Enhancing Microstructural Specificity in Diffusion MRI Analysis



Hossein Rafipoor

Wolfson College

University of Oxford

A thesis submitted for the degree of

Doctor of Philosophy

Trinity 2023

Acknowledgements

Apparently, I am on the verge of receiving a doctorate degree from Oxford University. Hold on, it's not a PhD, though; it's a D.Phil. My supervisor told me this just after I started, but he also assured me they are the same. Fingers crossed! Whatever it is, if I claim all the glory for it, the Oxford team might just realize their mistake and retract it back. After all, it's a fundamental scientific principle that substantial outcomes like this (if I may!), are the result of numerous contributing factors and not isolated incidents. Some of these factors I briefly mention below.

The first principal component, my family; firstly because they always let me do what I pleased. And what did I do with all this liberty? Developing a passion for science and reading books; Nothing says 'wild child' like a pulling all-nighters with intractable math questions! Also, in light of the terrible economic situation in Iran, their task was harder than solving a quadratic equation without a calculator, but they always made sure I had all I needed to live a good life, without letting me to feel the pressure. I am deeply thankful for all their love and support. And let's not forget my sister, my constant partner in crime, always there to keep things lively with sibling shenanigans and for stepping into my shoes for all family events and chores when work deadlines kept me busy. Special thanks also to my Aunt Fatemeh Chamandar - her support in several stages of my life and her inspirational qualities deserve an acknowledgement of their own.

The second component, my supervisors; trust me, it was not easy to work with them! It felt like having to explain calculus 101 to Einstein and Feynman every single week for three years. First up, Saad Jbabdi, always prepared to tackle intricate mathematics or coding challenges; while also being a grand master of big-picture thinking. He has always been there, even in his hardest times, to support and cheer me up with his sometime too blunt jokes! Nobody would ever believe those conversations was between me and my supervisor. Next up, Michiel Cottaar, the most genius person I've ever had the luck to work with. Not only did he generously shared his vast wealth of knowledge and skills to elevate my work and codes with incredible patience and top-quality explanations, but he also shared the results of his cooking courses, introduced me to his wonderful circle of friends, and led me on bike trips intense enough to forget about my hard feelings, especially in my cloudier days. My gratitude to both of them is boundless!

Next, to my colleagues, the smartest and kindest companions one could wish for - Daniel Korr, Ying Qiu Zhang, Fredrik Lange (i.e. the transformer), Paul McCarthy (i.e. the lord of scripts), Clemence Ligneul, Weikang Gong, Xi Chen, Miriam Szabo, Nicola Dinsdale, Istvan Huszar - our chats have been a source of pleasure and enlightening. Also, to my side panel of advisors and commentators - Amy Howard, Jesper Andersson, Mark Woolrich, Steve Smith - thanks for your priceless inputs. It's been a tremendous privilege to work alongside such brilliant intellects and genuine souls as you are.

I would also like to express my heartfelt appreciation to my viva examiners, Prof. Karla Miller and Dr. Chantal Tax, for taking on the hard task of reading my thesis and offering invaluable feedback. The discussions we had during my viva session were an absolute pleasure and provided me with invaluable learning experiences.

I cannot wrap this up without a nod to my funders at Wellcome Trust, the wonderful entities behind the scenes that sponsored this amazing journey. Without your generous support and hard work, turning my dream into this incredible experience wouldn't have been possible, thanks for putting your trust on me. I sincerely hope I've been a worthy investment, and I am committed to paying this forward in the most impactful way possible.

The FMRIB squad - what an exceptional team, always on standby with IT rescue missions when SSH decided to play hide and seek. The facilities team was always there to ensure everything ran smoothly, except for the tap water, and whiteboards – but hey, we can't all be perfect, I am really grateful to them, much more than I am for the HR team. More importantly, the culture that's been brewing in WIN over the years is genuinely a sparkling gem in the vast expanse of the world, nurturing not just academic growth in its best way, but also fostering all values of a great community. I've seen loads of tireless work being put into promoting inclusivity and equality here, so hats off to all these shining stars lighting up our prosperous corner of the universe! Also, a big high-five to Carinne from the public engagement team, who made the most cheerful noises when needed.

A detour now to the high school neuroscience program that set me on this fascinating path, made by a group of enthusiastic fellow Iranian young scientists - Mohsen Omrani, Maryam Vaziri, Arash Afraz, Baktash Babadi, and others - who left indelible marks on my career and mindset since then. Also, a hearty shout-out to Babak Araabi, Reza Rajimehr, Hamed Nili, and Rouzbeh Farhoudi; working with them prior to my PhD has been a great source of aspiration. None of this would've been possible without having the bless to know Abdolhossein Abbasian, whose inspiration reaches not just me, but the entire community of Iranian neuroscientists.

During my time here, I also had the opportunity to rub elbows with the ISMRM diffusion community, though mostly on a 2-d lattice, including some diffusion gurus

like Derek Jones, Dimitry Novikov, Markus Nilsson, Garry Zhang, and Marco Palombo. Thanks for letting me into your circle, the enlightening conversations, and bearing with the newbies like me!

Back in Iran, I owe a huge debt of gratitude to Saleh Hashemseresht, who rescued me from losing myself in chaos on several occasions. I also want to express my heartfelt thanks to my incredible team of remote supporters, including Bahar Sadeghi, Ata Mohammadi, Soheil Golara, Omid Sarayani, and Saghar Mirbagheri. Your unwavering support from afar has always been deeply appreciated.

In Oxford, I found a second family in the amazing individuals who include Maryam Aslani, Navid Aslani, Parisa Mansouri, Faranak Tavakolian, Emese Vegh, Phyllis Ferguson, and Rana Dasgupta. Despite my less-than-stellar ability to plan my own life, they always made sure I had a place to call home, helping me survive and thrive throughout this journey.

Now, it's hard to ignore the elephant in the room: the world's been on a bit of a roller-coaster ride since my D.Phil. journey started. A pesky little thing called COVID19 made a global pandemic that changed everyone's life. Picture this: weeks where I didn't see a single soul, flipped down sleeping schedules, and walking through hauntingly empty streets of Oxford! Enormous thanks to all the Wolfson College staff, who braved the storm to keep the gears running in the hardest days.

A heartfelt thanks to the people of Britain too. Despite the British weather, your respect and care made my experience as an international student here very comforting; well, obviously that does not include the food!

Also, a standing shout out to Chat GPT - the newborn AI companion that came out just at the right time to help me rewrite, proofread, and probably saved me from countless grammar mistakes. So here's to you, GPT, may you continue to save future generations of students from typos and syntax errors, while making sure not to take over our jobs!

And on a more solemn note - an ironic salute goes to the Iranian regime, who during my time here, continued to make storms of tragedy and unrest at home. From shooting down a passenger flight, claiming the lives of 176 innocent individuals, among them some of my beloved former classmates, Iman aghabali, Elnaz Nabiyi and Milad Ghasemi, to the brutal taking of young lives in the streets, and several executions based on ideological disputes or sham trials. These tragic events, while deeply painful, served as a reminder of a world of work that awaits, goals that need pursuing, and changes to be made.

Now, as I close this chapter and look to the future, I hope to carry forward this spirit of collaboration, camaraderie, and shared discovery that I gained. Here's to the next adventure!

Abstract

Enclosed within a protective bony shell, the human brain poses a significant challenge for direct physical examination. Therefore, non-invasive methods, such as Magnetic Resonance Imaging (MRI), have become very important tools for exploring the brain's structure and functionality. Nevertheless, these indirect techniques are not without their limitations, notably a lack of specificity. Although they are sensitive to microstructural variations, their capacity to link changes in signals to meaningful biophysical processes is limited.

Microstructural modelling has risen to address this problem, offering a means to correlate these otherwise abstract measurements with meaningful biophysical parameters. However, as the complexity of these models, either in terms of geometry or physical properties, improves, so does the number of free parameters required, posing a substantial challenge for reliable parameter estimation.

This thesis seeks to address two substantial issues within the field of microstructural modelling of diffusion MRI data. Firstly, it proposes a novel approach to identify the microstructural parameter alterations that can explain observed differences in diffusion MRI between diseased and control group. This approach enables the application of highly specific models without the associated concerns about parameter estimation. Secondly, it advocates for the development of a tool that extracts fibre-specific features from diffusion MRI data in a manner that promotes comparability across different subjects, facilitated by the use of hierarchical Bayesian models.

By offering these approaches to analyze diffusion MRI data, this research aims to circumventing the constraints imposed by existing microstructural modeling techniques, thus improve the precision of brain structure diagnosis and comprehension.

Contents

List of Figures **ix**

1 Introduction **1**

2 Background **4**

2.1 White matter . . . 5
2.1.1 Why white matter matters? . . . 5
2.1.2 Microstructure . . . 6
2.1.3 Anatomy . . . 7
2.1.4 Methods for studying the white matter . . . 8
2.1.5 White matter pathologies . . . 9
2.2 Diffusion MRI . . . 11
2.2.1 Acquisition . . . 11
2.2.1.1 MRI . . . 11
2.2.1.2 Diffusion . . . 12
2.2.1.3 Pulsed Gradient Spin Echo . . . 13
2.2.1.4 Diffusion weighted imaging . . . 16
2.2.2 Data analysis . . . 18
2.2.2.1 Preprocessing . . . 18
2.2.2.2 Diffusion Tensor metrics . . . 19
2.2.2.3 Multicompartment models . . . 21
2.2.2.4 Degeneracy problem . . . 22
2.2.2.5 Group analysis of diffusion MRI data . . . 23
2.2.2.6 Tractography . . . 26
2.3 Bayesian Modelling . . . 27
2.3.1 Bayes rule . . . 28
2.3.2 Computation problem . . . 29
2.3.3 Prior selection . . . 31
2.3.4 Why Bayesian modelling? . . . 32

3 Bayesian Estimation of Change **34**
3.1 Abstract . . . 35
3.2 Introduction . . . 36
3.3 Theory . . . 39
3.3.1 Inference on change in parameters . . . 39
3.3.2 Training models of change . . . 42
3.3.3 Biophysical model of diffusion . . . 45
3.3.4 Summary measurements . . . 47
3.4 Methods . . . 49
3.4.1 Simulations . . . 49
3.4.2 White matter hyperintensities . . . 52
3.5 Results . . . 53
3.5.1 Summary measurements . . . 53
3.5.2 Validation . . . 54
3.5.2.1 Comparison with model inversion . . . 54
3.5.2.2 Sensitivity to change in each parameter . . . 56
3.5.2.3 Estimating the amount of change . . . 56
3.5.3 White matter hyperintensities . . . 59
3.5.3.1 Model Inversion . . . 59
3.5.3.2 BENCH . . . 59
3.6 Discussion . . . 64
3.6.1 Comparison with model inversion . . . 64
3.6.2 Summary measures . . . 67
3.6.3 Future developments . . . 68
3.7 Software . . . 69
3.8 Appendix A. Toy example: Inferring changes in 2D . . . 69
3.9 Appendix B. Estimating Quality of Fit . . . 70

4 Hierarchical Modelling of Crossing Fibres in the White Matter **73**
4.1 Abstract . . . 74
4.2 Introduction . . . 74
4.3 Theory . . . 76
4.3.1 Crossing fibre model . . . 76
4.3.2 Hierarchical model . . . 78
4.3.3 Parameter Inference . . . 80
4.3.3.1 Determining the number of fibres in each voxel . . . 82
4.3.4 Fixel based threshold free cluster enhancement . . . 82
4.4 Methods . . . 85
4.4.1 Simulation settings . . . 85

4.4.2 Data sampling and voxel correspondence 85
4.4.3 Fixel template creation 86
4.4.4 GLM analysis 87
4.5 Results 87
4.5.1 Simulations 87
4.5.2 White matter fixel template 88
4.5.3 Fixel based GLM 93
4.6 Software implementation 98
4.7 Discussion 98
4.7.1 Fixel template creation and usage 99
4.7.2 Future Directions 101

5 Conclusion and Future Directions **103**
5.1 Bayesian estimation of change 105
5.2 Hierarchical modelling of crossing fibres in the white matter 107

References **109**

List of Figures

2.1 Illustration of white matter microstructure 6
2.2 The gradient spin echo pulse sequence. 15
2.3 Sample Diffusion Data 17
2.4 DTI Metrics 20
2.5 Model Degeneracy 23
2.6 Fixel-Based Analysis 25
2.7 fODF estimation 27

3.1 Illustration of the inversion-free inference on change (BENCH) . . . 37
3.2 Distribution of gradients. 43
3.3 Flowchart for training and inference using change models. 44
3.4 Compartments of the diffusion model. 48
3.5 Prior distributions for the parameters of the standard model 51
3.6 Maps of the summary measurements 54
3.7 The estimated amount of change in the summary measurements . . 55
3.8 Confusion matrices 57
3.9 Estimated amount of change in each parameter 58
3.10 NODDI parameter estimates 61
3.11 estimated amount of change in the measurements 62
3.12 Estimated change in each voxel 63
3.13 Parameter estimation 71
3.14 Distribution of differences for the correct change model 72

4.1 Axonal loss simulation 77
4.2 Generative forward model of diffusion MRI signals 78
4.3 Template creation pipeline 83
4.4 Fixel neighbourhood for TFCE 84
4.5 Quality of fit for parameter estimation 89
4.6 White matte fixel template 90
4.7 Fibre orientation estimates for an exemplar crossing fibre voxel . . . 91
4.8 Variability of the primary fiber population across subjects 92
4.9 Determining the number of fixels present in each voxel 93

4.10 Fixel t-statistics 94
4.11 Positive FA changes 95
4.12 Fixel p-values 96
4.13 Voxelwise vs. fixel-specific changes 97

The good life is one inspired by love and guided by knowledge. Neither love without knowledge nor knowledge without love can produce a good life.

— Bertrand Russell

1 Introduction

The brain is a critical organ, so that it is uniquely encapsulated in a complete bony shell, providing essential protection. This protective barrier introduces substantial challenges to direct physical access, at least while the owner is alive. As a result, non-invasive techniques such as Magnetic Resonance Imaging (MRI) [1] become essential tools. MRI is capable of probing brain structure and function in various ways[2–6]. However, the indirect nature of it, which relies on measuring magnetic fields, gives rise to its own set of limitations, the most significant of which is a lack of specificity. MRI measurements per-se are abstract and often devoid of meaningful interpretation despite their high sensitivity to changes in the brain.

Diffusion MRI (dMRI) specifically aims to probe brain microstructure by measuring water diffusion dynamics through the manipulation of the magnetic field.[3, 7, 8] Like other MRI modalities, dMRI is sensitive to changes but also notably unspecific. Directly linking an observed change in diffusion MRI signals to a meaningful biophysical alteration is far from straightforward. To address this gap, the field of microstructural modelling has evolved, producing numerous models that attempt to correlate these measurements with meaningful biophysical parameters [9–13]. The ultimate goal is to aid clinicians and researchers in diagnosing or understanding brain structure.

Despite their utility, these models present inherent challenges. A notable trade-off persists in the field: as the specificity of the models increases, so does the number of free parameters required to explain the data.[14] Consequently, a major challenge lies in reliably estimating these parameters to investigate microstructural changes resulting from various biological processes.To overcome this issue, researchers often incorporate more prior assumptions, which may lack validation, or acquire additional complementary data, which can be costly and not always feasible[9]. In this regard, the first contribution of this thesis is developing a new approach to address a more pertinent question: which microstructural parameter changes can best explain observed differences in diffusion MRI between two groups (e.g. patients vs controls)? This approach allows for the utilization of more specific models without the accompanying concern about parameter estimation.

Another challenge in microstructural modelling stems from the intricate architecture of white matter, which comprises millions of intertwined fiber bundles. While diffusion MRI data provides volumetric information about white matter microstructure at a millimetre-scale resolution, extracting fiber-specific features necessitates modelling. Numerous models exist that extract orientation and relative volume fraction of the fiber populations present within each brain location from diffusion MRI data.[15–17] However, individual variations in the location and orientation of fiber populations within the brain complicate the comparison of the same tracts across different individuals. The second contribution of this thesis is development of a new tool for extraction of fibre-specific information in a manner that enables meaningful comparisons across subjects. We employed hierarchical Bayesian models, which offers nested modelling of data from multiple subjects. This hierarchical framework offers two main advantages: it enables correspondence in model parameters across subjects, thereby allowing for comparability, and it provides robust parameter estimation for individual subjects, particularly for those with noisy data.

The second chapter of this thesis presents the relevant background of the research field including white matter research, diffusion MRI and Bayesian modelling. The

third chapter provides a description of the first developed method that addresses the identification of changes in microstructure from diffusion MRI data. The fourth chapter details the method for extracting fiber-specific metrics from diffusion data for the aim of group comparison. Lastly, the fifth chapter briefly discusses the limitations and future directions of the research presented in this thesis.

If I have seen further, it is by standing on the shoulders of Giants.

—Isaac Newton

2 Background

Contents

2.1 White matter 5
2.1.1 Why white matter matters? 5
2.1.2 Microstructure 6
2.1.3 Anatomy 7
2.1.4 Methods for studying the white matter 8
2.1.5 White matter pathologies 9
2.2 Diffusion MRI 11
2.2.1 Acquisition 11
2.2.1.1 MRI 11
2.2.1.2 Diffusion 12
2.2.1.3 Pulsed Gradient Spin Echo 13
2.2.1.4 Diffusion weighted imaging 16
2.2.2 Data analysis 18
2.2.2.1 Preprocessing 18
2.2.2.2 Diffusion Tensor metrics 19
2.2.2.3 Multicompartment models 21
2.2.2.4 Degeneracy problem 22
2.2.2.5 Group analysis of diffusion MRI data 23
2.2.2.6 Tractography 26
2.3 Bayesian Modelling 27
2.3.1 Bayes rule 28
2.3.2 Computation problem 29
2.3.3 Prior selection 31
2.3.4 Why Bayesian modelling? 32

2.1 White matter

2.1.1 Why white matter matters?

My ability to write this thesis, as well as your ability to read it, are both the result of billions of tiny processes carried out by the neurons in our brains. Although each neuron is a complex system that can have an extensive dynamic range of behaviour, such complex tasks are only possible by efficient interaction between billions of neurons scattered across the brain. The neuronal cell bodies are located in the outer surface of the brain or deep inside the brain (grey matter), and the axon bundles that connect them, just like wires, are placed inside the brain (white matter), as shown in Figure 2.1. These axons are coated with myelin, like the shielding of a wire, that is a fatty substance that gives white matter its characteristic colour.

