example of myriad possible applications of ConFiG for testing the validity and proving limits of existing dMRI modelling techniques.

For instance, ConFiG also offers us a new tool for validating existing models of dMRI microstructure. Previous models for measuring microstructural features such as axon diameter or orientation dispersion have used simple cylinder models in simulation validations [13,26,63,119], ConFiG enables us to test these models in realistic geometries. In principle, this offers us the possibility for a more comprehensive validation of dMRI modelling techniques, it will require some thought because certain microstructural features are not intuitively defined in these complex phantoms. For instance, the fibre diameter is simple to define when using cylinders, but when you have complex fibres that not only have a variable diameter along the fibres but also non-circular cross-sections, such a thing becomes more difficult to define. ConFiG phantoms may enable us to probe some of these questions, however, to find out which features can be best estimated using dMRI, for instance are we more sensitive to diameter variation along fibres (i.e beading) or the overall mean diameter.

The experiments presented in Chapter 8 demonstrate an example of this kind of application of ConFiG to investigate a dMRI model, demonstrating that certain assumptions in SD techniques may not hold true. Beyond this application to simply probe dMRI models, we may use this to develop improved models. For instance, it may be possible to include variability in the FRF into the modelling process to try to capture the variations in axonal morphology. Beyond SD techniques, there are many dMRI modelling techniques which take a similar approach to the representation of the dMRI signal, splitting it into a diffusion response term and an orientation distribution term. An example of this is neurite orientation dispersion and density imaging (NODDI) [13], which uses a Watson distribution for the orientation dispersion distribution and the MR response is diffusion in sticks, with an extracellular and cerebrospinal fluid (CSF) compartment. As in the SD case however, NODDI uses the same basic MR response for all the fibres (that is, they are all sticks with the same diffusivity), while we have shown that fibres have different responses. Again, it may be possible to attempt to account for this variability

when modelling dMRI using models such as NODDI.

Another potential application of ConFiG is in the development of a new computational model of WM microstructure. The idea here would be to estimate microstructural features by learning a mapping between dMRI signals and microstructure. This could be done by generating a wide range of WM numerical phantoms spanning the range of possible expected microstructural environments in the brain and simulating the dMRI signal within them to generate a dictionary of signals paired with ground-truth microstructure. As mentioned above, ConFiG phantoms may be used to identify the set of dMRI sequences that give us the best sensitivity to various microstructure features to give this computational model the best chance to work effectively. The benefit of developing this computational model would be that it can capture the relationship between complex microstructure and dMRI signal without relying on an explicit analytical model and further may be able to identify features which do not have a simple analytical expression such as axonal beading. Such approaches have recently been used to estimate membrane permeability, one such feature that is hard to model analytically [16, 18].

ConFiG itself can be also be improved. Currently the maximum size of a phantom is limited to around $50\,\mu\mathrm{m} \times 50\,\mu\mathrm{m} \times 50\,\mu\mathrm{m}$ which limits the diffusion time that can be explored without needing to extend phantoms with replicated copies. This is down to memory and computational time constraints from using a large number of points in the growth network. As mentioned in Section 5.5, this could be addressed by growing in layers enabling dense sampling of small regions of space using fewer points. Another option is to reconfigure the algorithm to remove the reliance on the network. Currently the network plays two main roles, one is defining the paths available for a fibre to grow and one is sampling the space so that growing fibres know where existing fibres are. A potentially better approach is to separate these two processes so that the sampling of the space is done by one structure (potentially a Delaunay/Voronoi based approach as currently or potentially something like an Octree or bounding volume hierarchy (BVH) [180]) and the paths are left separate to enable fibres to have more freedom in growth directions. This approach could potentially enable faster and more memory-efficient growth of ConFiG phantoms while giving the fibres more freedom in paths could enable even high density of phantoms.

The potential future directions presented here are only limited to applications of ConFiG to dMRI of WM. Since the growth idea behind ConFiG is quite general, it is easy to see how ConFiG could be adapted to apply to other tissues (for instance adding branching cells and vasculature) by changing the growth rules. Additionally, since the end product of ConFiG is a 3D surface mesh, these could be used in any other simulator that can handle a 3D mesh. Very basic examples of this are presented throughout the thesis as the meshes are used in Blender to render (effectively simulating visible light imaging) 3D images and 2D ‘virtual histology’ images. All of this means that ConFiG represents a potentially powerful new tool for the dMRI community and beyond.

Publications

Journal Articles

1. **Ross Callaghan**, Daniel C. Alexander, Marco Palombo and Hui Zhang. 2020. *Contextual fibre growth to generate realistic axonal packing for diffusion MRI simulation*. NeuroImage 220, 117107. https://doi.org/10.1016/j.neuroimage.2020.117107
2. Leevia 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. 2021. *Comparative analysis of signal models for microscopic fractional anisotropy estimation using q-space trajectory encoding*. NeuroImage 242, 118445.
3. **Ross Callaghan**, Daniel C. Alexander, Marco Palombo and Hui Zhang. 2022. *Impact of within-voxel heterogeneity in fibre geometry on spherical deconvolution*. (*Under Review at NeuroImage*)

Full Length Conference Articles

1. **Ross Callaghan**, Daniel C. Alexander, Hui Zhang and Marco Palombo. 2019. *Contextual fibre growth to generate realistic axonal packing for diffusion MRI simulation*. 26th International Conference on Information Processing in Medical Imaging (IPMI)

Conference Abstracts

1. **Ross Callaghan**, Noam Shemesh, Daniel C. Alexander, Hui Zhang, and Marco Palombo. 2020. *Contextual fibre growth to generate realistic axonal packing for diffusion MRI simulation*. 28th Annual Meeting of the International Society for Magnetic Resonance in Medicine (ISMRM)
2. **Ross Callaghan**, Noam Shemesh, Daniel C. Alexander, Hui Zhang, and Marco Palombo. 2019. *Towards a more realistic and flexible white matter numerical phantom generator for diffusion MRI simulation*. 27th Annual Meeting of the International Society for Magnetic Resonance in Medicine (ISMRM)
3. Andrada Ianus, **Ross Callaghan**, Daniel C. Alexander, Hui Zhang and Marco Palombo. 2019. *Effect of cell complexity and size on diffusion MRI signal: a simulation study*. 27th Annual Meeting of the International Society for Magnetic Resonance in Medicine (ISMRM)
4. **Ross Callaghan**, Daniel C. Alexander, Hui Zhang and Marco Palombo. 2019. *Contextual fibre growth to generate realistic axonal packing for diffusion MRI simulation*. 25th Annual Meeting of the Organization for Human Brain Mapping (OHBM)
5. **Ross Callaghan**, Bhavana Solanky, Sotirios Bisdas, Hui Zhang and Enrico De Vita. 2018. *Comparing MEGA editing techniques for in-vivo measurement of 2-hydroxyglutarate*. Proceedings of the 26th Annual Meeting of the International Society for Magnetic Resonance in Medicine (ISMRM)

Workshop Presentations

1. **Ross Callaghan**, Daniel C. Alexander, Hui Zhang and Marco Palombo. 2018. *Contextual fibre growth to generate realistic axonal packing for diffusion MRI simulation*. In School of Brain Cells & Circuits "Camillo Golgi": The Neural Bases of Action: from cellular microcircuits to large-scale networks and modelling

Colophon

This thesis is typeset using LaTeXand BibTeX using a custom template modified from Nuno Cardoso's PhD amazing thesis template (https://github.com/npcardoso/PhDThesis). The updated thesis template is available at https://github.com/rcallagh/ThesisTemplate.