The study of brain white matter holds significance for two primary reasons - clinical implications and understanding human brain functions. Firstly, abnormalities in white matter have been associated with a broad spectrum of neurological and psychiatric disorders, such as Alzheimer's disease, multiple sclerosis, depression, anxiety disorders, and schizophrenia.[18, 19] Measuring the structure and function of white matter in these conditions can provide insights into their underlying mechanisms of degeneration and promote the development of more efficient treatments.[20–22]

Additionally, understanding brain white matter is vital for comprehending normal brain function and the underlying mechanisms of brain activity. For instance, during the early stages of brain development, from infancy to childhood, white matter undergoes substantial changes.[23, 24]Tracking and quantifying these alterations can significantly contribute to our knowledge of brain development and healthy brain functioning. This knowledge later can facilitate the development of more effective interventions that target white matter pathways to enhance cognitive abilities and improve the quality of life for numerous individuals.

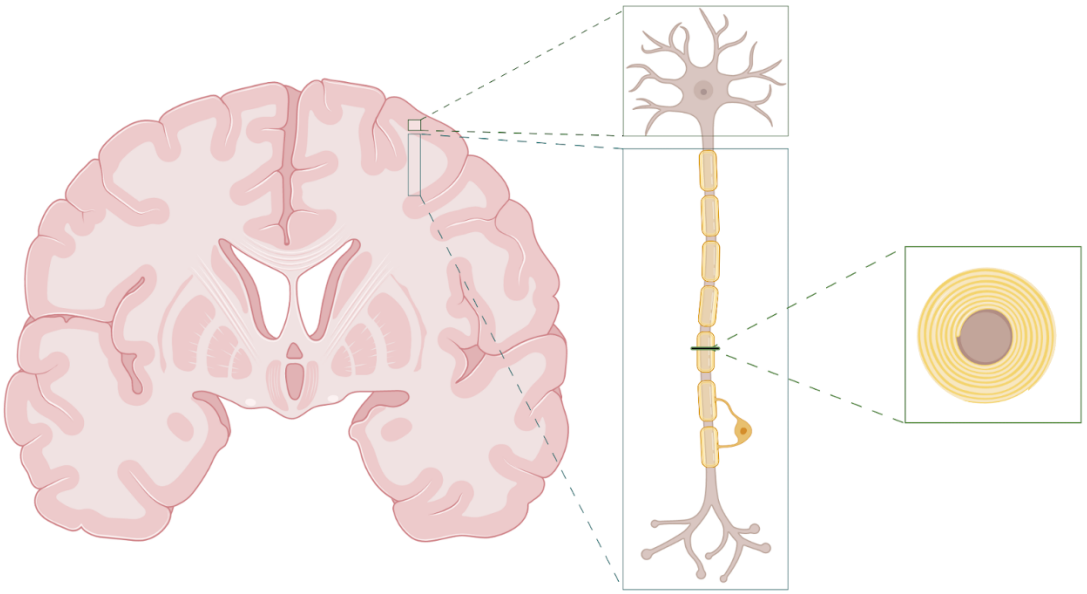


Figure 2.1: This schematic presents an overview of the structure of brain tissue, which consists of white matter and grey matter. The grey matter houses neural cell bodies and dendrites, while the white matter comprises axons that establish connections between distant neurons. The surrounding glial cells in the white matter are responsible for maintaining the neuronal environment functioning well. Oligodendrocytes are a critical type of glial cells that provide myelin sheath around the axons. Myelin, a fatty substance, enhances the speed and efficiency of electrical communication between neurons and gives white matter its characteristic appearance. The tissue is however not fully packed with these and there is plenty of space that is filled with water and capillaries. This figure is created with `BioRender.com`

2.1.2 Microstructure

The majority of white matter in the brain consists of myelinated and unmyelinated axons, which are the long, narrow fibres that stretch from grey matter and transmit signals to other brain regions. These axons are packed together into fibre bundles that create the brain's white matter tracts.

Myelin, the fatty layer that surrounds many axons and serves as an insulating layer, enables quick and effective neural signal transmission. Specialised cells called oligodendrocytes form the myelin sheath in the white matter (see Figure 2.1). Very little myelin is present in infants at birth and myelination starts during the early stages of brain development and lasts into early adulthood. Myelination can change during disease progression or due to learning and neuroplasticity.

In addition to myelinated axons, brain white matter also comprises other cells generally known as Glia, including oligodendrocytes, astrocytes, microglia, and endothelial cells. Oligodendrocytes are the main cells in the central nervous system that produce myelin. Microglia serve as the central nervous system's resident immune cells. Astrocytes give structural support and aid in regulating the brain's chemical environment. Endothelial cells are responsible for the formation of the blood-brain barrier, which controls the transport of substances to and from the brain.

Also, there is a complex network of blood vessels within the white matter, including arterioles, capillaries, and venules, to supply oxygen and nutrients to the axons and glial cells. In addition, the blood vessels facilitate the removal of waste and carbon dioxide. These vessels merge into veins that carry deoxygenated blood and waste away from the brain.

2.1.3 Anatomy

The bulk of white matter of the brain is made up of bundles of myelinated nerve fibres, known as tracts, which are responsible for transmitting signals between different regions of the brain and the rest of the body. These tracts make up a complex structural network that connects the cortex, cerebellum, brainstem, and spinal cord, allowing for the integration and processing of sensory, motor, and cognitive information throughout all of the brain. The major tracts of the white matter are categorised into smaller groups that are listed below:

Commissural fibres: nerve fibres that connect corresponding regions of the cortex in the two hemispheres. The corpus callosum is the largest bundle of commissural fibres.

Projection fibres: long nerve fibers that connect the cortex to other parts of the brain or spinal cord. They are categorized into two types: Efferent nerve fibers, which are axonal projections departing from cortex, and Afferent projections, which are those arriving at the cortex. Acoustic fibers and cerebrospinal tracts are examples of projection fibres.

Association fibres: nerve fibres that connect different parts of the cortex within the same hemisphere. Examples of accociation tracts are arcuate fasciculus, longitudinal fasciculus, frontal aslant and uncinate fasciculus.

Limbic fibres: fibre bundles that connect the limbic system within and to other brain regions. This includes the fornix and the cingulum bundle.

2.1.4 Methods for studying the white matter

The brain's white matter can be studied using a variety of techniques. Some of the most used methods include:

Tracer studies: Using tracers, substances that are injected into the brain to label certain nerve fibres with a specific colour, is one of the long-standing techniques of delineating the white matter pathways. The connections marked by the tracer are observable after the brain has been dissected. The use of tracers has greatly aided the mapping of brain connections for several species over the years. While the tracer approach is a direct method of observing connections, it has severe limitations on account of its invasive nature and artefacts caused by leaking, and the need for too many sacrifices.

Histology: Histology, the inspection of tissue under a microscope, is another method for investigating white matter. Like tracers, chemicals can be used to "stain" white matter samples, revealing the presence of certain substances like myelin, Nissl bodies, microglia, etc. Histology is especially beneficial for studying the microstructure of white matter and how it is altered as a result of diseases and pathologies. However, its application is restricted in many situations because of its invasive nature, small sample size, and quantification challenges.

Structural MRI : Certain Magnetic Resonance Imaging (MRI) modalities, specifically T_1, T_2, and FLAIR acquisitions, are employed to provide images of brain structures. These modalities have become integral components in clinical

practice, enabling the detection of various brain abnormalities including white matter lesions, tumors, and other abnormalities. The principal advantages of these MRI techniques lie in their non-invasive nature and capacity to capture comprehensive measurements of the entire brain. However, a notable limitation is their sensitivity, which often only allows detection of severe changes, frequently resulting in diagnoses in the later stages of disease progression.

Diffusion MRI (dMRI): Diffusion MRI is a form of magnetic resonance imaging (MRI) that is sensitive to the diffusion-induced movement of water molecules within tissue. By tracking the direction and speed of water molecules, this technique provides information about the local orientation and strength of white matter connections at each point in the brain. The advantage of dMRI is that it is non-invasive and can scan the entire brain at once, but the problem is it is very indirect and noisy.

Quantitative susceptibility mapping: The spatial distribution of magnetic susceptibility within biological tissues can be measured using the recently established MRI technique known as quantitative susceptibility mapping (QSM). This approach maps the magnetic field profile within the tissue that can reveal microstructural characteristics of the tissue, for instance the amount of myelin.

By integrating these approaches using computational modelling, the role of white matter in healthy brain function and neurological and psychiatric diseases can be better explored.

2.1.5 White matter pathologies

White matter structure and function can go wrong in a variety of conditions, leading to neurodegeneration, cognitive decline, and various neurological symptoms. Pathological processes that contribute to white matter degeneration include, but are not limited to, demyelination, inflammation, axonal loss, neural beading, and ischemia.

Demyelination is a common pathological process, characterized by the stripping away or damage of the myelin sheath that insulates nerve fibres. This impairs the conduction of signals in the affected nerves, causing impairment in sensation, movement, cognition, or other functions depending on which nerves are involved. Demyelinating diseases, such as multiple sclerosis, are representative examples of such pathology.

Axonal loss, another key type of white matter degeneration, refers to the death of axons. Diseases such as Alzheimer's and Parkinson's are associated with substantial axonal loss, profoundly impacting the functional capabilities of the brain's communication system.

Inflammation in the brain, often a response to injury, infection, or autoimmune disorders, can also lead to significant damage to white matter. Inflammatory cytokines and other substances can disrupt the blood-brain barrier, leading to a cascade of events that ultimately result in myelin damage and loss of white matter integrity.

Neural beading is often observed in the early stages of neuronal injury. It refers to the formation of bead-like swellings along the axons, which are typically associated with interrupted axonal transport, leading to local accumulation of cellular materials. This can induce further axonal damage or degeneration, contributing to the disruption of white matter.

Lastly, ischemia, a condition of reduced blood flow often induced by stroke or transient ischemic attacks, can lead to white matter injuries. Deprivation of essential nutrients and oxygen initiates a series of metabolic disturbances, ultimately leading to the death of neural cells and resulting in white matter lesions.

Understanding the pathological processes that contribute to white matter degeneration is vital for the development of therapeutic strategies. While each of these processes — demyelination, inflammation, axonal loss, neural beading, and ischemia— can independently contribute to the degradation of white matter, it is also important to note that they frequently co-occur and interact in many neurological

diseases. As such, the comprehensive study of these pathologies and their interplay remains a critical area of focus in the field of neurology and neuroscience.

2.2 Diffusion MRI

2.2.1 Acquisition

2.2.1.1 MRI

Magnetic Resonance Imaging (MRI) is a non-invasive medical imaging method that creates images of the body's internal structures and function by using a powerful magnetic field and radio frequency waves. The following are the primary stages of a basic magnetic resonance imaging procedure:

1. Spin alignment: By applying a very strong magnetic field, denoted as B0, the spin of a subset of hydrogen atoms within the medium get aligned with the magnetic field. This magnetic field causes the magnetised spins to precess at an angular frequency, known as the Larmor frequency, which is expressed as:

$$\omega = \gamma B_0 \tag{2.1}$$

   Where γ is the gyromagnetic ratio of hydrogen atoms.

2. Excitation: A short burst of radio frequency waves (RF pulses) is then applied to the field. This causes the spins that are aligned to the magnetic field to absorb energy and flip towards the plane perpendicular to the main field. The flip angle depends on the strength of the main magnetic field, length, magnitude and shape of the RF pulse.

3. Relaxation: When the RF pulse is turned off, the flipped spins start to return to their original state through a rotational movement and lose the energy they gained from RF pulses. The change of magnetization of the spins is governed by the Bloch equations[25]:

$$\frac{d}{dt}\begin{bmatrix} M_x \\ M_y \\ M_z \end{bmatrix} = \begin{bmatrix} -\frac{1}{T_2} & \gamma B_z & \gamma B_y \\ -\gamma B_z & -\frac{1}{T_2} & \gamma B_x \\ \gamma B_y & -\gamma B_x & -\frac{1}{T_1} \end{bmatrix} \begin{bmatrix} M_x \\ M_y \\ M_z \end{bmatrix} + \begin{bmatrix} 0 \\ 0 \\ \frac{M_0}{T_1} \end{bmatrix} \tag{2.2}$$

In this equation T_1 and T_2 are the rate of change in the magnetization parallel to the main magnetic field and perpendicular to it (transverse plane) respectively.

4. Spatial encoding: To generate an image, the MRI scanner applies gradients of magnetic fields to encode spatial information. These gradients cause the magnetic field strength to vary across the space, which in turn causes the spins to precess at different frequencies Eq. 2.1. We can later use these differences in the frequencies and phases to decode the location of spins.

5. Image reconstruction: The coils in the scanner measure RF signals emitted by spins inside the magnetic field. By altering the sampling frequency, taking into account the location-dependent frequencies, we probe the spatial arrangement of these spins. An inverse Fourier transform applied to these signals, forming what's called k-space, enables us to translate frequency-encoded signals into 3D spatial locations, known as 'voxels', with defined resolution. Various techniques for sampling k-space are available, each offering different trade-offs in speed, resolution, and signal-to-noise ratio.

2.2.1.2 Diffusion

Considering the molecular concentration of a substance as a function of both space and time $\rho(x,t)$, spatial rate of change (flux) of the substance is assumed to be proportional to the spatial gradient:

$$J(x,t) = -D\frac{\partial \rho(x,t)}{\partial x} \tag{2.3}$$

D is the diffusion coefficient that describes how fast the molecules transfer in space. According to the conservation of mass principle, the temporal rate of change of concentration is:

$$\frac{\partial \rho(x,t)}{\partial t} = -\frac{\partial J(x,t)}{\partial x} \tag{2.4}$$

Putting these equations together:

$$\frac{\partial \rho(x,t)}{\partial t} = D\frac{\partial^2 \rho(x,t)}{\partial x^2} \tag{2.5}$$

i.e. the rate of change of concentration in time is proportional to the second order rate of change in space. By solving this equation with the boundary conditions where $\rho(0,0) = 1$ and zero at all other points (a Dirichlet delta function), and $\frac{\partial \rho}{\partial t}(x, \infty) = 0$ (signifying a steady state), we can determine the concentration at a specific location 'x' at a given time 't'. This is represented by:

$$\rho(x,t) = \frac{1}{2\sqrt{\pi D t}} e^{-\frac{x^2}{4Dt}} \tag{2.6}$$

The diffusion equations can be generalised from a one-dimensional to a three-dimensional space. While the overall structure of the equation remains the same, the primary difference lies in the diffusion coefficient. In a one-dimensional context, this coefficient is a scalar quantity. However, in a three-dimensional space, it becomes a 3×3 symmetric matrix, referred to as the diffusion tensor. The generalised equation in three dimensions is as follows:

$$\rho(x,t) = \frac{1}{\sqrt{|D|(4\pi t)^3}} e^{-\frac{x^T D^{-1} x}{4\pi t}} \tag{2.7}$$

This equation can also be interpreted as the likelihood of a single particle that originated from the starting point being found at location 'x' after time 't'. In this context, this equation is referred to as the diffusion propagator, as it delineates how the substance spreads in space due to the diffusion process.

2.2.1.3 Pulsed Gradient Spin Echo

Stejskal and Tanner[26] proposed an MRI sequence designed to make the MRI signal sensitive to diffusion. The method employs two RF pulses and two magnetic field gradients. The initial pulse flips all magnetized spins into the transverse plane, and the subsequent refocusing pulse flips them back by 180 degrees(see Figure 2.2). This pair of RF pulses cancels out local variations in the magnetic field as spins that precess faster in one direction will acquire a positive phase during the first pulse and a negative phase after the refocusing pulse. This is known as a basic spin echo sequence.

The implementation of a magnetic field gradient in a specific direction after the excitation RF pulse and prior to the refocusing RF pulse causes spins to increase precession angular velocity based on their location. The application of a second identical gradient after the refocusing RF pulse results in spins to gain extra angular velocity based on their location but in exactly opposite direction. So the spins that did not move in the gradient's direction between the two gradient applications experiencing no net effect. Thus, the readout signal for these spins will be equivalent to the application of no gradients. Conversely, spins that moved in the direction of the gradient acquire additional phase due to varying speed up and down periods caused by the gradients.

In order to evaluate the impact of diffusion on the MR signal, assuming a short pulse duration such that spins remain static during the gradient application (a scenario referred to as the narrow band pulse regime), each spin's excess angular velocity caused by the first gradient application is $\omega_1 = \gamma G x_1$. In this expression, G denotes the gradient's intensity, and x_1 represents the spin's location when the first gradient is applied. Similarly, during the second gradient's application, the excess angular velocity is $\omega_2 = \gamma G x_2$. Considering that the refocusing pulse applied between the two gradients, the total extra phase due to both gradients is:

$$\phi = \gamma G \delta r \tag{2.8}$$

where δ denotes the pulse duration and $r = x_1 - x_2$.

The total magnetisation vector is the sum of all the spins' magnetization that are present within the medium. Considering that the phase of each spin depends on its movement between the two gradients, the total magnetization is:

$$S(t) = \int_0^\infty S_0(t) e^{-i\phi(r)} p(r) dr \tag{2.9}$$

Where $S_0(t)$ is the magnetization vector without any gradient encoding (that is also time dependent due to T_2 effects), and $p(r)$ is the probability that the spin has moved by distance r between application of the two gradients. In order to gain the maximum possible signal the readout time (T_e) is set to the double of

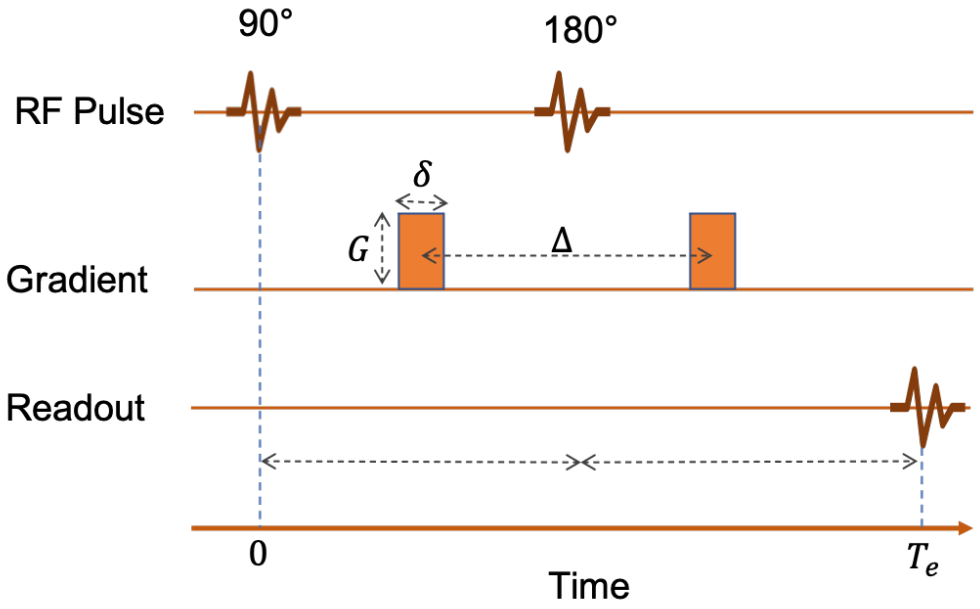


Figure 2.2: Initially proposed by Sketjal and Tanner, involves the transmission of two radiofrequency (RF) pulses. The first pulse flips the spins by 90 degrees, and the second pulse flips them by 180 degrees. The second RF pulse, also noted as refocusing pulse, eliminates differences in precession rates between spins due to the variations in the local magnetic field. To induce sensitivity to diffusion, two gradient vectors are applied on top of the main magnetic field, both with pulse width δ, one before and one after the refocusing pulse. The time interval between the application of the two gradients is Δ. If a spin does not move within the time frame between the two gradients, the gradients will cancel out, and the spin will not experience any net effect. However, if the spins move during this time period, they will undergo a phase difference from the stationary spins due to exposure to a different magnetic field. The readout signal is measured after the second gradient pulse to obtain the total signal. This pulse sequence allows for the detection of the diffusion of molecules in biological tissues, as the diffusion causes the spins to move and experience different magnetic fields, leading to a measurable phase difference in the signal.

time between the excitation and refocus RF pulse to let the stationary spins return back to their initial phase. This probability distribution is in fact the diffusion propagator for the total diffusion time. By replacing Eq. 2.8 into Eq. 2.9 and a change of variable $q = \gamma G\delta$ we have:

$$\frac{S(T_e)}{S_0(T_e)} = \int_0^\infty e^{-iqr} p(r) dr \tag{2.10}$$

The diffusion propagator is a Gaussian with variance Dt where t is the total diffusion time. Correcting for the spin movements during the application of gradients, the actual diffusion time is approximated as $t = \Delta - \frac{\delta}{3}$[27] where Δ is the time

interval between the application of two gradients (see Figure 2.2). Therefore, the signal attenuation due to diffusion is calculated from Eq. 2.10 as:

$$log(\frac{S}{S_0}) = -Dq^2(\Delta - \frac{\delta}{3}) \tag{2.11}$$

The signal attention is conventionally written as e^{-bD} where $b = (\gamma\delta G)^2(\Delta - \frac{\delta}{3})$ is referred to as b-value (accordingly the gradient direction is called b-vector). b-value reflects the sensitivity to the diffusion that is controlled by the strength and duration of the gradients used to generate diffusion-weighted images.