Written and compiled using GNU Emacs with the AUCTeX package (https://www.gnu.org/software/emacs/emacs, https://www.gnu.org/software/auctex).

References

[1] S. Ackerman, *Discovering the Brain*. Washington, DC: Institute of Medicine and National Academy of Sciences, The National Academies Press, 1992. (Cited on page 19)

[2] F. Crick and E. Jones, "Backwardness of human neuroanatomy," *Nature*, vol. 361, pp. 109–110, 1993. (Cited on page 19)

[3] H. Johansen-Berg and T. E. J. Behrens, *Diffusion MRI: From Quantitative Measurement to In vivo Neuroanatomy: Second Edition*. 2013. (Cited on page 19)

[4] P. J. Basser, J. Mattiello, and D. Lebihan, "Estimation of the Effective Self-Diffusion Tensor from the NMR Spin Echo," *Journal of Magnetic Resonance, Series B*, vol. 103, no. 3, pp. 247–254, 1994. (Cited on page 20)

[5] M. Palombo, A. Ianus, M. Guerreri, D. Nunes, D. C. Alexander, N. Shemesh, and H. Zhang, "SANDI: A compartment-based model for non-invasive apparent soma and neurite imaging by diffusion MRI," *NeuroImage*, vol. 215, no. April, p. 116835, 2020. (Cited on pages 20, 54, 66 and 126)

[6] J. R. Li, V. D. Nguyen, T. N. Tran, J. Valdman, C. B. Trang, K. V. 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, no. January, p. 116120, 2019. (Cited on pages 20, 54 and 66)

[7] I. O. Jelescu and M. D. Budde, "Design and Validation of Diffusion MRI Models of White Matter," *Frontiers in Physics*, vol. 5, no. November, 2017. (Cited on pages 20, 41, 54 and 66)

[8] M. Nilsson, S. Lasič, I. Drobnjak, D. Topgaard, and C. F. Westin, "Resolution limit of cylinder diameter estimation by diffusion MRI: The impact of gradient waveform and orientation dispersion," *NMR in Biomedicine*, vol. 30, no. 7, pp. 1–13, 2017. (Cited on pages 20, 46, 54 and 66)

[9] B. Scherrer, A. Schwartzman, M. Taquet, M. Sahin, S. P. Prabhu, and S. K. Warfield, "Characterizing brain tissue by assessment of the distribution of anisotropic microstructural environments in diffusion-compartment imaging (DIAMOND)," *Magnetic Resonance in Medicine*, vol. 76, no. 3, pp. 963–977, 2016. (Cited on pages 20, 54 and 66)

[10] M. Tariq, T. Schneider, D. C. Alexander, C. A. Gandini Wheeler-Kingshott, and H. Zhang, "Bingham-NODDI: Mapping anisotropic orientation dispersion of neurites using diffusion MRI," *NeuroImage*, vol. 133, pp. 207–223, 2016. (Cited on pages 20, 54, 66, 76 and 77)

[11] 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. (Cited on pages 20, 54 and 66)

[12] J. Xu, H. Li, K. D. Harkins, X. Jiang, J. Xie, H. Kang, M. D. Does, and J. C. Gore, "Mapping mean axon diameter and axonal volume fraction by MRI using temporal diffusion spectroscopy," *NeuroImage*, vol. 103, pp. 10–19, 2014. (Cited on pages 20, 44, 45, 54 and 66)

[13] 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. (Cited on pages 20, 41, 54, 66, 74, 76, 104, 122 and 145)

[14] M. Nilsson, E. Alerstam, R. Wirestam, F. Ståhlberg, S. Brockstedt, and J. Lätt, "Evaluating the accuracy and precision of a two-compartment Karger model using Monte Carlo simulations," *Journal of Magnetic Resonance*, vol. 206, no. 1, pp. 59–67, 2010. (Cited on pages 20, 21, 46, 47, 48, 54 and 66)

[15] G. Rensonnet, B. Scherrer, G. Girard, A. Jankovski, S. K. Warfield, B. Macq, J.-P. Thiran, and M. Taquet, "Towards microstructure fingerprinting: Estimation of tissue properties from a dictionary of Monte Carlo diffusion MRI simulations," *NeuroImage*, vol. 184, no. May 2018, pp. 964–980, 2018. (Cited on pages 20, 48, 54, 66 and 86)

[16] I. Hill, M. Palombo, M. Santin, F. Branzoli, and A.-c. Philippe, "Machine learning based white matter models with permeability : An experimental study in cuprizone treated in-vivo mouse model of axonal demyelination," *ArXiv*, 2019. (Cited on pages 20, 48, 54, 66, 86 and 146)

[17] M. Palombo, I. Hill, M. D. Santin, F. Branzoli, A.-C. Philippe, D. Wassermann, M.-S. Aigrot, B. Stankoff, H. Zhang, S. Lehericy, A. Petiet, D. C. Alexander, and I. Drobnjak, "Machine learning based estimation of axonal permeability: validation on cuprizone treated in-vivo mouse model of axonal demyelination," in *Proc. Joint Annual Meeting ISMRM-ESMRMB, Paris, France*, 2018. (Cited on pages 20, 48, 54, 66 and 86)

[18] G. L. Nedjati-Gilani, T. Schneider, M. G. Hall, N. Cawley, I. Hill, O. Ciccarelli, I. Drobnjak, C. A. Wheeler-Kingshott, and D. C. Alexander, "Machine learning based compartment models with permeability for white matter microstructure imaging," *NeuroImage*, 2017. (Cited on pages 20, 48, 54, 66, 83, 86 and 146)

[19] W. Li, B. Wu, A. V. Avram, and C. Liu, "Magnetic susceptibility anisotropy of human brain in vivo and its molecular underpinnings," *NeuroImage*, vol. 59, no. 3, pp. 2088–2097, 2012. (Cited on pages 21 and 54)

[20] J. Lee, K. Shmueli, M. Fukunaga, P. Van Gelderen, H. Merkle, A. C. Silva, and J. H. Duyn, "Sensitivity of MRI resonance frequency to the orientation of brain tissue microstructure," *Proceedings of the National Academy of Sciences of the United States of America*, vol. 107, no. 11, pp. 5130–5135, 2010. (Cited on pages 21 and 54)

[21] T. Xu, S. Foxley, M. Kleinnijenhuis, W. C. Chen, and K. L. Miller, "The effect of realistic geometries on the susceptibility-weighted MR signal in white matter," *Magnetic Resonance in Medicine*, vol. 79, no. 1, pp. 489–500, 2018. (Cited on pages 21, 54 and 86)

[22] E. Fieremans, D. S. Novikov, J. H. Jensen, and J. A. Helpern, "Monte Carlo study of a two-compartment exchange model of diffusion," *NMR in Biomedicine*, vol. 23, no. 7, pp. 711–724, 2010. (Cited on pages 21, 46 and 66)

[23] M. Nilsson, J. Lätt, E. Nordh, R. Wirestam, F. Ståhlberg, and S. Brockstedt, "On the effects of a varied diffusion time in vivo: is the diffusion in white matter restricted?," *Magnetic Resonance Imaging*, vol. 27, no. 2, pp. 176–187, 2009. (Cited on pages 21, 46, 66 and 75)