2.2.1.4 Diffusion weighted imaging

The microstructural characteristics of the brain tissue can be measured through diffusion signals by altering the acquisition parameters. In a pulsed gradient spin echo (PGSE) experiment, this entails altering the gradient's orientation and amplitude, diffusion interval, pulse width, and echo time. Changing any of these factors can reveal new information about tissue properties.

If the gradient is applied in only one direction at each time, it results in linear diffusion encoding that is only sensitive to diffusion along that orientation. By holding all other acquisition settings fixed and changing the orientation in several different directions, we acquire information about the geometrical structure, a scheme known as high angular resolution diffusion imaging (HARDI).

By adjusting the gradient's amplitude, duration, and time interval between two pulses, it is possible to change the b-value, which marks the overall signal sensitivity to diffusion scale. Because of the increased sensitivity to minor movements, the signal attenuation rate increases as the b-value increases for the same size of movement. A collection of measurements taken at different angles but with the same b value is conventionally referred to as a shell. Because low b-value shells have a high signal and high b-value shells are more sensitive to the geometry of the tissue, a full acquisition typically contains multiple b-values, also known as multi-shell data. An example diffusion dataset as acquired in the Human Connectome Project (HCP) [28, 29] is shown in Figure 2.3. This dataset constitutes three shells of diffusion data.

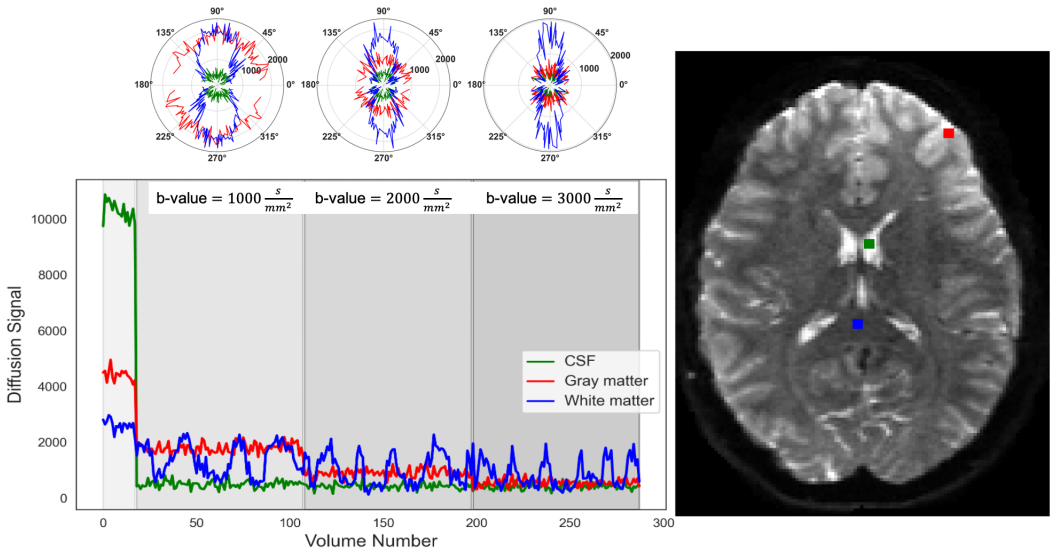


Figure 2.3: This figure shows diffusion data obtained from a subject within the Human Connectome Project's young adult dataset. The image on the right is showing a slice of a B0 image (no diffusion encoding). In this map, free water exhibits high signal intensity, appearing brighter than brain tissue, a characteristic similarly seen in T_2-weighted images. The plot on the left provides a representation of diffusion signal intensities across three distinct tissue types: CSF, white matter, and gray matter per volume of data acquisition (voxel coordinates are marked with their corresponding colour on the B0 image). Each volume in the dataset corresponds to a specific b-value and gradient direction, as indicated on the plot. Notably, the CSF signal is respectively high without diffusion encoding and undergoes significant decay even at low b-values, indicative of high diffusivity. In contrast, gray matter signal is considerably lower without diffusion encoding, yet it decays more gradually as b-values increase. The signal in white matter decays at a similar slow rate, but shows high dependence on the gradient orientation. The polar plots represent an alternative visualization of the same data, portraying the signal in relation to the gradient angle for each b-value with respect to the left-right axis on the map. It can be observed that white matter signal heavily depends on the gradient orientation, and it becomes more sensitive as b-value increases. In contrast, the signal in CSF and gray matter demonstrates minimal directional dependence.

Restricted water results in a distinct time-dependent behaviour of the signal, which may be distinguished by adjusting the time interval between the two diffusion gradients. If we consider that different types of tissues have different T_2 effects, it is useful to also alter the echo time in order to quantify T_2 effects independent of diffusion effects.

2.2.2 Data analysis

2.2.2.1 Preprocessing

Like any other data, Diffusion MRI data is often contaminated with artefacts that can reduce the quality of the data and affect the reliability of the final results. To address this, several preprocessing steps are typically performed to improve the quality of the data and correct for any artefacts that may have occurred during the acquisition process. While the exact steps may vary depending on the specific study design and software used for analysis, quality control, motion and eddy current correction are commonly performed to address the artefacts that can occur during the acquisition of dMRI data.

The first step in the preprocessing of dMRI data is quality control, which is necessary to ensure that the data is free of outliers and has an acceptable signal-to-noise ratio.

Secondly, the diffusion MRI can suffer from artefacts induced by variations in magnetic susceptibility of tissues, which can result in non-uniform magnetic field, and hence geometric distortions in the images. FSL's TOPUP[30, 31] is a popular tool for fixing this problem. TOPUP calculates the susceptibility induced off-resonance field by using repeated acquisitions with opposite k-space traverse directions. Since the distortion effect will be reversed when the k-space is traversed in the opposite direction, this knowledge allows for the mapping of field variations directly to the direction of distortion. TOPUP then finds the additional field that, when applied to these volumes, yields the highest similarity between the unwarped images. After estimating the field map, all of the diffusion volumes can be corrected for these artefacts by applying the estimated extra field to them.

Thirdly, the rapidly varying magnetic field during diffusion encoded acquisition induces an electrical current in the coils, referred to as eddy current. This electrical current generates an extra magnetic field that causes distortions in the measured signal. FSL Eddy[32] is a popular tool that utilises signal modelling with a Gaussian process to estimate the artefacts by exploiting the fact that diffusion signal does not change with a gradient in the opposite direction but eddy currents flip. FSL

Eddy estimates the distortions due to the extra induced field as well as to subject movements in each volume. It then unwarps the diffusion images using the estimated distortion map to correct for eddy currents and subject movements at the same time.

Even though these preprocessing steps work reasonably well in removing known artifcats, still care must be taken when reporting the results. Particualry, since diffusion MRI is sensitive to random movement of water in microscopic scales, it can be highly sensitive to subject movements [33, 34].

2.2.2.2 Diffusion Tensor metrics

The diffusion tensor model (DTI) is a signal representation for dMRI data that describes the diffusion in each location of the brain as an elongated Gaussian motion.[35, 36] The DTI is based on the concept of the diffusion tensor, and models the signal attenuation as an exponential decay:

$$S = S_0 e^{-bg^T Dg} \tag{2.12}$$

Here, b and g represent the b-value and gradient orientation (unit vector), respectively. S0 is the signal without diffusion encoding, and D is the diffusion tensor. The diffusion tensor has 6 free parameters due to its symmetric nature. These free parameters (diffusion tensor and S0) can be estimated from at least seven samples of diffusion data using non-linear optimisation techniques.

Once the diffusion tensor is estimated, it can be used to calculate a number of metrics that characterise the tissue's diffusion properties, including mean diffusivity (MD):

$$MD = \frac{\lambda_1 + \lambda_2 + \lambda_3}{3} \tag{2.13}$$

and fractional anisotropy (FA):

$$FA = \sqrt{\frac{3}{2}\frac{(\lambda_1 - MD)^2 + (\lambda_2 - MD)^2 + (\lambda_3 - MD)^2}{\lambda_1^2 + \lambda_2^2 + \lambda_3^2}} \tag{2.14}$$

Where $\lambda_1, \lambda_2, \lambda_3$ are eigen values of the diffusion tensor.

FA is the most commonly used [37, 38] DTI metric that reflects the degree of anisotropy in the diffusion tensor ranging from 0 (diffusion in all directions) to 1

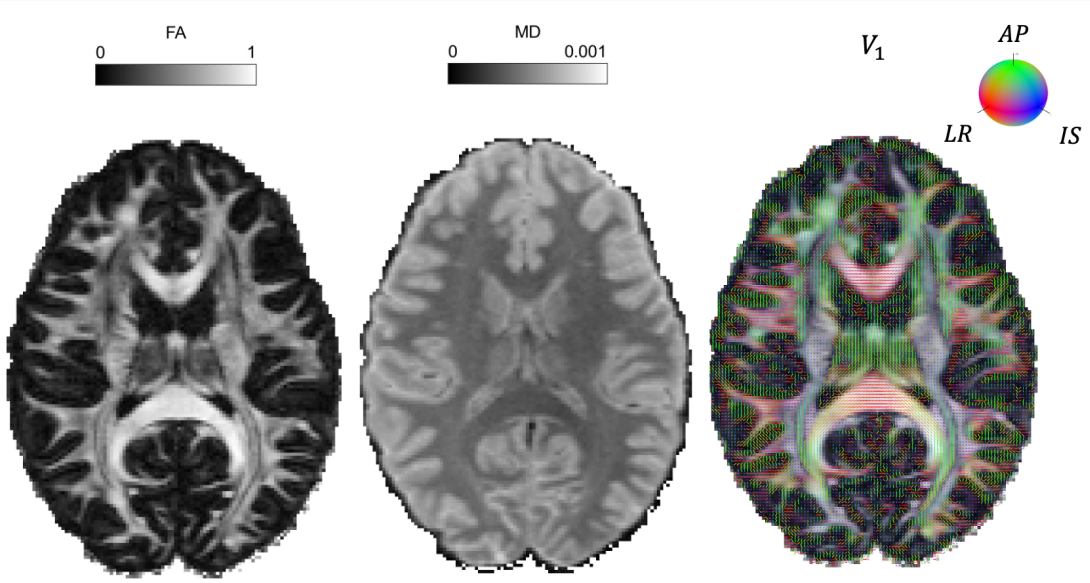


Figure 2.4: This figure illustrates three key metrics derived from Diffusion Tensor model (DTI). The left map displays Fractional Anisotropy (FA); regions with high anisotropy, primarily major fibre tracts, appearing brighter. The middle image presents a slice of Mean Diffusivity (MD). In this map, regions with higher overall diffusivity, including CSF, are represented by brighter intensity. The right map highlights the primary axis of the diffusion tensor, which corresponds to the eigenvector associated with the highest eigenvalue (V1), that is overlaid on the FA image. The orientation of the tensor is further clarified through color-coding for better visualization, in which red corresponds to Left-Right, blue corresponds to Inferior-Superior, and green corresponds to Anterior-Posterior orientations. This tensor orientation map can inform us about the direction of major fibre bundles at different brain locations. For instance, the corpus callosum, situated centrally, is distinctively depicted in red, signifying its left-to-right orientation.

(isotropic diffusion in only one direction). MD reflects the average diffusion rate across all directions in the voxel, which can be a proxy to tissue density. An example map of FA, MD and principal axis of diffusion tensor fitted to a sample subject from the HCP young adults dataset is shown in Figure 2.4.

Owing to its simplicity and sensitivity, the diffusion tensor model is the most popular model of diffusion, and the associated metrics have been widely employed in both research and clinical applications[39–43]. The model, however, is far from a genuine biophysical model, and the measurements are not explicit and directly interpretable[44].

2.2.2.3 Multicompartment models

The diffusion signal from white matter can be modelled by breaking down the white matter tissue into a series of basic compartments. The goal is to make it relevant from a biophysical perspective, so that changes in each parameter reveal interpretable information about the tissue.

In this framework, the diffusion MRI signal is modelled as a sum of contributions coming from various voxel-level compartments. These compartments may represent the intra axonal space, extra axonal space, and free water that make up the whole voxel water content. Each compartment's diffusion signal is represented by a particular type of diffusion tensor, which expresses the speed and direction of water diffusion there.

Table 2.1 provides a description of the most used compartments. The total dMRI signal is a weighted sum of the signals from each compartment, with the weights representing the contribution of each compartment's signal in the voxel. These proportions reflect a combination of volume and relaxation effects. To accommodate for more complex geometric structure, each compartment can also have a dispersion model. Dispersion is normally represented by spherical harmonics or a Watson or Bingham distribution. Watson distribution is the intersection of isotropic Gaussian distribution with the surface of a unit sphere. The Bingham distribution is the intersection of a Gaussian with a covariance structure with the surface of unit sphere.

Compartment	**Attenuation**	**Used for**
Dot	$A = 1$	Glia, trapped water
Ball	$A = e^{-bd}$	Free water, CSF
Stick	$A = e^{-bd(\mu^T g)^2}$	Intra axonal water
Tensor	$A = e^{-bg^T Dg}$	Intra- and Extra-axonal water
Sphere	GPD approximation[45]	Restricted water in glia/soma
Cylinder	GPD approximation[46]	Restricted water in axons

Table 2.1: Basic compartments

We can estimate the compartment parameters and relative proportions of each compartment in the voxel by fitting the multi-compartment model to the dMRI signal using non-linear optimization. Examples of these models include AxCaliber[11], NODDI[10], SANDI[47], VERDICT[48], etc. These parameter values can then be utilised to determine the extent to which tissue qualities have changed as a result of a pathology or disease.

2.2.2.4 Degeneracy problem

Multi-compartment models offer biophysically meaningful grounds for modelling white matter tissue using diffusion MRI data. Many microstructural characteristics of the tissue, such as size distribution, morphology, density, and orientation, have an impact on the dMRI signal. Microstructural modelling aims to estimate all these characteristics by fitting a large model with many parameters to the observed diffusion data. Therefore, it is possible that multiple combinations of the model parameters can generate the same diffusion signal. For instance, a model with sticks with high diffusivity and high dispersion could create a diffusion signal very similar to a ball compartment. Figure 2.5 presents another example of this phenomenon. It demonstrates that a sphere characterized by a combination of low diffusivity and small radius can account for the diffusion signal from gray matter as effectively as a sphere defined by high diffusivity and a large radius.

This leads to the degeneracy problem that is the challenge of reliably estimating microstructural parameters from the data because the models are not invertible. The degeneracy problem is a major challenge in microstructural modelling of diffusion data and confounds finding interpretations. This issue must be taken into consideration when analysing the outcomes of microstructural modelling research, and care must be taken when drawing inferences about the microstructure from the estimated parameters.

A number of approaches have been put out to deal with the degeneracy problem[9, 14], including taking into account biologically informed prior knowledge of the microstructure[10], utilising regularisation models that can combine many

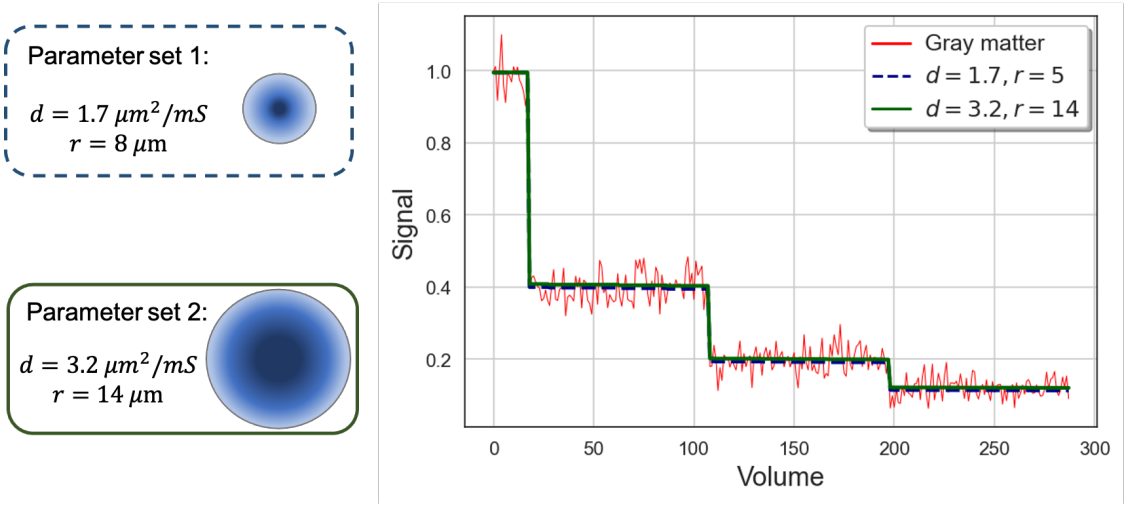


Figure 2.5: This figure provides an example of model degeneracy in the context of diffusion data. A spherical model was employed to fit the diffusion signal from gray matter, previously presented in Figure 2.3, depicted in red in the current plot. Two distinct sets of parameters exhibit equivalent efficacy in fitting the data – one entailing a small sphere with low diffusivity and the other involving a larger sphere with high diffusivity. The first model, characterized by a diffusivity of $1.7\mu m^2/mS$ and a radius of $8\mu m$, differs significantly with the alternative model that has a higher diffusivity value of $3.2\mu m^2/mS$, with a radius of $14\mu m$. Given the comparable fit of both parameter sets to the data, distinguishing between them is impossible. In this case, a diffusion sequence exhibiting variation in diffusion time (Δ) may provide valuable insights. This is due to the fact that water restricted within a sphere demonstrates a distinct pattern within a suitable range of diffusion times.

microstructural characteristics into a single parameter[49], and using complementary imaging techniques to provide extra information on the microstructure. Nonetheless, it is still an active area of research, and more progress is needed to increase the accuracy and robustness of microstructural modelling. This problem is tackled in the second part of the thesis via Bayesian modelling.