[24] J. C. Ford and D. B. Hackney, "Numerical model for calculation of apparent diffusion coefficients (ADC) in permeable cylinders - Comparison with measured ADC in spinal cord white matter," *Magnetic Resonance in Medicine*, vol. 37, no. 3, pp. 387–394, 1997. (Cited on pages 21, 39 and 66)

[25] 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. (Cited on pages 21, 37, 39, 45, 48, 56, 66, 75, 76, 105 and 126)

[26] 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. (Cited on pages 21, 45, 48, 66 and 145)

[27] A. Abdollahzadeh, B. Ilya, E. Jokitalo, J. Tohka, and A. Sierra, "Automated 3D Axonal Morphometry of White Matter," *Scientific Reports*, vol. 9, pp. 1–16, 2019. (Cited on pages 21, 48, 49, 66, 88, 90, 104 and 120)

[28] H. H. Lee, K. Yaros, J. Veraart, J. L. Pathan, F. Xia, L. Sungheon, D. S. Novikov, and E. Fieremans, "Along-axon diameter variation and axonal orientation dispersion revealed with 3D electron microscopy : implications for quantifying brain white matter microstructure with histology and diffusion MRI," *Brain Structure and Function*, vol. 0, no. 0, p. 0, 2019. (Cited on pages 21, 48, 49, 66, 84, 88, 89, 90, 93, 94, 96, 97, 100, 104, 107, 121 and 122)

[29] M. H. Levitt, *Spin Dynamics: Basics of Nuclear Magnetic Resonance*. 2008. (Cited on pages 25, 26, 27 and 33)

[30] P. B. Barker, A. Bizzi, N. De Stefano, R. Gullapalli, and D. D. M. Lin, *Clinical MR Spectroscopy*. 2009. (Cited on page 25)

[31] F. Bloch, "Nuclear Induction," *Physical Review*, vol. 652, no. 1938, pp. 460–474, 1946. (Cited on pages 25 and 28)

[32] E. M. Purcell, H. C. Torrey, and R. V. Pound, "Resonance Absorption by Nuclear Magnetic Moments in a Solid," *Physical Review*, vol. 69, no. 1-2, 1946. (Cited on page 25)

[33] N. M. Salibi and M. A. Brown, *Clinical MR spectroscopy : first principles*. Wiley-Liss, 1998. (Cited on page 25)

[34] E. Haacke, R. Brown, M. Thompson, and R. Venkatesan, "Magnetic resonance imaging: physical principles and sequence design. 1999," *New York: A John Wiley and Sons*, 1999. (Cited on pages 26, 28, 32 and 35)

[35] R. A. de Graaf, *In Vivo NMR Spectroscopy: Principles and Techniques: 2nd Edition*. 2007. (Cited on pages 29, 30, 31 and 33)

[36] E. L. Hahn, "Spin Echoes," *Physical Review*, vol. 80, no. 4, pp. 580–594, 1950. (Cited on page 33)

[37] M. W. Denny, *Air and Water: The Biology and Physics of Life's Media*. Air and Water: The Biology and Physics of Life's Media, Princeton University Press, 1993. (Cited on page 35)

[38] 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. (Cited on pages 36 and 37)

[39] W. Price, "Pulsed-field gradient nuclear magnetic resonance as a tool for studying translational diffusion: Part 1. Basic theory," *Concepts Magn. Reson.*, vol. 9, pp. 299–336, 1997. (Cited on pages 37, 38, 44 and 46)

[40] P. W. Kuchel, G. Pages, K. Nagashima, V. Sendhil, V. Vijayaragavan, V. Nagarajan, and K. H. Chuang, "Stejskal–Tanner Equation Derived in Full," *Concepts in Magnetic Resonance*, vol. 40, no. 5, pp. 205–214, 2012. (Cited on page 37)

[41] C. H. Neuman, "Spin echo of spins diffusing in a bounded medium," *The Journal of Chemical Physics*, vol. 60, no. 11, pp. 4508–4511, 1974. (Cited on page 37)

[42] B. Balinov, B. Jönsson, P. Linse, and O. Söderman, "The NMR self-diffusion method applied to restricted diffusion. simulation of echo attenuation from molecules in spheres and between planes," 1993. (Cited on page 37)

[43] P. T. Callaghan, "Pulsed-Gradient Spin-Echo NMR for Planar, Cylindrical, and Spherical Pores under Conditions of Wall Relaxation," 1995. (Cited on page 37)

[44] H. C. Torrey, "Bloch equations with diffusion terms," *Physical Review*, vol. 104, no. 3, pp. 563–565, 1956. (Cited on pages 38 and 39)

[45] P. T. Callaghan, *Principles of Nuclear Magnetic Resonance Microscopy*, vol. 19. 1991. (Cited on pages 38 and 44)

[46] W. S. Price, *NMR studies of translational motion*. 2009. (Cited on page 38)

[47] J. Jeener, "Macroscopic molecular diffusion in liquid NMR, revisited," *Concepts in Magnetic Resonance Part B: Magnetic Resonance Engineering*, vol. 14, no. 2, pp. 79–88, 2002. (Cited on page 38)

[48] L. Beltrachini, Z. A. Taylor, and A. F. Frangi, "An efficient finite element solution of the generalised Bloch-Torrey equation for arbitrary domains," *Mathematics and Visualization*, no. FEM, pp. 3–14, 2016. (Cited on pages 38 and 44)

[49] C. Grossmann and H.-G. Roos, *Numerical Treatment of Partial Differential Equations*. Springer, 2007. (Cited on page 39)

[50] H. Hagslätt, B. Jönsson, M. Nydén, and O. Söderman, "Predictions of pulsed field gradient NMR echo-decays for molecules diffusing in various restrictive geometries. Simulations of diffusion propagators based on a finite element method," *Journal of Magnetic Resonance*, vol. 161, no. 2, pp. 138–147, 2003. (Cited on pages 39 and 44)

[51] D. L. Logan, E. Veitch, C. Carson, K. R. Burrell, V. Gould, E. Wagner, D. L. Logan, E. Veitch, C. Carson, K. R. Burrell, V. Gould, and E. Wagner, *A First Course in the Finite Element Method Fourth Edition*, vol. 147. 2007. (Cited on page 39)

[52] A. Iserles, *A first course in the numerical analysis of differential equations*. 2009. (Cited on page 39)

[53] C.-H. Yeh, B. Schmitt, D. Le Bihan, J.-R. Li-Schlittgen, C.-P. Lin, and C. Poupon, "Diffusion Microscopist Simulator: A General Monte Carlo Simulation System for Diffusion Magnetic Resonance Imaging," *PLoS ONE*, vol. 8, no. 10, p. e76626, 2013. (Cited on pages 39 and 45)

[54] M. Nilsson, J. Lätt, F. Ståhlberg, D. van Westen, and H. Hagslätt, "The importance of axonal undulation in diffusion MR measurements: A Monte Carlo simulation study," *NMR in Biomedicine*, vol. 25, no. 5, pp. 795–805, 2012. (Cited on pages 39, 46, 66, 104 and 139)

[55] B. A. Landman, J. A. D. Farrell, S. A. Smith, D. S. Reich, P. A. Calabresi, and P. C. M. Van Zijl, "Complex geometric models of diffusion and relaxation in healthy and damaged white matter," *NMR in Biomedicine*, vol. 23, no. 2, pp. 152–162, 2010. (Cited on pages 39, 45 and 66)

[56] G. T. Balls and L. R. Frank, "A simulation environment for diffusion weighted MR experiments in complex media," *Magnetic Resonance in Medicine*, vol. 62, no. 3, pp. 771–778, 2009. (Cited on pages 39 and 45)

[57] a. Szafer, J. Zhong, and J. C. Gore, "Theoretical model for water diffusion in tissues.," *Magnetic resonance in medicine : official journal of the Society of Magnetic Resonance in Medicine / Society of Magnetic Resonance in Medicine*, vol. 33, no. 5, pp. 697–712, 1995. (Cited on pages 39 and 45)

[58] M. D. Budde and J. A. Frank, "Neurite beading is sufficient to decrease the apparent diffusion coefficient after ischemic stroke," *PNAS*, vol. 107, no. 32, pp. 14472–14477, 2010. (Cited on pages 39, 46, 66 and 104)

[59] 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 *MICCAI*, pp. 404–411, 2010. (Cited on pages 39, 75 and 126)

[60] M. Palombo, D. C. Alexander, and H. Zhang, "A generative model of realistic brain cells with application to numerical simulation of diffusion- weighted MR signal," *NeuroImage*, vol. 188, no. June 2018, pp. 391–402, 2019. (Cited on pages 39 and 85)

[61] J. Rafael-Patino, G. Girard, D. Romascano, M. Barakovic, G. Rensonnet, J.-P. Thiran, and A. Daducci, "Realistic 3D Fiber Crossing Phantom Models for Monte Carlo Diffusion Simulations," in *26th Annual Meeting of the International Society for Magnetic Resonance in Medicine (ISMRM)*, 2018. (Cited on pages 39, 49, 54 and 66)

[62] 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, pp. 1–53, 2019. (Cited on pages 41 and 104)

[63] D. C. Alexander, T. B. Dyrby, M. Nilsson, and H. Zhang, "Imaging brain microstructure with diffusion MRI: Practicality and applications," *NMR in Biomedicine*, no. November 2016, pp. 1–26, 2017. (Cited on pages 41, 53, 66, 104, 139 and 145)

[64] P. Basser, J. Mattiello, and D. LeBihan, "MR diffusion tensor spectroscopy and imaging," *Biophysical Journal*, vol. 66, pp. 259–267, Jan. 1994. (Cited on page 41)

[65] 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*, vol. 53, pp. 1432–1440, June 2005. (Cited on page 41)

[66] 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, pp. 1377–1386, Dec. 2005. (Cited on page 41)

[67] T. Behrens, M. Woolrich, M. Jenkinson, H. Johansen-Berg, R. Nunes, S. Clare, P. Matthews, J. Brady, and S. Smith, "Characterization and propagation of uncertainty in diffusion-weighted MR imaging," *Magnetic Resonance in Medicine*, vol. 50, pp. 1077–1088, Nov. 2003. (Cited on page 41)

[68] Y. Assaf and P. J. Basser, "Composite hindered and restricted model of diffusion (CHARMED) MR imaging of the human brain," *NeuroImage*, vol. 27, pp. 48–58, Aug. 2005. (Cited on page 41)

[69] D. M. Healy, H. Hendriks, and P. T. Kim, "Spherical Deconvolution," *Journal of Multivariate Analysis*, vol. 67, pp. 1–22, Oct. 1998. (Cited on page 42)

[70] 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. (Cited on pages 42, 100 and 123)

[71] D. Alexander, G. Barker, and S. Arridge, "Detection and modeling of non-Gaussian apparent diffusion coefficient profiles in human brain data," *Magnetic Resonance in Medicine*, vol. 48, pp. 331–340, Aug. 2002. (Cited on pages 42 and 123)

[72] F. Dell'Acqua and J. D. Tournier, "Modelling white matter with spherical deconvolution: How and why?," *NMR in Biomedicine*, vol. 32, no. 4, pp. 1–18, 2019. (Cited on pages 42, 100 and 120)

[73] K. G. Schilling, Y. Gao, I. Stepniewska, V. Janve, B. A. Landman, and A. W. Anderson, "Histologically derived fiber response functions for diffusion MRI vary across white matter fibers—An ex vivo validation study in the squirrel monkey brain," *NMR in Biomedicine*, vol. 32, no. 6, pp. 1–17, 2019. (Cited on pages 42 and 120)

[74] D. Raffelt, J. D. Tournier, S. Rose, G. R. Ridgway, R. Henderson, S. Crozier, O. Salvado, and A. Connelly, "Apparent Fibre Density: A novel measure for the analysis of diffusion-weighted magnetic resonance images," *NeuroImage*, vol. 59, no. 4, pp. 3976–3994, 2012. (Cited on pages 42 and 125)

[75] 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. (Cited on pages 42, 75, 100 and 126)

[76] K. G. Schilling, A. Daducci, K. Maier-Hein, C. Poupon, J. C. Houde, V. Nath, A. W. Anderson, B. A. Landman, and M. Descoteaux, "Challenges in diffusion MRI tractography – Lessons learned from international benchmark competitions," *Magnetic Resonance Imaging*, vol. 57, no. October 2018, pp. 194–209, 2019. (Cited on pages 42, 100, 120, 136 and 139)

[77] K. H. Maier-Hein, P. F. Neher, J. C. Houde, M. A. Côté, E. Garyfallidis, J. Zhong, M. Chamberland, F. C. Yeh, Y. C. Lin, Q. Ji, W. E. Reddick, J. O. Glass, D. Q. Chen, Y. Feng, C. Gao, Y. Wu, J. Ma, H. Renjie, Q. Li, C. F. Westin, S. Deslauriers-Gauthier, J. O. O. González, M. Paquette, S. St-Jean, G. Girard, F. Rheault, J. Sidhu, C. M. Tax, F. Guo, H. Y. Mesri, S. Dávid, M. Froeling, A. M. Heemskerk, A. Leemans, A. Boré, B. Pinsard, C. Bedetti, M. Desrosiers, S. Brambati, J. Doyon, A. Sarica, R. Vasta, A. Cerasa, A. Quattrone, J. Yeatman, A. R. Khan, W. Hodges, S. Alexander, D. Romascano, M. Barakovic, A. Auría, O. Esteban, A. Lemkaddem, J. P. Thiran, H. E. Cetingul, B. L. Odry, B. Mailhe, M. S. Nadar, F. Pizzagalli, G. Prasad, J. E. Villalon-Reina, J. Galvis, P. M. Thompson, F. D. S. Requejo, P. L. Laguna, L. M. Lacerda, R. Barrett, F. Dell'Acqua, M. Catani, L. Petit, E. Caruyer, A. Daducci, T. B. Dyrby, T. Holland-Letz, C. C. Hilgetag, B. Stieltjes, and M. Descoteaux, "The challenge of mapping the human connectome based on diffusion tractography," *Nature Communications*, vol. 8, no. 1, 2017. (Cited on pages 42 and 120)

[78] P. F. Neher, F. B. Laun, B. Stieltjes, and K. H. Maier-hein, "Computational Diffusion MRI and Brain Connectivity," no. 1, pp. 105–114, 2014. (Cited on page 42)

[79] S. Jbabdi and H. Johansen-Berg, "Tractography: Where Do We Go from Here?," *Brain Connectivity*, vol. 1, no. 3, pp. 169–183, 2011. (Cited on pages 42 and 120)

[80] 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*, vol. 47, no. 3, pp. 455–460, 2002. (Cited on pages 43 and 45)