2.2.2.5 Group analysis of diffusion MRI data

Diffusion MRI can disclose valuable information about tissue structure for individuals, but it is challenging to do group analysis using diffusion metrics [50]. This is mostly due to the fact that each subject's fibre bundles are slightly displaced and oriented differently, which complicates inter-subject alignment. In other words, it is highly possible that one compares diffusion metrics that are from different parts of the tissue across subjects.

One method for group analysis of diffusion MRI data is directly performing a GLM on orientation-invariant measures that summarise the entire voxel, such as FA, after registering these images into a standard common space [51]. While this method can shed light on alterations in the white matter at the local level, it has significant limitations [52]. The voxelwise measurements are often highly unspecific because they summarise all the data into a single metric. Second, this method heavily depends on having an accurate volumetric registration. Finally, a simple shift in one tract can skew results in all neighbouring locations.

Another method is TBSS [53], which is based on projecting diffusion measurements onto a skeleton of white matter tracts and analysing the projected data statistically. TBSS creates a "skeleton" of white matter by extracting the centres of the major tracts that are shared by all subjects from FA images. A group-wise mean skeleton is then produced using a non-linear registration method that aligns the skeletons into a common space. The DTI metrics are then projected onto the skeleton, to be captured along the coregistered skeleton of all individuals. Following that, the projected data is examined using conventional statistical techniques to compare between-group variations in diffusion metrics along the skeleton. TBSS is a powerful approach and has been widely utilised in the literature [54–58] to investigate changes in white matter microstructure. However, it has a few drawbacks, including the fact that it can only be used in skeletal regions and the fact that it does not characterise crossing fibres.

The white matter tracts obtained using tractography may also be utilised for tract-based microstructural analysis, a technique known as tractometry. This involves estimating diffusion properties such as MD, FA, etc along the length of the tracts. Following that, the diffusion measures can be compared between subjects or linked to behavioural or clinical characteristics using statistical analysis.

Fixel-Based Analysis (FBA), a technique that has gained popularity more recently, analyses microstructure on the level of individual fibres, or "fixels," rather than the entirety of a voxel. A fixel is a small patch of white matter that is contained within a voxel with a certain orientation that reflects a single fibre population

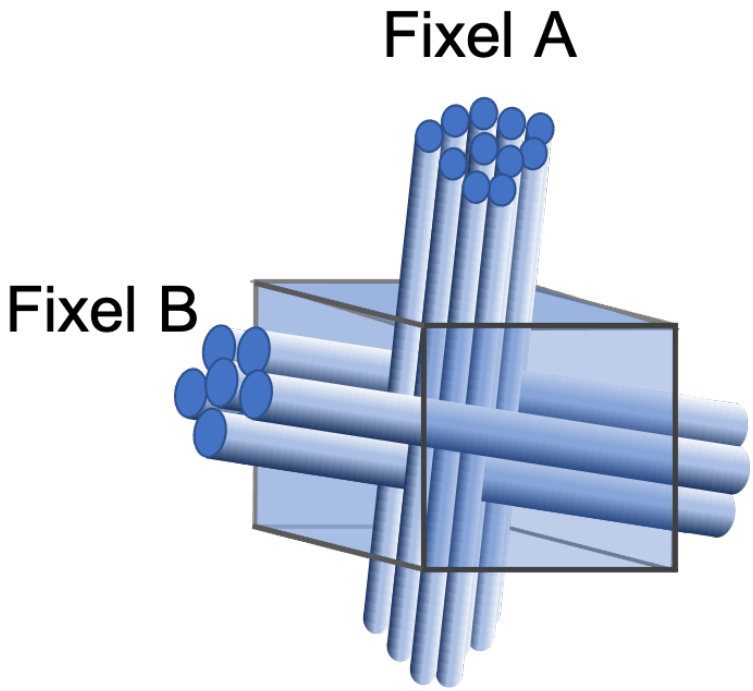


Figure 2.6: This figure provides a schematic illustration of a voxel from a white matter region that contains crossing fibers. In the field of diffusion MRI, the term 'fixel'('FIber'+'voXEL') is utilized to describe a discrete element of a fibre bundle that exists within a single voxel's spatial extent. The fixels occupying each voxel may poses heterogeneous microstructural properties, which might be missed when the data is analysed at a voxel level. However, a critical challenge intrinsic to the analysis of diffusion MRI data with a fixel-based approach pertains to the potential positional and orientational variation of fibres across different subjects. Even slight shifts in fibre positions or alterations in their orientations can complicate the process of identifying corresponding fixels across multiple individuals.

(FIbre+ voXEL), as shown in Figure 2.6. The most used implementation of FBA, as implementerd in the MRTRIX software package [59, 60] uses deconvolution techniques to estimate the fibre orientation distribution function (fODF) within each voxel. Following that, the subject level ODFs are registered to a common space to estimate group level ODFs, where group fixels are extracted. Then, for the orientations that are matched with a group level fixel, microstructural metrics such as fibre density (FD), fibre cross section (FC), and fibre density and cross section combined (FDC) are estimated for each subject.

The third chapter of this thesis is about development and implementation of a new tool for fixel based analysis.

2.2.2.6 Tractography

Tractography is the technique of reconstructing white matter pathways in the brain from diffusion MRI data. Diffusion MRI provides information about the direction and scale of water diffusion in all locations of the brain. The myelin layer that shields the water within the fibres makes it more likely that the water will diffuse along the direction of the fibres than across them; thus, the diffusion signal can be used to infer the fibre orientations at each point in the brain.

To be able to run tractography one first needs to estimate local orientation of fibres in each voxel from diffusion data, which can be done in various ways (see Figure 2.7). One such method is spherical deconvolution[61], where each fibre is assumed to have the same fibre response function (FRF) that is estimated using few exemplar voxels. Diffusion MRI data is then deconvolved with the estimated FRF to calculate the fibre orientation distribution function (fODF) for each voxel.

Another approach is using the ball and sticks model, which explicitly models crossing fibres as a combination of the ball compartment representing free diffusion and a few sticks compartments representing the fibre populations. To estimate the orientation of the fibres in each voxel the model is fitted to the data using nonlinear optimisation or sampling techniques such as MCMC.

The next step of tractography involves connecting the local orientation estimates to rebuild the pathways of white matter fibres throughout the brain. There are several different algorithms that can do so. These algorithms use a given starting mask to track the paths of the fibres, and can be configured in several ways to follow different types of fibres, such as those connecting specific brain regions or avoiding them. Tractography algorithms can be deterministic or probabilistic. Deterministic tractography follows the most likely path of a fibre through the brain. This type of tractography is suitable for identifying well-defined major fibre tracts but may have difficulty resolving small tracts or crossing fibres. Probabilistic tractography estimates the probability of a fibre path based on the distribution of direction of diffusion at each point. This type of tractography can better resolve crossing

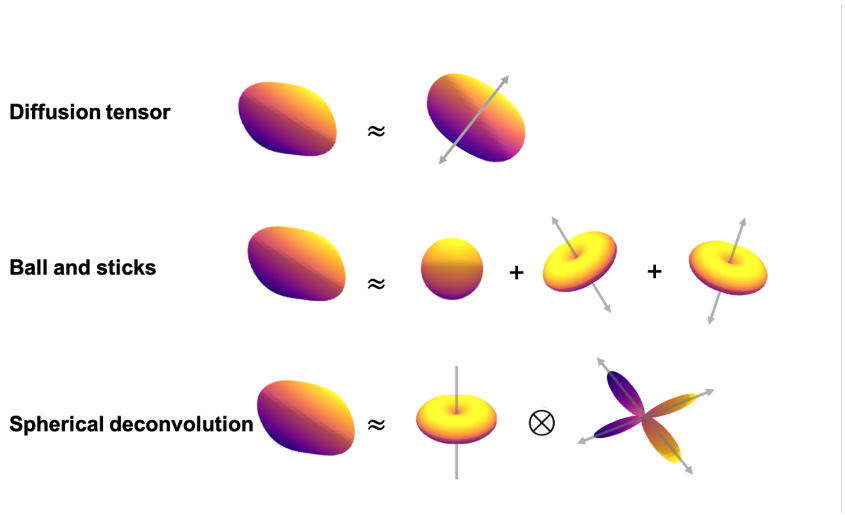


Figure 2.7: Three distinct ways of estimating the orientations of local fibres using diffusion MRI data. The principal axis of the diffusion tensor model can be used to extract the major orientation of diffusion in the tissue. However, this approximation might not be reliable in the crossing fibre regions. The ball and sticks model calculates the orientations by non-linearly fitting exponential decays for fully anisotropic compartments (representing water inside axons) and isotropic compartments (representing free water). The spherical deconvolution method uses a small number of sample voxels to estimate the diffusion signal for a single fibre population. The estimated response function is then used as a kernel to deconvolve the diffusion MRI signal in order to estimate the fibre orientation distribution in all voxels. The orientation distributions in this figure are created using fsl-dive.

fibres and provide more detailed information about the connectivity of the brain, but it may suffer from too many false positives.

After constructing the tracts using streamlines, they can be used to quantify the structural connectivity between brain regions.

2.3 Bayesian Modelling

Brain imaging data is a powerful tool for studying the brain's structure and function non-invasively. However, this data is often too indirect and can be challenging to interpret. To overcome this challenge, we try to translate this data into understandable components that can help us predict and explain the brain's characteristics and behaviour that aren't immediately visible.

To achieve this, we construct models that rely on mechanistic laws and simple components that can explain or generate the observed data. These models help us

establish a connection between the observed data and the parameters of interest that reveal the underlying structure and dynamics of the system.

However, empirical data is frequently confounded by noise, outliers and other sources of uncertainty, necessitating the consideration of model parameters as random variables rather than deterministic fixed values. Additionally, the indirect nature of the observations presents substantial challenges, making it difficult, if not impossible, to accurately establish the relationships and dependencies among these parameters solely based on the available measurements. In such cases, we use probabilistic models that use random variables to describe the parameters of the model, instead of using fixed values. Moreover, the unpredicted or unmodeled contributions to the observations are considered as noise. This probabilistic approach enables us to capture the inherent randomness and complexity of the system, which can lead to a more robust and accurate understanding of the data.

2.3.1 Bayes rule

Bayes' rule is centred on the equation that describes how to update the probabilistic model of a phenomenon based on new evidence or information. It is named after the British statistician Thomas Bayes, who first proposed the idea. The Bayes rule states that:

$$P(model|data) = \frac{P(data|model)P(model)}{P(data)} \tag{2.15}$$

To give an example to understand each part, consider we want to know how likely it is that someone has COVID (model) if they got a positive test result (data).

$P(model|data)$**:** is referred to as the posterior probability, that is the probability of this person having COVID given a positive test result. This is what we want to calculate in the end.

$P(data|model)$**:** is referred to as likelihood, that is the probability of getting a positive test result when someone has COVID. This is the sensitivity of the test that can be learnt from previous experiments.

$P(model)$**:** is referred to as prior, that is our belief that this person has COVID, which may be based on factors such as age, social interactions, and symptoms; but not directly related to the test results.

$P(data)$**:** is referred to as evidence, that is the probability of getting a positive test result, which may depend on the accuracy and reliability of the test, but not coming from the current patient condition.

There are two concepts that are used with Bayes rule. Conditioning, refers to the process of updating the probability distribution of a variable given new evidence or data. Marginalisation refers to the process of integrating over a variable in a joint probability distribution to obtain the distribution of other variables. It is used to perform inference on a subset of the model's parameters, which can be useful for making predictions without the need to know all other ones. We have used these in the second chapter of this thesis in order to infer the likely microstructural changes in a diseased tissue without needing to know all other parameters of the model.

2.3.2 Computation problem

The Bayes equation appears to be straightforward and uncomplicated; however, in practice, it can be challenging to compute, primarily due to difficulties in calculating the evidence. This is because, to determine the evidence, we need to consider all possible models that can explain the data:

$$P(data) = \int_{all-possible-models} P(data|model)P(model)d.model \tag{2.16}$$

Performing this integration over all potential models can be extremely computationally intensive. There are ways to address this issue, that are briefly mentioned below.

One approach involves using specific forms of prior distribution that yields a closed-form solution for the posterior distribution. When the posterior distribution belongs to the same family of distributions as the prior distribution, the prior and the likelihood are termed conjugates. This relationship simplifies the calculation of the posterior distribution because the algebraic form of the posterior distribution is

known. As a result, we only need to update the parameters of the prior instead of solving the entire integral every time we get a new data. However, in many cases, particularly when dealing with complex likelihood functions, a conjugate prior for the likelihood function may not exist, leading to what is known as the conjugacy problem.

An alternative method to address this issue is to use sampling techniques to generate random samples from the posterior distribution rather than computing the posterior distribution itself. One prevalent technique is Markov Chain Monte Carlo (MCMC), which generates a Markov chain of samples such that the chain's normalised distribution converges to the target posterior distribution. The Metropolis-Hastings algorithm, a widely-used MCMC algorithm, iteratively proposes new samples based on the previous sample using a proposal distribution. Then, it accepts or rejects the new sample based on a probability determined by the ratio of the unnormalized posterior distribution at the proposed sample and the current sample. By utilising the ratio of the posterior at two locations, the evidence is cancelled out, eliminating the need for its calculation. Sampling techniques are very strong and popular tools in a variety of problems, however, when the number of parameters is large or the target distribution has irregular shapes their performance drastically drops. Also, these methods are characterised by a considerable computational demand and a slow convergence rate. Furthermore, they are markedly sensitive to the selection of hyperparameters, including initial conditions, proposal distribution, and other associated parameters.

Another approach involves variational methods, which reformulate the integration in computing the evidence into an optimization problem. This method assumes a specific closed-form solution for the posterior distribution and seeks to find its parameters to make it as close as possible to the true posterior distribution. The similarity between the approximate distribution and the true distribution is measured using the KL-divergence. It is straightforward to show that minimising the KL divergence is equivalent to maximising another function that does not require computing the evidence. Thus, the problem of solving the integral transforms into optimising a cost function over the parameters of the approximating distribution

without the need for integration. Variational methods offer advantages over the previous approaches, such as not requiring specific form of priors and being faster and less computationally demanding than sampling techniques. However, their reliance on certain types of distributions with closed-form solutions can be limiting.

Finally, normalising flows[62], employ regression models from machine learning to generate parameterized, yet free form probability distributions, instead of relying on a limited set of probability distributions. This allows the approximating probability distribution in the variational approach to have a more flexible shape capable of covering a wide range of posterior distributions. These approaches are very strong and have been employed in cutting-edge machine learning applications. Nevertheless, due to the highly intricate nature of the cost function, training these models is challenging, and conventional optimization techniques may not work well. Moreover, in order to achieve a good performance, these methods necessitate proper samples of unbiased data to cover all the possibilities, which consequently results in a considerable computational demand during the training process.

2.3.3 Prior selection

In Bayesian analysis, prior distributions serve to provide our prior knowledge or assumptions regarding the model before observing the data. In certain situations, we may have minimal or no information about the model. In these cases, it is desirable to employ uninformative priors, which are prior distributions that impose minimum assumptions on the model parameters and allow the data to derive the posterior distribution. The selection of an uninformative prior relies on the specific problem and the nature of the parameters. However, there are a few general strategies for making uninformative priors.

A straightforward but naive workaround is to employ improper priors, which are functions that do not integrate to a finite value, hence can't be a proper probability distribution. A common example of such a function is a constant function everywhere, representing an infinitely wide uniform distribution. Although this prior appears to impose no assumptions on the model parameters, it may

covertly favour a specific range of parameters, especially if the parameters can vary across scales. Using improper priors is generally discouraged, as they violate the fundamental assumptions of Bayesian analysis and render the resulting posterior distribution challenging to interpret as a proper probability distribution.

Another strategy is utilising Jeffrey's prior, which is based on the principle that the prior distribution should remain invariant under reparameterization of the model. While this approach is effective for simple distributions, it exhibits strong limitations when dealing with multiple parameters, particularly in the context of the conjugacy problem.

The alternative approach is employing reference priors, which seeks to identify a function that maximises a certain measure of distance or divergence between the posterior and prior distributions. By doing so, this method tries to minimise the impact of the prior on the inference, thereby allowing the data to play a more prominent role in shaping the posterior distribution. However, estimating such a function can be highly challenging, particularly when the likelihood function is complex. Moreover, concentrating exclusively on minimising the influence of the prior makes it susceptible to reparameterization and complicates the integration of available prior knowledge.

An alternative approach involves the utilization of hierarchical priors, which involve incorporating prior distributions at multiple levels within a hierarchical model. This methodology enables the estimation of parameters at both the population level and the individual level. By employing hierarchical priors, the inherent uncertainty and variability in the parameters across different levels are captured, resulting in a more comprehensive representation of the underlying data.

2.3.4 Why Bayesian modelling?

Bayesian modelling refers to any statistical method constructed on Bayes's rule, which permits the incorporation of uncertainty and prior beliefs into the modelling process. This framework is different from black-box machine learning algorithms, which typically seek to categorise or predict particular features from the data.

Bayesian modelling has several advantages over alternative approaches in machine learning. Firstly, Bayesian methods are clear about their assumptions, while other methods make assumptions covertly. Secondly, Bayesian modelling provides a direct way to integrate uncertainty into the modelling process and the predictions. That is to say, rather than point estimates, Bayesian inference provides samples or probability distributions to explain/predict the outcomes. Thirdly, Bayesian models are conceptually straightforward and logically consistent, making them a modular and adaptable method. Because of the modular structure, Bayesian models can utilise the relationships and interactions between variables in order to make more efficient inference, such as in graphical models. In the second chapter of this thesis, we have employed the techniques of Bayesian analysis, specifically conditioning and marginalisation, to isolate parameters that are not relevant to our objective of estimating changes in microstructure. The utilisation of Bayesian modelling has enabled us to transform the challenge of fitting a complex model to limited data into a tractable and more clinically relevant problem. Hierarchical models are a class of Bayesian models that allow for the incorporation of multiple levels of variation and structure in the data. In a hierarchical model, the data are assumed to be generated from a series of nested distributions, with each level representing a different source of variation or structure. For instance, when modelling individual data, one can employ population data within a hierarchical framework. This approach enables us to utilise data from other subjects to make more accurate predictions for individuals and also the development of models capable of explaining sub populations. In the third chapter of this thesis, we applied this method to estimate fibre orientations in individual brains such that they corresponded to population-level tracts.

"Change is the only constant in life."

— Heraclitus

3

Bayesian Estimation of Change

Contents

3.1 Abstract 35
3.2 Introduction 36
3.3 Theory 39
3.3.1 Inference on change in parameters 39
3.3.2 Training models of change 42
3.3.3 Biophysical model of diffusion 45
3.3.4 Summary measurements 47
3.4 Methods 49
3.4.1 Simulations 49
3.4.2 White matter hyperintensities 52
3.5 Results 53
3.5.1 Summary measurements 53
3.5.2 Validation 54
3.5.2.1 Comparison with model inversion 54
3.5.2.2 Sensitivity to change in each parameter 56
3.5.2.3 Estimating the amount of change 56
3.5.3 White matter hyperintensities 59
3.5.3.1 Model Inversion 59
3.5.3.2 BENCH 59
3.6 Discussion 64
3.6.1 Comparison with model inversion 64
3.6.2 Summary measures 67
3.6.3 Future developments 68
3.7 Software 69
3.8 Appendix A. Toy example: Inferring changes in 2D . . 69
3.9 Appendix B. Estimating Quality of Fit 70

3.1 Abstract

Biophysical models that attempt to infer real-world quantities from data usually have many free parameters. This over-parameterisation can result in degeneracies in model inversion and render parameter estimation ill-posed. However, in many applications, we are not interested in quantifying the parameters *per se*, but rather in identifying changes in parameters between experimental conditions (e.g. patients vs controls). Here we present a Bayesian framework to make inference on changes in the parameters of biophysical models even when model inversion is degenerate, which we refer to as Bayesian EstimatioN of CHange (BENCH).

We infer the parameter changes in two steps; First, we train models that can estimate the *pattern of change* in the measurements given any hypothetical direction of change in the parameters using simulations. Next, for any pair of real data sets, we use these pre-trained models to estimate the probability that an observed difference in the data can be explained by each model of change.

BENCH is applicable to any type of data and models and particularly useful for biophysical models with parameter degeneracies, where we can assume the change is sparse. In this study, we apply the approach in the context of microstructural modelling of diffusion MRI data, where the models are usually over-parameterised and not invertible without injecting strong assumptions. Using simulations, we show that in the context of the standard model of white matter our approach is able to identify changes in microstructural parameters from conventional multi-shell diffusion MRI data. We also apply our approach to a subset of subjects from the UK-Biobank Imaging to identify the dominant standard model parameter change in areas of white matter hyperintensities under the assumption that the standard model holds in white matter hyperintensities.

3.2 Introduction

Modelling diffusion MRI (dMRI) data comes in two flavours. Phenomenological models, such as diffusion tensor imaging (DTI) [36] and diffusion kurtosis imaging (DKI) [63]) attempt to describe the diffusion signal in a structured mathematical form,while (bio)physical models such as the standard model [64], NODDI [10], Ball and Rackets [15] and AxCaliber [11]) attempt to infer properties of the tissue microstructure given the data. This active field of research relies on the inversion of biophysical forward models, but it is also notoriously difficult to overcome model degeneracies [14]. To resolve these degeneracies, the conventional approach is to constrain a subset of the parameters and only make inferences on the remaining parameters [10]. However, when the assumptions are not accurate (e.g., in diseased tissue), they will bias the estimated model parameters and cause errors in interpretation. As a result, not only is there a limit to the number of microstructural parameters that can be estimated, but the reliability of the estimated parameters can also be questionable [14, 65, 66].

It is worth mentioning that there are efforts on acquiring complementary information using for example multiple diffusion encoding [67–69], as well as introducing more biophysically informed priors to limit the search space, to provide enough constraints to uniquely estimate the parameters of the standard model. Here we adopt the standard model of white matter fitted to conventional multi-shell diffusion MRI data as a well-studied degenerate model merely as a toy example to illustrate the concept.

However, in many real-world applications, the model parameters may not be of direct interest. Rather, we are often interested in the "change" in the parameters under different experimental conditions. For example, to study mechanisms underlying a disease one might compare the parameter estimates of biophysical models between patient and control groups. Nevertheless, the parameter estimation is only tractable when the model of interest is invertible given the data. This limits one to simple biophysical models or requires injection of prior assumptions.

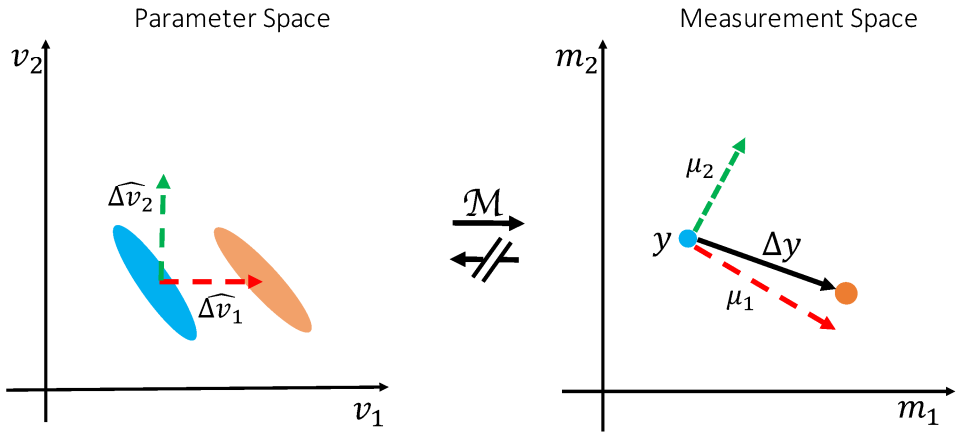


Figure 3.1: Illustration of the inversion-free inference on change (BENCH). Consider a toy model with two parameters and two measurements $\mathcal{M}(\nu_1, \nu_2) = [m_1, m_2]$. The dimensions m_1 and m_2 are two independent measurements that summarise the data. For example, in the context of diffusion signals, these could correspond to mean diffusivity and fractional anisotropy. Also, y symbolizes a point within the space defined by m_1 and m_2, which represents a specific set of measurements. Each oval in the parameter space (left) corresponds to a single point in the measurement space (right) with the same colour; signifying that there is a one to many mapping from measurements to parameters (i.e., the model is degenerate). Despite the degeneracies we are able to estimate which of the parameters best explains the change in the measurements. We do so by comparing the observed change (Δy) with the expected change in the measurements (μ_1, μ_2) as a result of each hypothesised *pattern of change* ($\Delta\hat{\nu}_1$, $\Delta\hat{\nu}_2$), please refer to Table **??** for clarification about the notations.

In this study, we show that it is possible to make precise inferences on the change in model parameters even in complex degenerate models. We argue that, using a sparsity assumption on the *pattern of change*, we can limit the hypothesis space, and so circumvent the degeneracy in the parameter estimation (see Figure 3.1, also refer to 3.8 for more details about directly inferring changes). Our approach proceeds in two steps: First, we use simulated data generated from a forward model to train models that calculate how each parameter affects the measurements. Once these models of change have been trained for all hypothetical patterns of change, we use them to infer the posterior probability of which *pattern of change* in parameter(s) can best explain the change between real datasets. We call this approach BENCH, which stands for Bayesian EstimatioN of CHange.

When confronted with a degenerate biophysical model, BENCH makes a different set of assumptions from the traditional approach of fixing some parameters and

identifying any change in the remaining free parameters. When comparing patients and controls, the traditional approach assumes that the prior values for the fixed parameters hold across the region of interest in both groups. Hence, any change of signal across the region of interest between the two groups is assumed to be fully explained by the predetermined set of free parameters. In contrast, by not relying on model inversion, BENCH can work directly with the degenerate biophysical model without fixing any parameters. However, this comes at the price of limiting the change to some predetermined set of possible patterns set by the user (e.g., parameter A could change, or parameter B increases by the same amount as parameter C decreases). While the number of such proposed microstructural changes can be large, each of them has to be sparse (i.e., they have fewer degrees of freedom than the number of free parameters that could be estimated using the conventional approach). In this study, we limit ourselves to changes of just one parameter at a time for the sake of simplicity of explanation.

BENCH is applicable to any situation where we are interested in comparing the parameters of a generative (bio)physical model across different conditions. Here we apply the framework to diffusion MRI microstructure modelling. As an example use case, we studied microstructural changes in White Matter Hyperintensities (WMH), which are extra bright regions that are commonly seen in T2-weighted images at specific brain regions in elderly people. Despite the abundance and clinical implications of WMHs [18, 70], the underlying changes in the histopathology and microstructure remain unknown [71, 72].

The structure of this chapter is as follows. In the Theory section, we present the general inference method and how we train the models of change. In the Methods section, we cover the diffusion-specific materials including the computation of summary measurements that are used to represent diffusion data and the microstructural model for diffusion MRI. In the Results section, we first demonstrate the ability of our model in detecting the underlying parameter changes using simulated data. We then apply the method to study microstructural changes in white matter hyperintensities as an example application. In the Discussion

section, the potential applications, limitations, and possible future directions of this work are presented.

3.3 Theory

Table 3.1: Description of notations for theory section

Notation	Description	Dimension
$\mathcal{M}$	Generative biophysical model	$R^p \rightarrow R^d$
y	Baseline measurement	d
ν	Parameters of the model $\mathcal{M}$	p
Δy	Observed change in the measurement	d
$\hat{\Delta\nu}$	*Pattern of change*, unit vector in the parameters space*	p
$\lvert\nu\rvert$	The amount of change in the direction defined by $\hat{\Delta\nu}$	scalar
$\nabla_{\hat{\Delta\nu}}\mathcal{M}(\nu)$	The derivative of $\mathcal{M}$ in the direction of $\hat{\Delta\nu}$ at point ν	d
ϵ	Measurement noise Gaussian with zero mean and covariance Σ_n	d
$\mu_{\hat{\Delta\nu}}$	The average derivative of y in the direction $\hat{\Delta\nu}$	d
$\Sigma_{\hat{\Delta\nu}}$	The covariance of derivatives y in the direction $\hat{\Delta\nu}$	$d \times d$

* Hat notation is used to indicate unit vector.

3.3.1 Inference on change in parameters

Given a baseline vector of measurement (y), an observed change in the measurement (Δy), and a generative biophysical model ($\mathcal{M}$) which maps a vector of parameters (ν) to measurements, we introduce the concept of a *pattern of change* ($\hat{\Delta\nu}$). This is defined as a unit vector within the parameter space, that indicates a direction of change in parameters. In this section, our objective is to quantitatively estimate the probability that a *pattern of change* can explain the observed change in the measurements, i.e. $P(\hat{\Delta\nu} \mid y, \Delta y)$. Please note that all subsequent probability distributions are inherently conditioned on the baseline measurements, but for the sake of readability, they are only explicitly indicated when necessary.

A *pattern of change* is a unit vector in the parameter space, e.g. it can be a change in a single parameter or any linear combination of the model parameters. The user can specify any set of hypothetical patterns of change they deem plausible, and perform subsequent calculations for each one independently. For simplicity

of the explanations and notation, we only assume a single parameter change in the rest of this chapter, but all the equations apply to any linear combination of the parameters. Thus, considering that the model comprises p free parameters, we define p distinct patterns of change and compute the posterior probability for each one. If the model is invertible, we may directly estimate the parameter change ($\Delta\nu$, that entails the direction $\hat{\Delta\nu}$ and amount $|\Delta\nu|$) by inverting the model on y and $y + \Delta y$ to get the corresponding parameter estimates and calculate the differences. Alternatively, in BENCH we estimate the probability $P(\hat{\Delta\nu} \mid \Delta y)$ for the set of hypothesised patterns of change. Using Bayes' rule:

$$P(\hat{\Delta\nu} \mid \Delta y) = \frac{P(\Delta y \mid \hat{\Delta\nu})P(\hat{\Delta\nu})}{\sum_{\hat{\Delta\nu}'} P(\Delta y \mid \hat{\Delta\nu}')P(\hat{\Delta\nu}')} \tag{3.1}$$

Where $\hat{\Delta\nu}'$ is the set of all hypothetical patterns of change that are tested. We assume no prior preference between the patterns of change (i.e. $P(\hat{\Delta\nu}) = 1$ for all $\hat{\Delta\nu}$). However, it can be set to any arbitrary distribution by the user if they have prior knowledge about the type of change. To estimate the desired posterior probability in equation 3.1 up to a scale we need to estimate the likelihood term $P(\Delta y|\hat{\Delta\nu})$. The *pattern of change* $\hat{\Delta\nu}$ represents the direction but not the amount of the change in the parameters. We therefore marginalize the likelihood with respect to the amount of change ($|\Delta\nu|$):

$$P(\Delta y \mid \hat{\Delta\nu}) = \int P(|\Delta\nu|)P(\Delta y \mid \hat{\Delta\nu}, |\Delta\nu|)d|\Delta\nu| \tag{3.2}$$

We assume that the prior distribution for the amount of change follows a log-normal distribution with a fixed mean and scale parameter (adjustable hyper parameters). The choice of a log-normal distribution facilitates the prior to accommodate changes spanning various orders of magnitude.

The likelihood term inside the integral, $P(\Delta y|\hat{\Delta\nu}, |\Delta\nu|)$, defines how the measurements change as a result of a fully characterised vector of change in the parameters with the given direction ($\hat{\Delta\nu}$) and amount ($|\Delta\nu|$). To relate this parameter change to a change in data one also needs to know the baseline parameters (ν), as

$$\Delta y = \mathcal{M}(\nu + |\Delta\nu|\,\hat{\Delta\nu}) - \mathcal{M}(\nu) + \epsilon \tag{3.3}$$

where ϵ is the measurement noise. However, for a degenerate biophysical model, we cannot estimate a unique set of baseline parameters v for which we can estimate equation 3.3. While, one could integrate over all possible values of ν, this is a very high-dimensional integral, which would be very computationally expensive. Instead, we propose an alternative way to avoid the need of estimating the baseline parameters to estimate the likelihood in equation 3.2.

Assuming that $|\Delta\nu|$ is reasonably small, and $\mathcal{M}$ is behaving smoothly w.r.t ν, using a first order Taylor expansion we can express Δy as:

$$\Delta y = \nabla_{\hat{\Delta\nu}}\mathcal{M}(\nu) \times |\Delta\nu| + \epsilon \tag{3.4}$$

Where $\nabla_{\hat{\Delta\nu}}\mathcal{M}(\nu)$ is the derivative of $\mathcal{M}$ in the direction of $\hat{\Delta\nu}$ at point ν, and ϵ is the measurement noise. Given the baseline measurements (y), but not the baseline parameters (ν), there can be an infinite number of $\nabla_{\hat{\Delta\nu}}\mathcal{M}(\nu)$ for a degenerate model (Figure 3.2). To account for all instances of the derivatives, we model them as a random variable with the probability distribution $P(\nabla_{\hat{\Delta\nu}}\mathcal{M} \mid y)$. This distribution could be estimated from the posterior probability $P(\nu \mid y)$, however, if the model is degenerate it would be difficult to estimate the posterior. Instead, we model it as a Gaussian distribution with mean $\mu(y)$ and covariance $\Sigma(y)$, i.e.

$$P(\nabla_{\hat{\Delta\nu}}\mathcal{M} \mid y) = Normal(\mu_{\hat{\Delta\nu}}(y), \Sigma_{\hat{\Delta\nu}}(y)) \tag{3.5}$$

where $\mu_{\hat{\Delta\nu}}$ represents the average expected change in the measurements as a result of change in parameters in the direction $\hat{\Delta\nu}$, $\Sigma_{\hat{\Delta\nu}}$ represents the uncertainty around this expectation due to the unknown baseline parameters (Figure 3.2). This particular formulation facilitates the transition of uncertainty from the baseline parameters to the measurement space, thus enabling its effective modelling and prediction.In the next section we will describe a method for estimating $\mu_{\hat{\Delta\nu}}(y)$ and $\Sigma_{\hat{\Delta\nu}}(y)$ by training regression models using simulated data. Once we compute these hyperparameters, by inserting equation 3.5 back into equation 3.4 we can compute the likelihood term inside the integral by

$$P(\Delta y \mid \hat{\Delta\nu}, |\Delta\nu|) = Normal(|\Delta\nu|\,\mu_{\hat{\Delta\nu}}(y), |\Delta\nu|^2\,\Sigma_{\hat{\Delta\nu}}(y) + \Sigma_n) \tag{3.6}$$

where Σ_n is the noise covariance matrix.

Finally, by numerical approximation of the integral over the amount of the change in equation 3.2, we are able to approximate the likelihood function $P(\Delta y \mid \hat{\Delta\nu})$ which we can then use in equation (3.1) yielding the desired posterior distribution on the patterns of change. Figure **??** b) provides a schematic representation of the necessary steps to compute the posterior as detailed in equation 3.2. Moreover, using the approximation of the likelihood function in equation 3.6 the posterior probability of the amount of change for each direction is proportional to

$$P(|\Delta\nu| \mid \Delta y, \hat{\Delta\nu}) \propto P(\Delta y \mid \hat{\Delta\nu}, |\Delta\nu|) P(|\Delta\nu|) \tag{3.7}$$

Note that this likelihood is unnormalized so a high or low value doesn't necessarily reflect the quality of the change model in explaining the data. For such measure please refer to appendix 3.9. We can still estimate the most likely amount of change in the parameter given the measurements by finding the $|\Delta\nu|$ that maximizes the above posterior probability (maximum a posteriori estimation).

3.3.2 Training models of change

In this section we describe how to train a regression model to estimate the hyperparameters of the distribution $P(\nabla_{\hat{\Delta\nu}}\mathcal{M} \mid y)$, namely the average $(\mu_{\hat{\Delta\nu}}(y))$ and uncertainty $(\Sigma_{\hat{\Delta\nu}}(y))$ of change in the measurement (y) for a *pattern of change* $(\hat{\Delta\nu})$.