[81] S. N. Hwang, C. L. Chin, F. W. Wehrli, and D. B. Hackney, "An image-based finite difference model for simulating restricted diffusion," *Magnetic Resonance in Medicine*, vol. 50, no. 2, pp. 373–382, 2003. (Cited on page 43)

[82] J. Xu, M. D. Does, and J. C. Gore, "Numerical study of water diffusion in biological tissues using an improved finite difference method Numerical study of water diffusion in biological tissues," *Physics in medicine and biology*, vol. 52, pp. N111–N126, 2007. (Cited on page 44)

[83] J. Xu, M. D. Does, and J. C. Gore, "Sensitivity of MR diffusion measurements to variations in intracellular structure: Effects of nuclear size," *Magnetic Resonance in Medicine*, vol. 61, no. 4, pp. 828–833, 2009. (Cited on page 44)

[84] B. F. Moroney, T. Stait-Gardner, B. Ghadirian, N. N. Yadav, and W. S. Price, "Numerical analysis of NMR diffusion measurements in the short gradient pulse limit," *Journal of Magnetic Resonance*, vol. 234, pp. 165–175, 2013. (Cited on page 44)

[85] D. V. Nguyen, J. R. Li, D. Grebenkov, and D. Le Bihan, "A finite elements method to solve the Bloch-Torrey equation applied to diffusion magnetic resonance imaging," *Journal of Computational Physics*, vol. 263, no. FEM, pp. 283–302, 2014. (Cited on page 44)

[86] H.-g. Lipinski, "Monte Carlo simulation of extracellular diffusion in brain tissues Monte Carlo simulation of extracellular diffusion in brain tissues," *Physics in medicine and biology*, vol. 35, p. 441, 1990. (Cited on page 45)

[87] G. J. Stanisz, A. Szafer, G. A. Wright, and R. M. Henkelman, "An Analytical Model of Restricted Diffusion in Bovine Optic-Nerve," *Magnetic Resonance in Medicine*, vol. 37, pp. 103–111, 1997. (Cited on page 45)

[88] A. Duh, A. Mohorič, and J. Stepišnik, "Computer simulation of the spin-echo spatial distribution in the case of restricted self-diffusion," *Journal of Magnetic Resonance*, vol. 148, no. 2, pp. 257–266, 2001. (Cited on page 45)

[89] P. a. Cook, Y. Bai, K. K. Seunarine, M. G. Hall, G. J. Parker, and D. C. Alexander, "Camino: Open-Source Diffusion-MRI Reconstruction and Processing," *14th Scientific Meeting of the International Society for Magnetic Resonance in Medicine*, vol. 14, p. 2759, 2006. (Cited on pages 45, 61, 75, 105 and 126)

[90] J. R. Stiles, D. Van Helden, T. M. Bartol, E. E. Salpeter, and M. M. Salpeter, "Miniature endplate current rise times less than 100 microseconds from improved dual recordings can be modeled with passive acetylcholine diffusion from a synaptic vesicle.," *Proceedings of the National Academy of Sciences*, vol. 93, no. 12, pp. 5747–5752, 1996. (Cited on pages 45 and 86)

[91] J. R. Stiles and T. M. Bartol, *Monte Carlo Methods for Simulating Realistic Synaptic Microphysiology using MCELL*. 2001. (Cited on pages 45 and 86)

[92] R. A. Kerr, T. M. Bartol, B. Kaminsky, M. Dittrich, J.-C. J. Chang, S. B. Baden, T. J. Sejnowski, and J. R. Stiles, "Fast Monte Carlo Simulation Methods for Biological Reaction-Diffusion Systems in Solution and on Surfaces.," *SIAM J Sci Comput*, vol. 30, no. 6, p. 3126, 2008. (Cited on page 45)

[93] G. T. Baxter and L. R. Frank, "A computational model for diffusion weighted imaging of myelinated white matter," *NeuroImage*, vol. 75, pp. 212–220, 2013. (Cited on page 45)

[94] P. N. Sen and P. J. Basser, "A model for diffusion in white matter in the brain," *Biophysical Journal*, vol. 89, no. 5, pp. 2927–2938, 2005. (Cited on page 45)

[95] J. Kärger, H. Pfeifer, and W. Heink, "Principles and Application of Self-Diffusion Measurements by Nuclear Magnetic Resonance," *Advances in Magnetic and Optical Resonance*, vol. 12, no. C, pp. 1–89, 1988. (Cited on page 46)

[96] M. Lin, H. He, G. Schifitto, and J. Zhong, "Simulation of changes in diffusion related to different pathologies at cellular level after traumatic brain injury," *Magnetic Resonance in Medicine*, vol. 76, no. 1, pp. 290–300, 2016. (Cited on page 47)

[97] T. A. G. M. Huisman, L. H. Schwamm, P. W. Schaefer, W. J. Koroshetz, N. Shetty-Alva, Y. Ozsunar, O. Wu, and A. G. Sorensen, "Diffusion Tensor Imaging as Potential Biomarker of White Matter Injury in Diffuse Axonal Injury," *American Journal of Neuroradiology*, vol. 25, no. 3, pp. 370–376, 2004. (Cited on page 47)

[98] J. J. Bazarian, J. Zhong, B. Blyth, T. Zhu, V. Kavcic, and D. Peterson, "Diffusion Tensor Imaging Detects Clinically Important Axonal Damage after Mild Traumatic Brain Injury: A Pilot Study," *Journal of Neurotrauma*, vol. 24, no. 9, pp. 1447–1459, 2007. (Cited on page 47)

[99] D. R. Rutgers, P. Fillard, G. Paradot, M. Tadié, P. Lasjaunias, and D. Ducreux, "Diffusion tensor imaging characteristics of the corpus callosum in mild, moderate, and severe traumatic brain injury," *American Journal of Neuroradiology*, vol. 29, no. 9, pp. 1730–1735, 2008. (Cited on page 47)

[100] M. Lin, H. He, Q. Tong, Q. Ding, X. Yan, T. Feiweier, and J. Zhong, "Effect of myelin water exchange on DTI-derived parameters in diffusion MRI: Elucidation of TE dependence," *Magnetic Resonance in Medicine*, vol. 1660, pp. 1650–1660, 2017. (Cited on page 47)

[101] W. W. Lam, S. Jbabdi, and K. L. Miller, "A model for extra-axonal diffusion spectra with frequency-dependent restriction," *Magnetic Resonance in Medicine*, vol. 73, no. 6, pp. 2306–2320, 2015. (Cited on page 47)

[102] J. Stepišnik, "Time-dependent self-diffusion by NMR spin-echo," *Physica B: Physics of Condensed Matter*, vol. 183, no. 4, pp. 343–350, 1993. (Cited on page 47)

[103] G. Rensonnet, D. Jacobs, B. Macq, and M. Taquet, "A hybrid method for efficient and accurate simulations of diffusion compartment imaging signals," in *Proc. of SPIE*, vol. 9681, p. 968107, 2015. (Cited on pages 47 and 48)

[104] D. Grebenkov, "Laplacian Eigenfunctions in NMR. I. A Numerical Tool," *Concepts in Magnetic Resonance Part A*, vol. 32, no. 4, pp. 277–301, 2008. (Cited on page 47)

[105] G. Rensonnet, B. Scherrer, S. K. Warfield, B. Macq, and M. Taquet, "Assessing the validity of the approximation of diffusion-weighted-MRI signals from crossing fascicles by sums of signals from single fascicles," *Magnetic Resonance in Medicine*, vol. 2345, pp. 2332–2345, 2017. (Cited on pages 47, 66 and 75)