Given some baseline parameters (ν) one can calculate the baseline measurements as $y = \mathcal{M}(\nu)$ and approximate the derivative in direction $\hat{\Delta\nu}$ using

$$\nabla_{\hat{\Delta\nu}}\mathcal{M}(\nu) \approx \lim_{t \to 0} \frac{\mathcal{M}(\nu + t\hat{\Delta\nu}) - \mathcal{M}(\nu)}{t} \tag{3.8}$$

Therefore, by sampling ν from the parameter space using a prior distribution, we generate a simulated dataset of pairs $[y, \nabla_{\hat{\Delta\nu}}\mathcal{M}]$ that we use for training regression models.

We use a regression model parametrised by $w_{\mu_{\hat{\Delta\nu}}}$ to estimate $\mu_{\hat{\Delta\nu}}$ as:

$$\mu_{\hat{\Delta\nu}}(y; w_{\mu_{\hat{\Delta\nu}}}) = F(y).w_{\mu_{\hat{\Delta\nu}}} \tag{3.9}$$

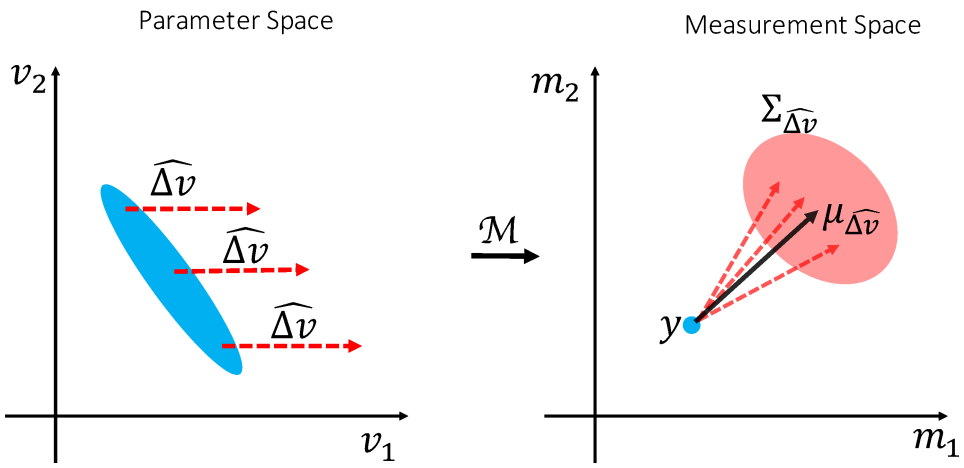


Figure 3.2: Distribution of gradients. The way measurements change as a result of a particular change in the parameters can only be calculated if we know the baseline parameters. When we are only given the measurements, there are several instances of equally likely gradient directions depending on the underlying baseline parameters. We model all of these gradients given the baseline measurements as a random variable with a presumed distribution. This allows us to transfer the uncertainty due to the inverse model degeneracy into the measurement space. The blue oval in the parameter space (left) represents all the parameter settings that map onto the same blue point in the measurement space(right). Each of these parameter settings can produce a different gradient direction in the measurements space. The collection of such gradients of change $\hat{\Delta\nu}$ for the measurement y are modelled as a Gaussian distribution with mean $\mu_{\hat{\Delta\nu}}(y)$ and covariance $\Sigma_{\hat{\Delta\nu}}(y)$.

where $F(y)$ is the design matrix, which depends on arbitrary affine or non-linear transformations of y. Note that the subscript $\mu_{\hat{\Delta\nu}}$ of the weights indicates that each *pattern of change* in the parameters has its own set of weights.

We also employ a regression model for the uncertainty hyperparameter $\Sigma_{\hat{\Delta\nu}}$ parametrised by $w_{\Sigma_{\hat{\Delta\nu}}}$. However, $\Sigma_{\hat{\Delta\nu}}$ must be positive definite, which would not be guaranteed when directly estimating $\Sigma_{\hat{\Delta\nu}}$ by training an element-wise regression model. To account for the positive definite nature of $\Sigma_{\hat{\Delta\nu}}$, we instead train regression models for elements of the lower triangular matrix of its Cholesky decomposition (L). Also, since the diagonal elements of the lower-triangular matrix in Cholesky decomposition must be non-negative, we use their log-transform in the regression model. Hence

$$\Sigma_{\hat{\Delta\nu}}(y; w_{\Sigma_{\hat{\Delta\nu}}}) = \mathcal{T}(F(y).w_{\Sigma_{\hat{\Delta\nu}}}) \tag{3.10}$$

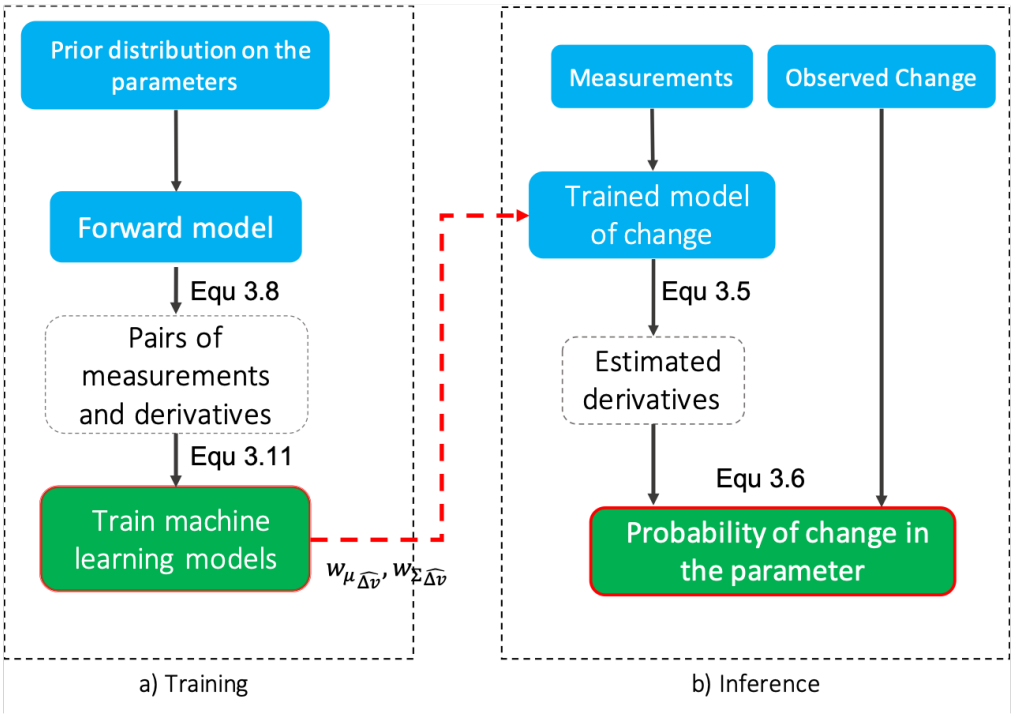


Figure 3.3: Flowchart for training and inference using change models. The blue, white and green blocks indicate user defined inputs, intermediate variables and outputs respectively. In the training phase for each parameter change, samples that are drawn from the provided prior distribution are passed through the forward model to estimate pairs of measurements and derivatives. Then, regression models are trained to estimate the distribution of derivatives given the measurements using a maximum likelihood estimation. This phase does not require real data and needs to be done only once. In the inference stage using these trained models we estimate the distribution of the derivatives for any given baseline measurements. We then calculate the posterior probability that change in each parameter caused the change in the measurements using the derivative distributions. The relevant equation for each operation is shown next to it.

where $\mathcal{T}$ denotes the transformation of the regressed vector to the full covariance matrix that includes the arrangement of elements, exponentiation of the diagonals, and the matrix multiplication for inverse Cholesky decomposition.

Putting back the above regression models into equation 3.5 the likelihood of observing pairs of baseline measurements and gradients in terms of the parameters of regression models is:

$$L(w_{\mu_{\hat{\Delta\nu}}}, w_{\Sigma_{\hat{\Delta\nu}}}) = \prod_i Normal(\nabla_{\hat{\Delta\nu}}\mathcal{M}_i; F(y_i).w_{\mu_{\hat{\Delta\nu}}}, \mathcal{T}(F(y_i).w_{\Sigma_{\hat{\Delta\nu}}})) \tag{3.11}$$

Accordingly, we estimate the optimal weights $w_{\mu_{\hat{\Delta\nu}}}, w_{\Sigma_{\hat{\Delta\nu}}}$ by maximizing the above likelihood function for the simulated pairs of $[y_i, \nabla_{\hat{\Delta\nu}}\mathcal{M}_i]$ using a combination

of the BFGS and Nelder-Mead methods as implemented in SciPy [73].

This procedure is repeated for each hypothetical *pattern of change*, yielding two sets of weights for the average and uncertainty of change, which we refer to as a " change model". Once we estimated these weights, for any given baseline measurement we use the regression models in equations 3.9 and 3.10 to estimate the distribution of derivatives and then the desired probability distributions. Figure **??** shows a schematic overview of the inputs, outputs and steps that are required to train a change model, as well as how to use them to infer the change in parameters.

In this work, we used a second degree polynomial function of the data for the regression models that estimate the mean change ($\mu_{\hat{\Delta\nu}}$) from the baseline measurements. For the uncertainty parameter ($\Sigma_{\hat{\Delta\nu}}$) a first degree (linear) model is chosen as we expect less variability across samples for this hyperparameter. The weights for the regression models were estimated using a maximum likelihood optimization and a training dataset with 100,000 simulated samples.

3.3.3 Biophysical model of diffusion

In this section we explain the biophysical model of diffusion that we used to model brain microstructure with diffusion MRI data. The diffusion signal S in the brain is conventionally modelled as the sum of signals from multiple compartments. We will here adopt the three-compartment standard model [64] consisting of an isotropic free water (denoted by the subscript "iso"), an intra-axonal ("in"), and an extra-axonal ("ex") compartment:

$$S = S_{iso}A_{iso} + S_{in}A_{in} + S_{ex}A_{ex} \tag{3.12}$$

where S_i represents the baseline signal contribution (at $b = 0$), and A_i represents the signal attenuation due to the diffusion weighting in each compartment (Figure 3.4).

The attenuation for the isotropic compartment is modelled as an exponential decay:

$$A_{iso} = e^{-bd_{iso}} \tag{3.13}$$

where d_{iso} is the diffusion coefficient of free water.

The intra-axonal compartment is modelled as a set of dispersed identical sticks with no perpendicular diffusivity. The stick response function for gradient direction g and b-value b is given by

$$R(b, g; \mu, d_{a_{in}}) = e^{-bd_{a_{in}}(\mu^T g)^2} \tag{3.14}$$

where $d_{a_{in}}$ is the diffusion coefficient along the orientation of the stick μ.

The fibre Orientation Distribution Function (fODF) is modelled with a Watson distribution, which is defined as

$$f(x) = \frac{1}{c} e^{\kappa(\mu^T x)^2} \tag{3.15}$$

where μ is the average orientation, κ is the concentration coefficient and c is a normalization constant. To assimilate the dispersion coefficient to the notion of variance and limit it to a bounded range, we use the change of variable from κ to Orientation Dispersion Index (ODI) as $ODI = \frac{2}{\pi} \arctan(\frac{1}{\kappa})$. Unlike κ which is unbounded, ODI is limited to the range $(0, 1)$, where higher ODI values correspond to more dispersion. So, the diffusion signal for this compartment is the spherical convolution of the fibre response function with the Watson ODF:

$$A_{in} = \iint_{S^2} e^{-bd_{a_{in}}(g^T n)^2} \frac{1}{c} e^{cot(\frac{\pi}{2} ODI)(\mu^T n)^2} dn \tag{3.16}$$

where the integral is over the surface of the unit sphere S^2 representing all possible fibre orientations in 3D.

The extra-axonal compartment is modelled similar to the intra-axonal compartment, with the addition of a non-zero diffusion perpendicular to the fibre orientation. The fibre response function in this case is given by

$$R = e^{-b[d_{a_{ex}}(\mu^T g)^2 + d_{ex,r}(1-(\mu^T g)^2)]} \tag{3.17}$$

where $d_{ex,r} \leq d_{a_{ex}}$ are the radial and axial diffusion coefficients. To avoid this dependence between the diffusivity parameters, the parameter τ defined as the ratio of perpendicular to parallel diffusivity is used as a substitute to $d_{ex,r}$. The free parameter τ - subject to $\tau \in [0, 1]$ to maintain the inequality constraint for

the diffusivities - can be considered as a measure of tortuosity as it measures the extent to which water diffusion perpendicular to the fibre orientation is hindered with respect to the parallel diffusion. Therefore, the fibre response function for the extra axonal compartment is

$$R = e^{-bd_{a_{ex}}[(\mu^T g)^2 + \tau(1-(\mu^T g)^2)]} \tag{3.18}$$

As the compartments share the same geometry, the same fibre orientation distribution is used. Accordingly, the signal attenuation for extra-axonal compartment is given by

$$A_{ex} = \iint_{S^2} e^{-bd_{a_{ex}}[(\mu^T g)^2 + \tau(1-(\mu^T g)^2)]} \frac{1}{c} e^{cot(\frac{\pi}{2}ODI)(\mu^T n)^2} dn \tag{3.19}$$

We use the confluent hypergeometric function of the first kind with matrix argument to compute the integrals for both intra and extra axonal compartments similar to [15].

Table 3.2 summarises all the free parameters of the described biophysical model along with their valid range.

Parameter	Description	Range
s_{iso}	Signal fraction for isotropic (free water) diffusion compartment	$[0, 1]$
s_{in}	Signal fraction for intra-axonal compartment	$[0, 1]$
s_{ex}	Signal fraction for extra-axonal compartment	$[0, 1]$
d_{iso}	Isotropic (free water) diffusivity coefficient	$[0, \infty]$
$d_{a_{in}}$	Parallel diffusivity for the intra-axonal compartment	$[0, \infty]$
$d_{a_{ex}}$	Parallel diffusivity for the extra-axonal compartment	$[0, \infty]$
τ	radial to axial diffusivity ratio for the extra-axonal compartment	$[0, 1]$
ODI	Orientation dispersion index	$[0, 1)$

Table 3.2: Microstructural parameters of the diffusion model. All diffusion coefficients are in $\mu m^2/ms$

3.3.4 Summary measurements

Diffusion MRI data are usually measured in multiple shells to capture tissue properties that are sensitive to diffusion of water molecules at various spatial

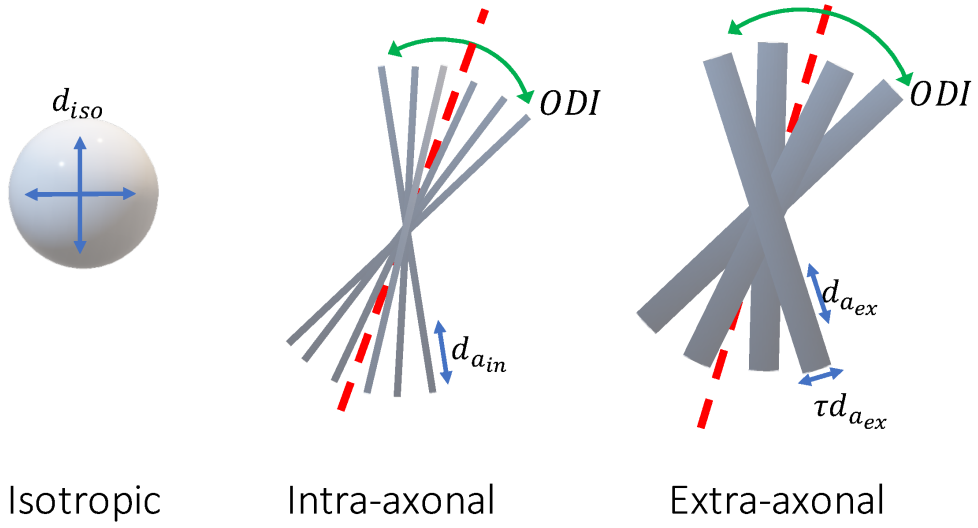


Figure 3.4: Compartments of the diffusion model. We use a three compartment model that can describe diffusion MRI signals from various brain tissues namely CSF, white matter and gray matter. The isotropic compartment models unrestricted diffusion of water molecules outside of tissue (CSF) with a single free parameter d_{iso}. The intra-axonal compartment models the diffusion of water within axons as several sticks with identical parallel diffusivity parameter $d_{a_{in}}$, and zero radial diffusivity, that are dispersed by a Watson distribution with orientation dispersion index ODI. The extra-axonal compartment is also a Watson dispersed zeppelin with parallel diffusivity $d_{a_{ex}}$ and perpendicular diffusivity $d_{ex,r} = \tau d_{a_{ex}}$. Including the signal fraction parameters $(s_{iso}, s_{in}, s_{ex})$ this model has 8 free parameters, which are more than that can be fitted to a conventional dMRI data.

scales. Within each shell, gradients are applied in several directions to measure the geometrical structure of the tissue. However, since we are only interested in the microstructural characteristics, any orientation-related information is irrelevant. We therefore need summary measurements from each shell that are invariant to orientations. We create these summary measurements using real spherical harmonics, which are analogous to the Fourier transform for the spherical domain.

Spherical harmonics are a complete set of orthonormal functions over the surface of a unit sphere. That is to say, any bounded real function that is defined over the unit sphere can be represented by a unique linear combination of these functions with real coefficients. Each real spherical harmonic is denoted by $Y_{l,m}(\theta, \phi)$ where $l = 0, 1, 2, ...$ is the degree and $m = -l, ..., l$ is the order, and $\theta \in [0, \pi]$, $\phi \in [-\pi, \pi]$ are the polar and longitudinal angles in standard spherical coordinate system

respectively. The diffusion signal at each shell is decomposed as:

$$S(\theta, \phi) = \sum_{l=0}^{\infty} \sum_{m=-l}^{l} C_{l,m} Y_{l,m}(\theta, \phi) \tag{3.20}$$

Since the harmonics are a linear basis, one can easily calculate the coefficients for the signal in each shell by inverting the design matrix formed by the harmonics sampled at the gradient directions.

The coefficients are not orientationally invariant. However, the total power in each degree, which is defined as the vector norm of all the corresponding coefficients, is rotationally invariant [12, 74, 75]. Also, since the diffusion signal is symmetric around the origin and the harmonics of odd degree are odd functions (anti-symmetric w.r.t origin), all odd degrees have zero coefficients.

Consequently, for each shell of diffusion data, we calculate the mean squares of all coefficients for degrees $l = 0, 2, 4, ...$ as the orientationally-invariant summary measurements.

$$y_l = \frac{1}{2l+1} \sum_{m=-l}^{l} C_{l,m}^2 \tag{3.21}$$

The mean is chosen over the norm to make the scale equal across all degrees. For the case of $l = 0$, we simply use the only coefficient (without the square), so that it represents the mean signal. The higher order summary measurements quantify the signal anisotropy; with greater l being more sensitive to sharper changes. We used a logarithm transformation on the power of the coefficients to make the distribution of the measurements for real data closer to a Gaussian and also being more sensitive to smaller changes.