[106] R. A. Salo, I. Belevich, E. Manninen, E. Jokitalo, O. Gröhn, and A. Sierra, "Quantification of anisotropy and orientation in 3D electron microscopy and diffusion tensor imaging in injured rat brain," *NeuroImage*, vol. 172, no. February, pp. 404–414, 2018. (Cited on pages 48, 66, 88 and 104)

[107] M. Andersson, H.-M. Kjer, J. Rafael-Patino, A. Pacureanu, B. Pakkenberg, J.-P. Thiran, M. Ptito, M. Bech, A. B. Dahl, V. A. Dahl, and T. B. Dyrby, "Axon morphology is modulated by the local environment and impacts the non-invasive investigation of its structure-function relationship," *bioRxiv*, 2020. (Cited on pages 48, 116, 120 and 144)

[108] 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. (Cited on pages 49, 54 and 66)

[109] 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, Mar. 2020. (Cited on page 49)

[110] K. Ginsburger, F. Poupon, J. Beaujoin, D. Estournet, F. Matuschke, J.-F. Mangin, M. Axer, and C. Poupon, "Improving the realism of white matter numerical phantoms: A step toward a better understanding of the influence of structural disorders in diffusion MRI," *Frontiers in Physics*, vol. 5, no. FEB, pp. 1–18, 2018. (Cited on pages 49, 54, 64 and 66)

[111] K. Ginsburger, F. Matuschke, F. Poupon, J.-F. Mangin, M. Axer, and C. Poupon, "MEDUSA : A GPU-based tool to create realistic phantoms of the brain microstructure using tiny spheres," *NeuroImage*, vol. 193, no. February, pp. 10–24, 2019. (Cited on pages 49, 54, 66, 72 and 77)

[112] F. Matuschke, K. Ginsburger, C. Poupon, K. Amunts, and M. Axer, "Dense Fiber Modeling for 3D-Polarized Light Imaging Simulations," *arXiv:1901.10284 [physics]*, Jan. 2019. (Cited on page 49)

[113] D. J. Price, A. P. Jarman, J. O. Mason, and P. C. Kind, *Building brains - An introduction to neural development*. 2017. (Cited on pages 54, 66, 68, 71 and 72)

[114] L. A. Lowery and D. V. Vactor, "The trip of the tip: Understanding the growth cone machinery," *Nature Reviews Molecular Cell Biology*, vol. 10, no. 5, pp. 332–343, 2009. (Cited on pages 54, 66 and 68)

[115] K. V. Mardia and P. E. Jupp, *Directional Statistics*. John Wiley & Sons, Ltd., 2008. (Cited on pages 56, 76, 91, 106, 122 and 125)

[116] E. Stockley, H. Cole, A. Brown, and H. Wheal, "A system for quantitative morphological measurement and electronic modelling of neurons: three-dimensional reconstruction," *J Neurosci Methods*, vol. 47, no. 1-2, pp. 39–51, 1993. (Cited on page 58)

[117] R. Callaghan, D. C. Alexander, H. Zhang, and M. Palombo, "Contextual Fibre Growth to Generate Realistic Axonal Packing for Diffusion MRI Simulation," in *Information Processing in Medical Imaging. IPMI2019. Lecture Notes in Computer Science, vol 11492* (A. Chung, J. Gee, P. Yushkevich, and S. Bao, eds.), Springer, Cham, 2019. (Cited on pages 64, 67, 69 and 70)

[118] H. Zhang, T. B. Dyrby, and D. C. Alexander, "Axon diameter mapping in crossing fibers with diffusion MRI.," *Medical image computing and computer-assisted intervention : MICCAI ... International Conference on Medical Image Computing and Computer-Assisted Intervention*, 2011. (Cited on page 66)

[119] H. Zhang, P. L. Hubbard, G. J. M. Parker, and D. C. Alexander, "Axon diameter mapping in the presence of orientation dispersion with diffusion MRI," *NeuroImage*, vol. 56, no. 3, pp. 1301–1315, 2011. (Cited on pages 66 and 145)

[120] P. F. Neher, F. B. Laun, B. Stieltjes, and K. H. Maier-Hein, "Fiberfox: Facilitating the creation of realistic white matter software phantoms," *Magnetic Resonance in Medicine*, vol. 72, no. 5, pp. 1460–1470, 2014. (Cited on page 66)

[121] M. Palombo, C. Ligneul, E. Hernandez-Garzon, and J. Valette, "Can we detect the effect of spines and leaflets on the diffusion of brain intracellular metabolites?," *NeuroImage*, 2018. (Cited on page 66)

[122] J. Brabec, S. Lasič, and M. Nilsson, "Time-dependent diffusion in undulating thin fibers: Impact on axon diameter estimation," *NMR in Biomedicine*, no. February 2019, pp. 1–19, 2019. (Cited on pages 66, 104, 105, 114, 126 and 139)

[123] L. Brusini, G. Menegaz, and M. Nilsson, "Monte Carlo simulations of water exchange through myelin wraps: Implications for diffusion MRI," *IEEE Transactions on Medical Imaging*, vol. 38, no. 6, pp. 1438–1445, 2019. (Cited on pages 66 and 86)

[124] E. W. Dent, S. L. Gupton, and F. B. Gertler, "The growth cone cytoskeleton in Axon outgrowth and guidance," *Cold Spring Harbor Perspectives in Biology*, vol. 3, no. 3, pp. 1–39, 2011. (Cited on page 68)

[125] D. Mortimer, T. Fothergill, Z. Pujic, L. J. Richards, and G. J. Goodhill, "Growth cone chemotaxis," *Trends in Neurosciences*, vol. 31, no. 2, pp. 90–98, 2008. (Cited on page 68)

[126] F. Polleux and W. Snider, "Initiating and growing an axon.," *Cold Spring Harbor perspectives in biology*, vol. 2, no. 4, pp. 1–20, 2010. (Cited on page 68)

[127] P. Rauch, P. Heine, B. Goettgens, and J. A. Käs, "Different modes of growth cone collapse in NG 108-15 cells," *European Biophysics Journal*, vol. 42, no. 8, pp. 591–605, 2013. (Cited on pages 68 and 70)

[128] T. Sakisaka and Y. Takai, "Cell adhesion molecules in the CNS," *Journal of Cell Science*, vol. 118, no. 23, pp. 5407–5410, 2005. (Cited on pages 68 and 71)

[129] D. Šmít, C. Fouquet, F. Pincet, M. Zapotocky, and A. Trembleau, "Axon tension regulates fasciculation/defasciculation through the control of axon shaft zippering," *eLife*, vol. 6, pp. 1–49, 2017. (Cited on pages 68 and 73)

[130] J. Barry, Y. Gu, and C. Gu, "Polarized targeting of L1-CAM regulates axonal and dendritic bundling in vitro," *European Journal of Neuroscience*, vol. 32, no. 10, pp. 1618–1631, 2010. (Cited on page 73)

[131] A. G. Voyiadjis, M. Doumi, E. Curcio, and T. Shinbrot, "Fasciculation and defasciculation of neurite bundles on micropatterned substrates," *Annals of Biomedical Engineering*, vol. 39, no. 1, pp. 559–569, 2011. (Cited on page 73)

[132] D. C. Van Essen, K. Ugurbil, E. Auerbach, D. Barch, T. E. Behrens, R. Bucholz, A. Chang, L. Chen, M. Corbetta, S. W. Curtiss, S. Della Penna, D. Feinberg, M. F. Glasser, N. Harel, A. C. Heath, L. Larson-Prior, D. Marcus, G. Michalareas, S. Moeller, R. Oostenveld, S. E. Petersen, F. Prior, B. L. Schlaggar, S. M. Smith, A. Z. Snyder, J. Xu, and E. Yacoub, "The Human Connectome Project: A data acquisition perspective," *NeuroImage*, vol. 62, no. 4, pp. 2222–2231, 2012. (Cited on pages 74, 75 and 122)

[133] S. N. Sotiropoulos, S. Jbabdi, J. Xu, J. L. Andersson, S. Moeller, E. J. Auerbach, M. F. Glasser, M. Hernandez, G. Sapiro, M. Jenkinson, D. A. Feinberg, E. Yacoub, C. Lenglet, D. C. Van Essen, K. Ugurbil, and T. E. Behrens, "Advances in diffusion MRI acquisition and processing in the Human Connectome Project," *NeuroImage*, vol. 80, pp. 125–143, 2013. (Cited on pages 75 and 138)

[134] H.-h. Lee, A. Papaioannou, S.-L. Kim, D. S. Novikov, and E. Fieremans, "A time-dependent diffusion MRI signature of axon caliber variations and beading," *Communications Biology*, no. 2020, 2019. (Cited on pages 75, 107, 114, 116, 120 and 144)

[135] E. Fieremans and H. H. Lee, "Physical and numerical phantoms for the validation of brain microstructural MRI: A cookbook," *NeuroImage*, vol. 182, no. June, pp. 39–61, 2018. (Cited on page 75)

[136] 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*, vol. 202, no. January, p. 116137, 2019. (Cited on pages 75, 126 and 138)

[137] S. N. Sotiropoulos, S. Moeller, S. Jbabdi, J. Xu, J. L. Andersson, E. J. Auerbach, E. Yacoub, D. Feinberg, K. Setsompop, L. L. Wald, T. E. Behrens, K. Ugurbil, and C. Lenglet, "Effects of image reconstruction on fiber orientation mapping from multichannel diffusion MRI: Reducing the noise floor using SENSE," *Magnetic Resonance in Medicine*, vol. 70, no. 6, pp. 1682–1689, 2013. (Cited on pages 81, 122 and 138)

[138] M. Palombo, C. Ligneul, C. Najac, J. Le Douce, J. Flament, C. Escartin, P. Hantraye, E. Brouillet, G. Bonvento, and J. Valette, "New paradigm to assess brain cell morphology by diffusion-weighted MR spectroscopy in vivo," *Proceedings of the National Academy of Sciences*, vol. 113, no. 24, pp. 6671–6676, 2016. (Cited on page 86)

[139] C. Ophus, "A fast image simulation algorithm for scanning transmission electron microscopy," *Advanced Structural and Chemical Imaging*, vol. 3, no. 1, pp. 1–11, 2017. (Cited on page 86)

[140] L. Grella, G. Lorusso, and D. L. Adler, "Simulations of Scanning Electron Microscopy Imaging and Charging of Insulating Structures," *Scanning*, vol. 25, no. 6, pp. 300–308, 2003. (Cited on page 86)

[141] S. Babin, S. S. Borisov, H. Ito, A. Ivanchikov, and M. Suzuki, "Simulation of scanning electron microscope images taking into account local and global electromagnetic fields," *Journal of Vacuum Science & Technology B, Nanotechnology and Microelectronics: Materials, Processing, Measurement, and Phenomena*, vol. 28, no. 6, pp. C6C41–C6C47, 2010. (Cited on page 86)

[142] F. Matuschke, "Dense Fiber Modeling for 3D-Polarized Light Imaging Simulations," *ArXiv*, 2019. (Cited on page 86)

[143] M. Menzel, K. Michielsen, H. De Raedt, J. Reckfort, K. Amunts, and M. Axer, "A Jones matrix formalism for simulating three-dimensional polarized light imaging of brain tissue," *Journal of the Royal Society Interface*, vol. 12, no. 111, 2015. (Cited on page 86)

[144] R. A. Kerr, T. M. Bartol, B. Kaminsky, M. Dittrich, J. C. J. Chang, S. B. Baden, T. J. Sejnowski, and J. R. Stiles, "Fast Monte Carlo simulation methods for biological reaction-diffusion systems in solution and on surfaces," *SIAM Journal on Scientific Computing*, vol. 30, no. 6, pp. 3126–3149, 2008. (Cited on page 86)

[145] S. Plimpton, "Short-Range Molecular Dynamics," *Journal of Computational Physics*, vol. 117, no. 6, pp. 1–42, 1997. (Cited on page 86)

[146] M. Helmstaedter, K. L. Briggman, and W. Denk, "3D structural imaging of the brain with photons and electrons," *Current Opinion in Neurobiology*, vol. 18, no. 6, pp. 633–641, 2008. (Cited on page 88)

[147] W. Denk and H. Horstmann, "Serial block-face scanning electron microscopy to reconstruct three-dimensional tissue nanostructure," *PLoS Biology*, vol. 2, no. 11, 2004. (Cited on page 88)

[148] R. S. Womersley, "Efficient spherical designs with good geometric properties," *Contemporary Computational Mathematics - A Celebration of the 80th Birthday of Ian Sloan*, pp. 1243–1285, 2018. (Cited on page 89)

[149] C. T. Rueden, J. Schindelin, M. C. Hiner, B. E. DeZonia, A. E. Walter, E. T. Arena, and K. W. Eliceiri, "ImageJ2: ImageJ for the next generation of scientific image data," *BMC Bioinformatics*, vol. 18, no. 1, pp. 1–26, 2017. (Cited on page 90)

[150] D. Legland, I. Arganda-Carreras, and P. Andrey, "MorphoLibJ: Integrated library and plugins for mathematical morphology with ImageJ," *Bioinformatics*, vol. 32, no. 22, pp. 3532–3534, 2016. (Cited on page 90)

[151] J. Schindelin, I. Arganda-Carreras, E. Frise, V. Kaynig, M. Longair, T. Pietzsch, S. Preibisch, C. Rueden, S. Saalfeld, B. Schmid, J. Y. Tinevez, D. J. White, V. Hartenstein, K. Eliceiri, P. Tomancak, and A. Cardona, "Fiji: An open-source platform for biological-image analysis," *Nature Methods*, vol. 9, no. 7, pp. 676–682, 2012. (Cited on page 90)

[152] C. A. Schneider, W. S. Rasband, and K. W. Eliceiri, "NIH Image to ImageJ: 25 years of image analysis," *Nature Methods*, vol. 9, no. 7, pp. 671–675, 2012. (Cited on page 90)

[153] E. G. Baxi, J. DeBruin, D. M. Tosi, I. V. Grishkan, M. D. Smith, L. A. Kirby, H. J. Strasburger, A. N. Fairchild, P. A. Calabresi, and A. R. Gocke, "Transfer of myelin-reactive Th17 cells impairs endogenous remyelination in the central nervous system of cuprizone-fed mice," *Journal of Neuroscience*, vol. 35, no. 22, pp. 8626–8639, 2015. (Cited on page 93)

[154] P. J. Paine, S. P. Preston, M. Tsagris, and A. T. Wood, "An elliptically symmetric angular Gaussian distribution," *Statistics and Computing*, vol. 28, no. 3, pp. 689–697, 2018. (Cited on page 94)