3.4 Methods

3.4.1 Simulations

For all the simulations we used the acquisition protocol conducted by the UK Biobank (UKB) [76, 77] which includes two shells of diffusion ($b = 1, 2 \frac{\text{ms}}{\mu\text{m}^2}$) with linear diffusion encoding. Each shell consists of 50 gradient directions distributed

uniformly over the surface of the unit sphere, in addition to 5 acquisitions with $b = 0$, yielding a total of 105 measurements.

We used the rotationally invariant summary measurements computed from spherical harmonics for signal representation. The summary measurements for each shell are norms of coefficients at $l = 0$ (absolute value) and $l = 2$ (log mean squared). This produces 5 rotational invariant summary measurements from a diffusion data, namely *b0-mean, b1-mean, b1-l2, b2-mean*, and *b2-l2*.

The described standard model for diffusion is used for both simulated test data and for training models of change. The prior distributions for the parameters are shown in figure 3.5. We note that these priors are not used for constraining the model parameters but rather they are used to generate training samples for the regression models. The choice of the prior distributions is arbitrary as long as they can reflect all hypothetical parameter combinations that can produce measurements similar to real data.

The standard model is not invertible given a conventional multishell diffusion data with linear diffusion encoding [14, 64] . Typically, additional constraints are imposed to render the model invertible, e.g. in NODDI [10], the diffusion coefficients are fixed to a prior value as follows:

$$d_{iso} = 3\frac{\mu m^2}{ms}, d_{a_{in}} = d_{a_{ex}} = 1.7\frac{\mu m^2}{ms}$$

Additionally, the tortuosity parameter τ is coupled to the signal fractions:

$$\tau = \frac{s_{in}}{s_{in} + s_{ex}} \tag{3.22}$$

Accordingly, this constrained model has four free parameters: s_{iso}, s_{in}, s_{ex} and ODI.

For both the constrained and unconstrained models, we generated a test dataset containing pairs of simulated diffusion signals, such that in each pair at most one microstructural parameter is different. To generate each pair, we sample a baseline parameter setting from the prior distributions and change one of the parameters by an effect size of 0.1. We also generate pairs of data where no parameter changes and the difference between the two samples is only due to the addition of noise. We

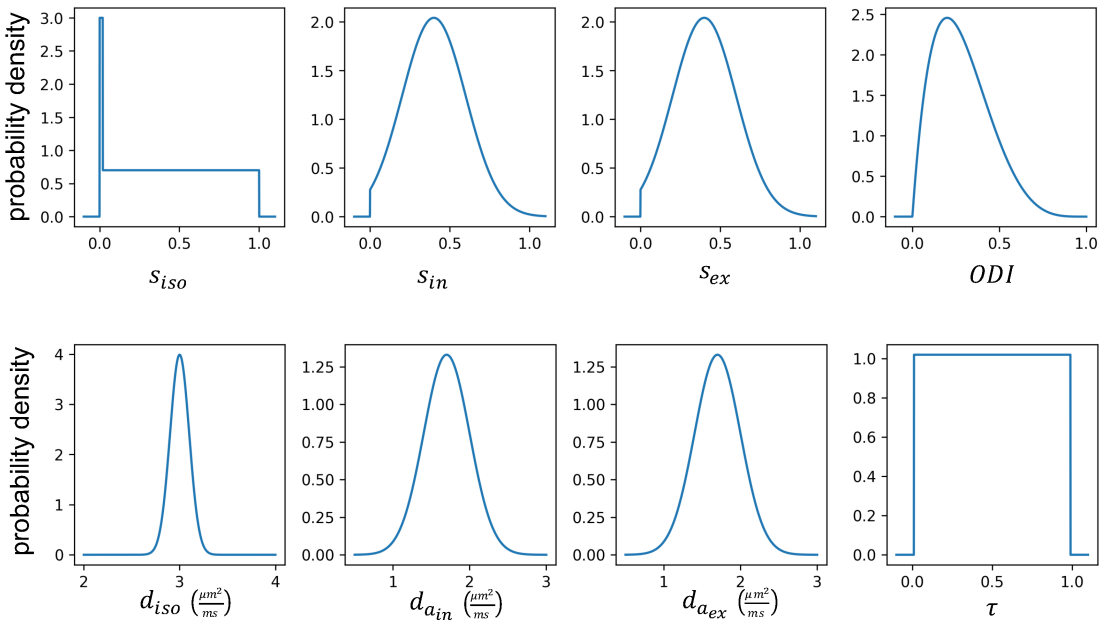


Figure 3.5: Prior distributions for the parameters of the standard model. These priors are used for generating pairs of measurements and gradients for training the models of change. Also, the same priors are used for simulating test datasets. The priors are chosen such that they contain all probable parameter combinations that can produce measurements similar to real data. The delta function along with uniform distribution in the isotropic signal fraction is used to model pure tissue types as well as partial volume effect. In the training phase, the signal fractions are normalized to sum up to 1. A beta (shape parameters $\alpha = 2, \beta = 5$) distribution is used for ODI to impose a nearly uniform distribution for effective fibre dispersion. The prior for isotropic and axial diffusivities are normal distributions with mean 3 and 1.7 ($\frac{\mu m^2}{ms}$) and standard deviation 0.1 and 0.3 respectively; as we expect faster diffusion as well as less variability in the free water component.

then apply the forward model to both parameter settings to produce diffusion MRI signals. Gaussian noise with standard deviation $\sigma_n = 0.01$ (SNR=100) is added to all diffusion signals. For a typical single-subject SNR of around 30-40, an SNR of around 100 would already be reached when averaging across ∼10 subjects.

The signal fraction parameters are constrained to sum up to 1 for training models of change. Note that whilst this imposes a constraint that the *b0-mean* for the baseline measurement is equal to 1, it does not constrain a *change* in that summary measurement. Accordingly, all the summary measurements (both in the baseline and the change) are normalized by the *b0-mean* of the baseline measurement for any real data. This differs from the parameterization in conventional NODDI, where

there is a constraint on the signal fractions to sum up to 1, and add a separate b0 parameter that is directly estimated from b0 signal. Instead, here we assume all the signal fraction parameters can change independently.

For the direct inversion approach, a maximum a posteriori algorithm is employed to estimate the parameters of the constrained model from each diffusion signal separately. Then using a z-test across the parameter estimates in each pair, we calculate a p-value for the change in each parameter (corrected for multiple comparisons across parameters). The parameter with the minimum p-value is identified as the changed parameter. All the cases with minimum $p > 0.05$ are identified as no change.

We also used BENCH for identifying change on the same dataset. To estimate the noise covariance in the summary measurements Σ_n, 100 noisy instances of signals were generated, and the sample covariance of the difference between summary measurements in each pair was estimated. We then estimated the posterior probability of change in each parameter using the trained models of change. The *no change* model has a zero mean and covariance Σ_n everywhere. The change model with the maximum posterior probability is selected as the predicted change.

3.4.2 White matter hyperintensities

We investigate the possible microstructural changes in white matter hyperintensities (WMH) using BENCH and model inversion. In this experiment, we used diffusion MRI of 3000 randomly selected subjects from the UK biobank dataset. To account for the variability in overall intensity across subjects, we divided each subject's diffusion data by the average intensity of the b0 image across the brain's white and grey matter extracted using FSL FAST [78]. We then computed the spherical harmonics-based summary measurements from the diffusion MRI data for each subject and interpolated these measures into the standard MNI space using non-linear transformations estimated by FSL FNIRT [79, 80].

Segmentations of the WMHs were generated from T2 FLAIR images using FSL's BIANCA [81] as part of the UK Biobank pipeline [76]. We computed the average

summary measurements for Normally Appearing White Matter (NAWM) that are voxels within the white matter mask not classified as WMH and the WMHs for all voxels that included more than 10 subjects with WMH. For each voxel, subjects were split into two groups according to whether the voxel has been classified as WMH or not. Averaging the summary measures within groups provides us with the baseline measurement (y) and the observed change (Δy) related to WMH. The noise covariance (Σ_n) in each voxel was estimated using the within group covariance matrix divided by the number of subjects in the normal appearing white matter group.

3.5 Results

3.5.1 Summary measurements

A representative axial slice of the normalized summary measurements from a single subject are shown in Figure 3.6. The *"mean"* summary measures represent the normalised average signal. The $l2$ measures quantify the anisotropy in each voxel (similar to Fractional Anisotropy maps in DTI).

The bottom panels of Figure 3.6 show histograms of the summary measurements across the brain for the same subject, as well as distributions of simulated data based on prior distributions over the model parameters. The distribution for the generated samples fully covers the range of the data and follows a very similar density distribution. This verifies that the prior distributions are wide enough to capture the full range of real data.

Figure 3.7 shows estimated derivatives of the summary measurements at baseline data representative of putative voxels in the white matter and grey matter. The error bars show estimated standard deviations of the derivatives (the square root of diagonals of the estimated covariance matrix). This variance is reflecting the uncertainty in the underlying parameters that can generate these measurements, as well as residuals of the regression model for the mean.

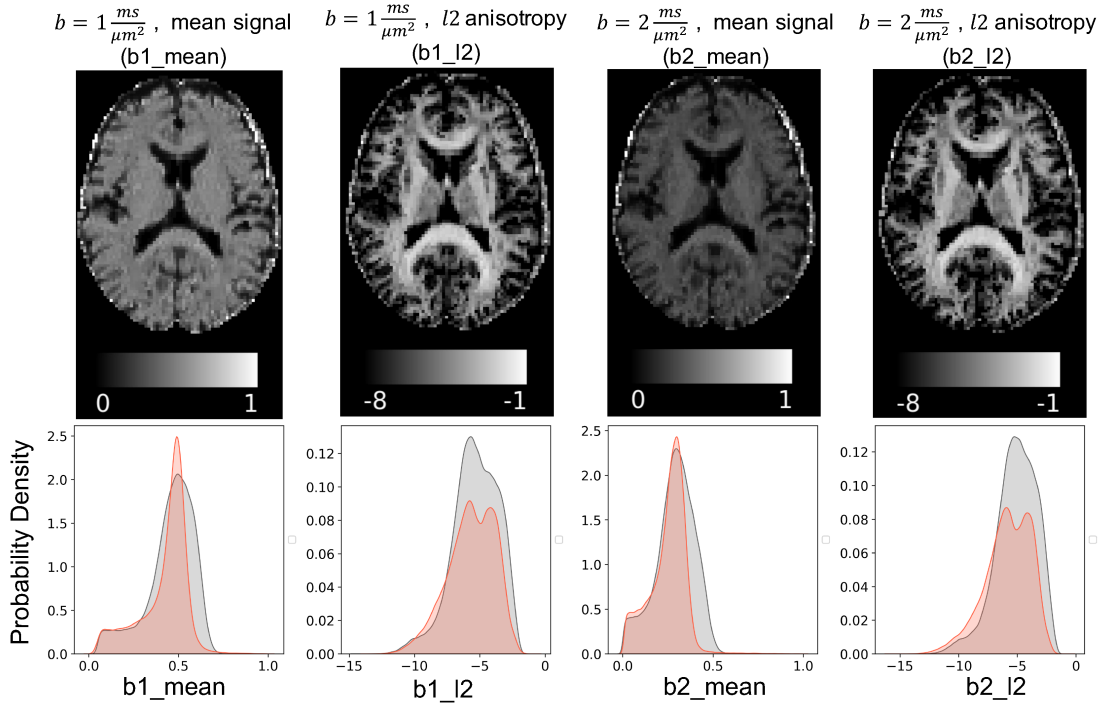


Figure 3.6: Maps of the summary measurements for a sample subject in the UK biobank dataset (top) and their histogram (bottom). The mean summary measurements is reflecting the average (across directions) diffusivity in each shell. The $l2$ summary measurements estimate the anisotropy, which is similar to the fractional anisotropy (FA), but computed with a linear transformation of the signal. Histograms show the distribution of these measurements across the brain; as well as the distribution of simulated data using the standard model and provided prior distributions. This shows that the simulations capture the full range of the summary measures from real data.

3.5.2 Validation

We first employed simulated data to evaluate the performance of the proposed approach in inferring microstructural changes from diffusion MRI data. The details of experiment parameters are provided in the methods section.

3.5.2.1 Comparison with model inversion

Figure 3.8a shows the confusion matrix using model inversion (left), and our inversion-free approach (right) for an invertible model with only 4 free parameters. Each element of these matrices represents the percentage of times a change in the parameter represented at the corresponding column is identified as a change in the corresponding row. Both approaches were able to detect the true parameter

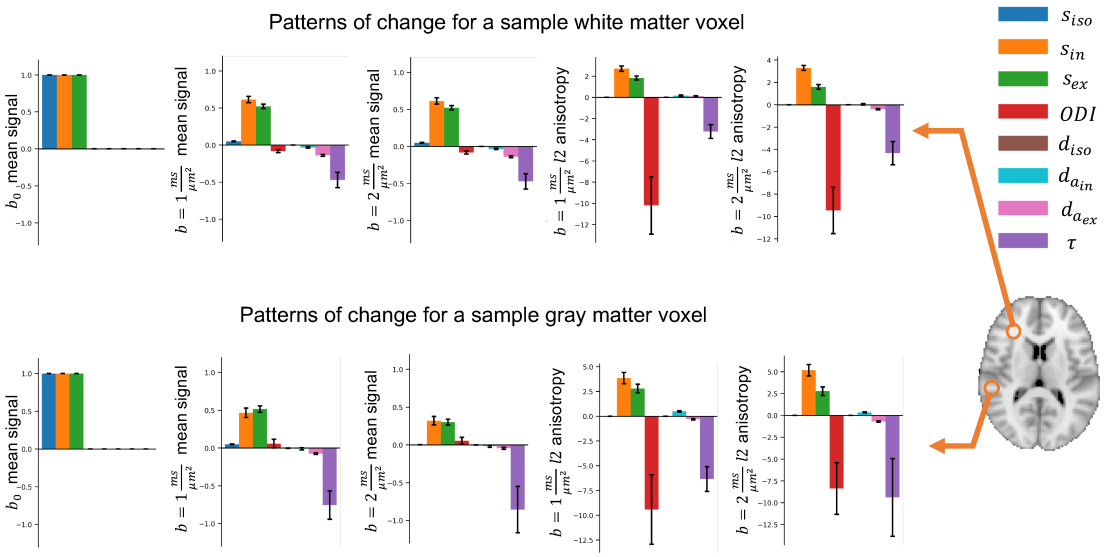


Figure 3.7: The estimated amount of change in the summary measurements as a result of a unit change in each parameter ($\mu_{\hat{\Delta\nu}}$) for a sample white matter and grey matter voxel. The error bars show the estimated standard deviation of change. colours correspond to parameters and columns indicate summary measurements. Due to differences in the baseline, each voxel can have a different change vector for the same parameter change. This added degree of freedom can model the variability of parameters (e.g. diffusivities) across the brain, which is not considered in constrained models; e.g. NODDI.

change in most of the cases.

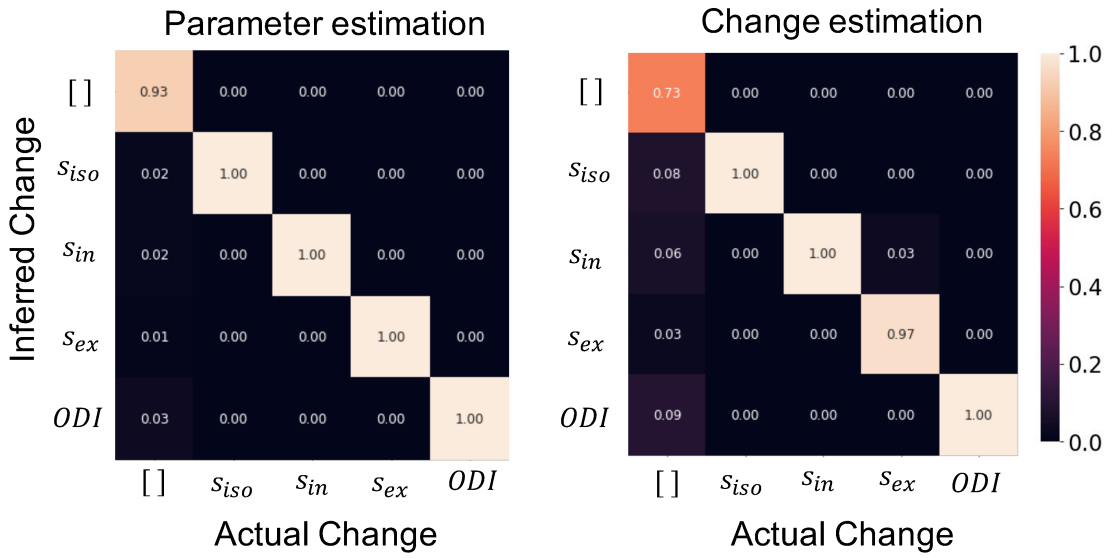
For the standard model with all 8 free parameters, Figure 3.8b shows the confusion matrices using the direct model inversion (left) and change estimation (right). Since the uncertainties of the parameter estimates are very large due to the model degeneracies, almost all of the changes are confused with *no change* when using direct inversion. However, the inversion-free approach is able to identify changes in s_{iso}, s_{in}, s_{ex} and ODI. Although, there is confusion between the remaining parameters compared to the restricted model, here we do not make any strong assumptions on the value of those parameters. Also, most of the confusions for these parameters are between them, meaning that we are able to distinguish a change in those parameters (e.g. the diffusivity parameters) from others. Change in isotropic diffusivity is mostly confused with the *no change* model. This is due to the b-values in the UKB protocol which are too high for this parameter; a change in this parameter has minimal effect on the signal.