[155] 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, July 2018. (Cited on page 104)

[156] E. Kaden, F. Kruggel, and D. C. Alexander, "Quantitative mapping of the per-axon diffusion coefficients in brain white matter: Quantitative Mapping of the Per-Axon Diffusion Coefficients," *Magnetic Resonance in Medicine*, vol. 75, pp. 1752–1763, Apr. 2016. (Cited on pages 104 and 137)

[157] E. Panagiotaki, T. Schneider, B. Siow, M. G. Hall, M. F. Lythgoe, and D. C. Alexander, "Compartment models of the diffusion MR signal in brain white matter: A taxonomy and comparison," *NeuroImage*, vol. 59, pp. 2241–2254, Feb. 2012. (Cited on page 104)

[158] L. J. Edwards, K. J. Pine, I. Ellerbrock, N. Weiskopf, and S. Mohammadi, "NODDI-DTI: Estimating Neurite Orientation and Dispersion Parameters from a Diffusion Tensor in Healthy White Matter," *Frontiers in Neuroscience*, vol. 11, p. 720, Dec. 2017. (Cited on page 104)

[159] M. Andersson, H. M. Kjer, J. Rafael-Patino, A. Pacureanu, B. Pakkenberg, J.-P. Thiran, M. Ptito, M. Bech, A. Bjorholm Dahl, V. Andersen Dahl, and T. B. Dyrby, "Axon morphology is modulated by the local environment and impacts the noninvasive investigation of its structure–function relationship," *Proceedings of the National Academy of Sciences*, vol. 117, pp. 33649–33659, Dec. 2020. (Cited on page 104)

[160] B. Dhital, M. Reisert, E. Kellner, and V. G. Kiselev, "Intra-axonal diffusivity in brain white matter," *NeuroImage*, vol. 189, pp. 543–550, Apr. 2019. (Cited on pages 104, 117 and 126)

[161] J. Veraart, E. Fieremans, and D. S. Novikov, "On the scaling behavior of water diffusion in human brain white matter," *NeuroImage*, vol. 185, pp. 379–387, Jan. 2019. (Cited on page 104)

[162] E. T. McKinnon, J. A. Helpern, and J. H. Jensen, "Modeling white matter microstructure with fiber ball imaging," *NeuroImage*, vol. 176, pp. 11–21, Aug. 2018. (Cited on page 104)

[163] A. Howard, A. Khrapitchev, J. Mollink, R. Mars, N. Sibson, J. Sallet, S. Jbabdi, and K. L. Miller, "Estimating intra-axonal axial diffusivity with diffusion MRI in the presence of fibre orientation dispersion," in *ISMRM Virtual Conference & Exhibition*, 2020. (Cited on pages 104, 108, 116 and 117)

[164] D. S. Novikov, J. H. Jensen, J. A. Helpern, and E. Fieremans, "Revealing mesoscopic structural universality with diffusion," *Proceedings of the National Academy of Sciences of the United States of America*, vol. 111, no. 14, pp. 5088–5093, 2014. (Cited on page 114)

[165] M. Catani and M. Thiebaut de Schotten, *Atlas of Human Brain Connections*. 2013. (Cited on page 120)

[166] H. Johansen-Berg and T. E. Behrens, "Just pretty pictures? What diffusion tractography can add in clinical neuroscience," 2006. (Cited on page 120)

[167] C. Thomas, F. Q. Ye, M. O. Irfanoglu, P. Modi, K. S. Saleem, D. A. Leopold, and C. Pierpaoli, "Anatomical accuracy of brain connections derived from diffusion MRI tractography is inherently limited," *Proceedings of the National Academy of Sciences*, vol. 111, pp. 16574–16579, Nov. 2014. (Cited on page 120)

[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, Oct. 2020. (Cited on pages 121, 122 and 128)

[169] C. Brechbühler, G. Gerig, and O. Kübler, "Parametrization of Closed Surfaces for 3-D Shape Description," *Computer Vision and Image Understanding*, vol. 61, no. 2, pp. 154–170, 1995. (Cited on page 123)

[170] D. A. Raffelt, J. D. Tournier, R. E. Smith, D. N. Vaughan, G. Jackson, G. R. Ridgway, and A. Connelly, "Investigating white matter fibre density and morphology using fixel-based analysis," *NeuroImage*, vol. 144, pp. 58–73, 2017. (Cited on page 125)

[171] H.-H. Lee, A. Papaioannou, S.-L. Kim, D. S. Novikov, and E. Fieremans, "Probing axonal swelling with time dependent diffusion MRI," *ArXiv*, 2020. (Cited on pages 125, 126, 128, 136 and 139)

[172] E. B. Saff and A. B. J. Kuijlaars, "Distributing many points on a sphere," *The Mathematical Intelligencer*, vol. 19, pp. 5–11, Dec. 1997. (Cited on page 126)

[173] D. Christiaens, S. Sunaert, P. Suetens, and F. Maes, "Convexity-constrained and nonnegativity-constrained spherical factorization in diffusion-weighted imaging," *NeuroImage*, vol. 146, pp. 507–517, Feb. 2017. (Cited on page 137)

[174] M. Cottaar, M. Bastiani, N. Boddu, M. F. Glasser, S. Haber, D. C. van Essen, S. N. Sotiropoulos, and S. Jbabdi, "Modelling white matter in gyral blades as a continuous vector field," *NeuroImage*, vol. 227, p. 117693, Feb. 2021. (Cited on page 138)

[175] D. C. Van Essen, S. Jbabdi, S. N. Sotiropoulos, C. Chen, K. Dikranian, T. Coalson, J. Harwell, T. E. Behrens, and M. F. Glasser, "Mapping Connections in Humans and Non-Human Primates," in *Diffusion MRI*, pp. 337–358, Elsevier, 2014. (Cited on page 138)

[176] C. H. Yeh, J. D. Tournier, K. H. Cho, C. P. Lin, F. Calamante, and A. Connelly, "The effect of finite diffusion gradient pulse duration on fibre orientation estimation in diffusion MRI," *NeuroImage*, vol. 51, no. 2, pp. 743–751, 2010. (Cited on page 137)

[177] D. Christiaens, J. Veraart, L. Cordero-Grande, A. N. Price, J. Hutter, J. V. Hajnal, and J. D. Tournier, "On the need for bundle-specific microstructure kernels in diffusion MRI," *NeuroImage*, vol. 208, no. November 2019, p. 116460, 2020. (Cited on page 139)

[178] C. F. Westin, H. Knutsson, O. Pasternak, F. Szczepankiewicz, E. Özarslan, D. van Westen, C. Mattisson, M. Bogren, L. J. O'Donnell, M. Kubicki, D. Topgaard, and M. Nilsson, "Q-space trajectory imaging for multidimensional diffusion MRI of the human brain," *NeuroImage*, vol. 135, pp. 345–362, 2016. (Cited on page 145)

[179] L. Kerkelä, F. Nery, R. Callaghan, F. Zhou, N. G. Gyori, F. Szczepankiewicz, M. Palombo, G. J. Parker, H. Zhang, M. G. Hall, and C. A. Clark, "Comparative analysis of signal models for microscopic fractional anisotropy estimation using q-space trajectory encoding," *NeuroImage*, vol. 242, p. 118445, Nov. 2021. (Cited on page 145)

[180] C. Ericson, *Real-Time Collision Detection*. CRC Press, zeroth ed., Dec. 2004. (Cited on page 146)