3.5.2.2 Sensitivity to change in each parameter

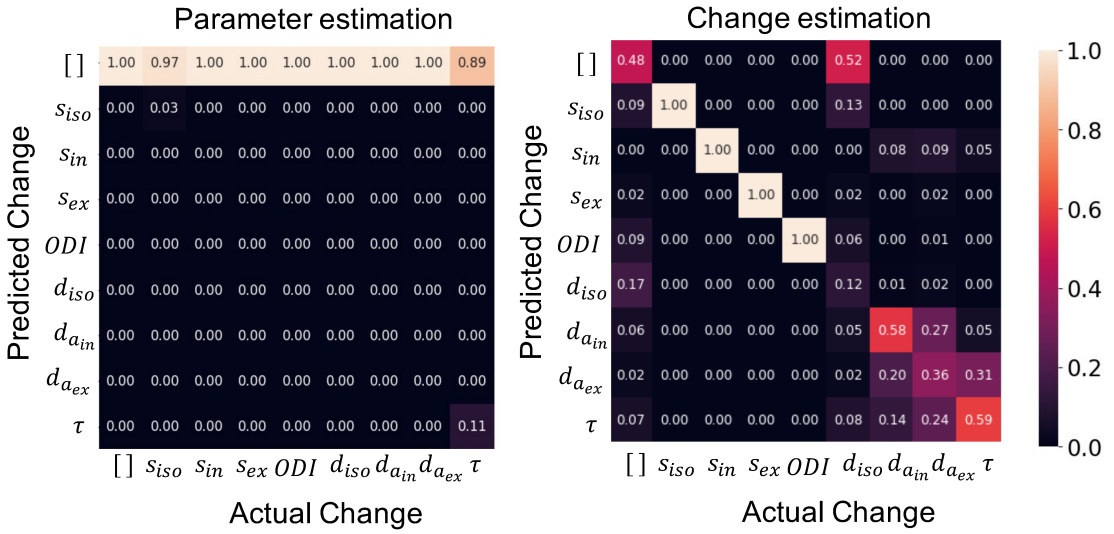
To evaluate the sensitivity of the approach to the amount of change in each parameter, we generated test datasets with variable effect sizes starting from 0 to 0.1 with step sizes of 0.01. Figure 3.9a shows the average posterior probability of change in each parameter versus the effect size. In all types of change, at very small effect sizes (< 0.01) the change is confused with no change, but as the effect size increases the probability of identifying the true change (red curves) increases. Changes in all signal fraction parameters and in the fibre dispersion are identified with high accuracy even at very small effect sizes. However, changes in diffusivity parameters are confused with each other (but not with signal fraction parameters) even at larger effect sizes. It is worth mentioning that effect size and SNR are two important factors (both unknown in real data) that affect the performance of detection in a similar way. So, when SNR is higher the approach can be more sensitive to smaller changes. Here we show the results for SNR=100.

3.5.2.3 Estimating the amount of change

So far we have only examined the posterior probabilities relating to the identity of the parameters that can best explain a change. However our framework also allows us to estimate the posterior probability on the *amount* of change for each parameter $P(|\Delta\nu| \mid dy, \hat{\Delta v})$ (Equ.3.7). Figure 3.9b shows the estimated (maximum a posteriori estimation) versus actual change in each parameter for different effect sizes.

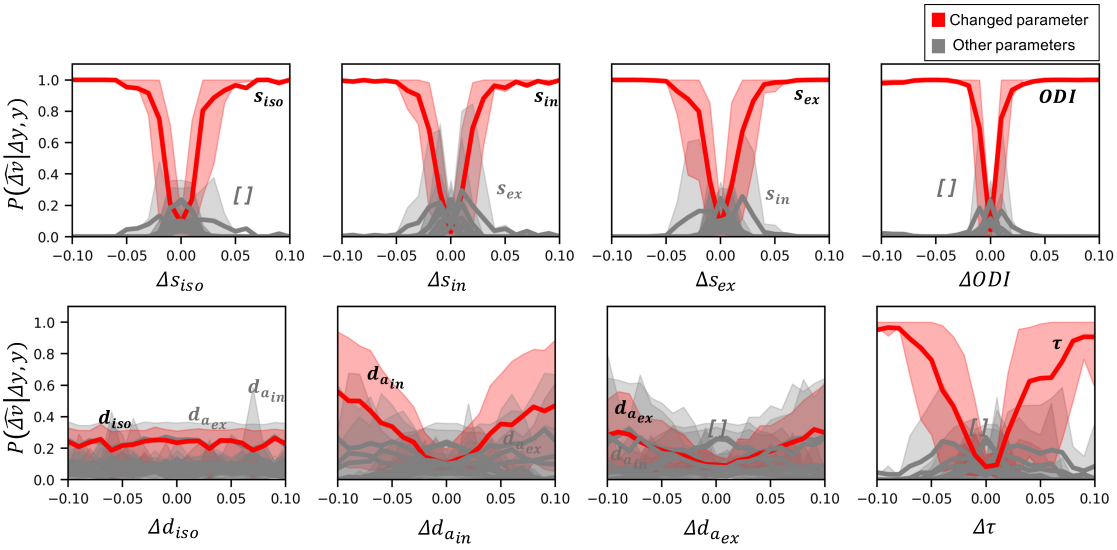


(a) Confusion matrices for the constrained model using parameter estimation and BENCH.

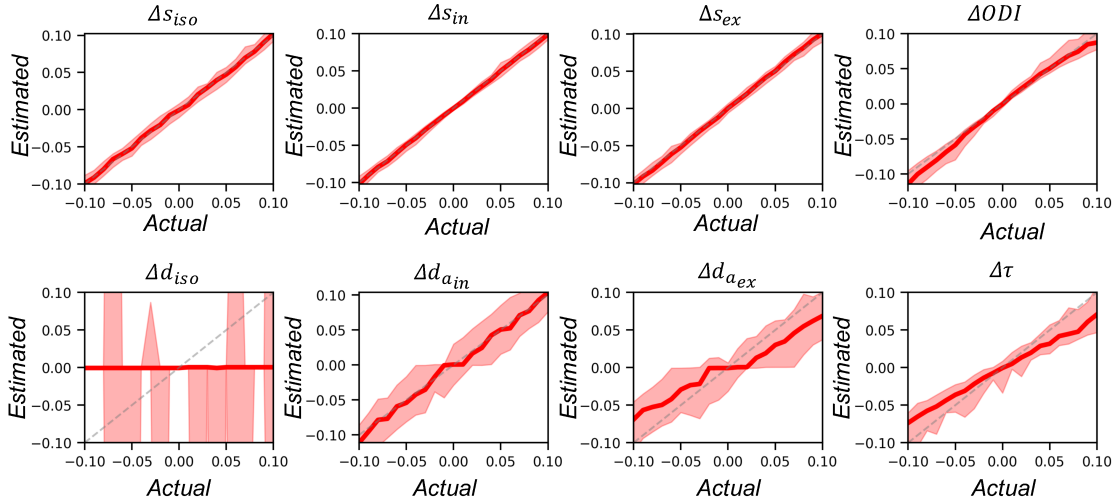


(b) Confusion matrices for the full model using parameter estimation and change estimation.

Figure 3.8: a) The numbers indicate the percentage of time a change in the corresponding column is identified as a change in the corresponding row. The diagonal elements show the accuracy in identifying true change. a) Both of the approaches performed near to ideal in detecting the true change in the case of constrained model. The change estimation has more false positives, but unlike the inversion approach, we did not explicitly define a false positive rate threshold. b) Given diffusion data at few shells, the full model is not invertible, i.e. the parameter estimates have a high variance. Therefore, almost no significant change is detected using parameter estimates. On the other hand, the change estimation approach can still identify changes in all the parameters of the restricted model. Although there remains confusion between a subset of the parameters since these have similar effects on the diffusion signal.



(a) Estimated probability of change in each parameter.



(b) Estimated amount of change in each parameter.

Figure 3.9: a) Each plot shows the estimated probabilities when the corresponding parameter on the $x-axis$ has changed between two datasets. Red curves show the average posterior probability of change in the actually changed parameter versus the amount of change. The gray curves show the probability for other parameters. Shaded areas show the 10 to 90 percentile range. Larger absolute amount of change results in higher posterior probability for the true parameter change. Change in the signal fraction parameters and ODI is distinguishable for effect sizes as small as 0.05. However, changes in diffusivity parameters even at very large effect sizes is cluttered with other parameters. **b)** Each plot shows the maximum a posteriori estimation of the amount of change vs the actual change in the parameter. The shaded areas show the 10 to 90 interval. The estimated change in the signal fractions follow the identity line (dashed gray line). The estimated change in d_{iso} is mostly around zero with a high variance as the posterior distribution is very flat and symmetric around zero. The change in $d_{a_{ex}}\tau$ and ODI is systematically biased at higher effect sizes.

3.5.3 White matter hyperintensities

3.5.3.1 Model Inversion

We inverted the NODDI model using non-linear fit implemented in DMIPY [82] in all subjects and ran a voxel wise glm to estimate the differences between white matter hyperintensities and normally appearing white matter (NAWM). Unlike in BENCH, NODDI requires fixing the diffusivity parameters. Usually, they are fixed to $d_{a,in} = d_{a,ex} = 1.7\frac{\mu m^2}{ms}$ (and $d_{iso} = 3.0\frac{\mu m^2}{ms}$). However, it has been recently suggested that the axial diffusivity should be higher based on several studies attempting to directly measure their value [83, 84]. We have therefore run the same analysis also with $d_{a,in} = d_{a,ex} = 2.5\frac{\mu m^2}{ms}$.

The z-maps for the contrast of WMH vs the baseline for all the parameters are shown in Figure 3.10. The strongest changes are seen in f_{intra} and it is consistent in both high and low diffusivity regimes. The direct inversion also suggests changes in the other two parameters (f_{iso} and ODI). However, interestingly, changing the pre-specified diffusivities in NODDI alters the story for f_{iso} and ODI which go in opposite directions(see scatter plots in 3.10 many points(voxels) lie in the 2nd or 4th quarter). These results demonstrate that the choice of fixed parameter values can affect the inferred change in other parameters.

3.5.3.2 BENCH

We used the trained models of change on the parameters of the full (standard) model to infer changes in WMH. Figure 3.11a shows the observed change in the summary measurements (normalized by b0 mean of the baseline) in white matter hyperintensities (dashed line) as well as predictions from each model of change(coloured bars) for average data from a small patch of white matter. For each parameter, the best amount of change given the baseline, the observed change and the noise covariance is estimated using equation 3.3. In other words, the bars indicate the closest change in the measurements that can be produced when only that parameter has changed.

This plot suggest that the observed change in WMH is an increase in the *b0_mean* and *b1_mean* as well as an increase in anisotropy for the b1 shell. This *pattern of change* is better aligned with a positive change in s_{ex} than in any other parameter.

Figure 3.11b shows the estimated probability of change $P(\hat{\Delta\nu} \mid \Delta y)$ for each parameter of the standard model for an axial slice of the brain in voxels that included more than 10 WMH samples(subjects). These probabilities are normalized to sum up to 1 for each voxel. The colours indicate the probability that a change in the corresponding parameter can explain the observed changes in WMHs.

Figure 3.12a shows the best explaining model of change in each voxel in a few axial slices of the brain. To check for the reproducibility of the results, we have divided subjects in two batches of equal size (1500 each) and repeated the whole pipeline. The inferred changes were highly similar in the two batches with average error of 0.4% in the estimated probability of change.

In more than 65% of the voxels, that are mostly in deep white matter, the best model is a change in s_{ex}. However, in voxels adjacent to the ventricles, all other models compete and there is not a dominantly winning model. This might be due to a true difference in microstructure in these periventricular voxels, or may be caused by high variability across subjects due to CSF partial volume effects.

Figure 3.12b shows the estimated amount of change in s_{ex} in voxels where this was the most probable parameter. In most of the voxels an increase in s_{ex} between 0 and 0.4 explains the observed change in WMH. The bottom right panel shows that the amount of change increases with distance from the ventricles, whereas in deep white matter the average amount of change remains relatively constant.

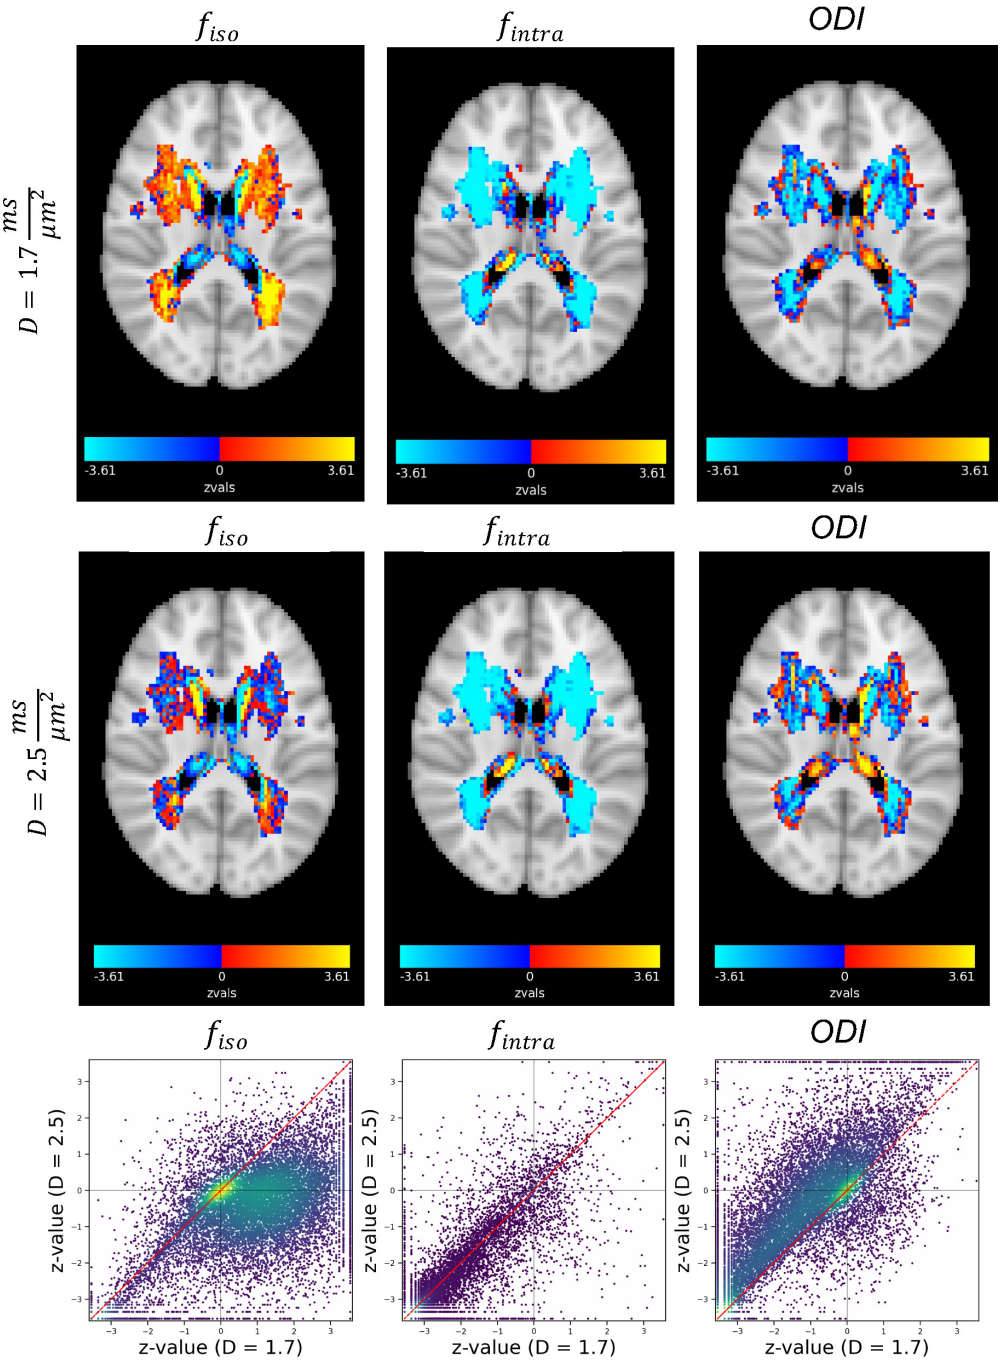
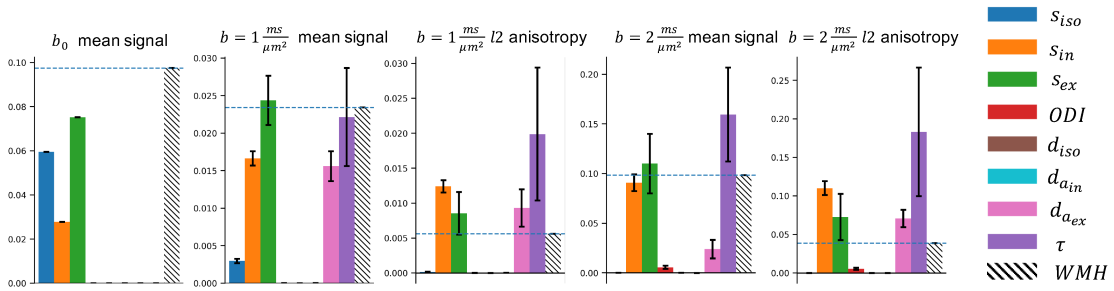
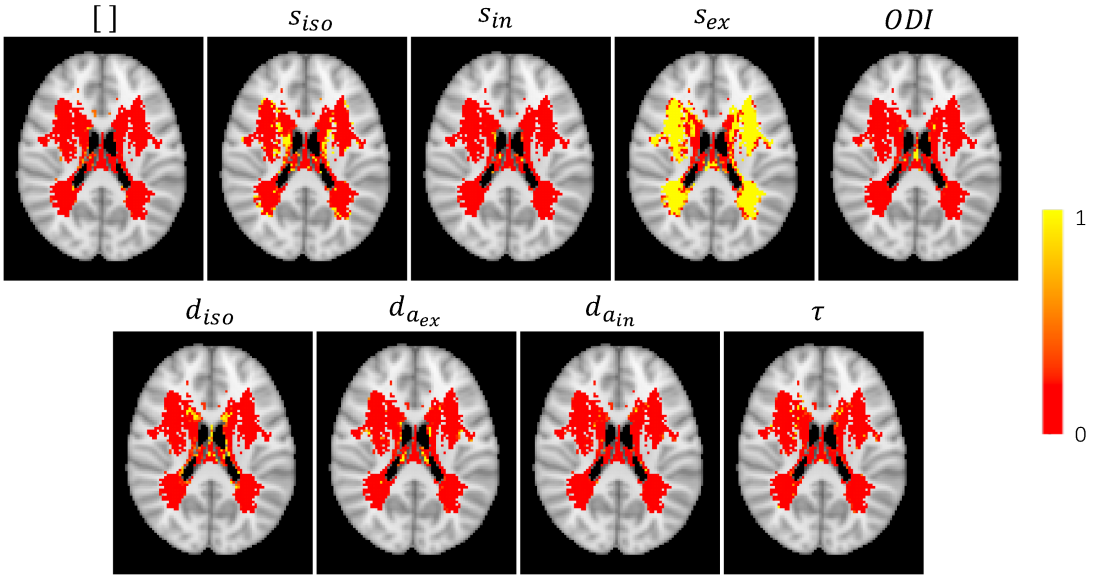


Figure 3.10: NODDI parameter estimates. Top) z-maps for the difference between WMH and normally appearing white matter with the assumption $d_{intra} = d_{extra} = D = 1.7\frac{\mu m^2}{ms}$. **Middle)** The same maps with the assumption $d_{intra} = d_{extra} = D = 2.5\frac{\mu m^2}{ms}$. **Bottom)** Scatter plot of the z-values for the two cases. The results show the assumed fixed value for the diffusivity significantly affects the estimated change between WMH and normal tissue for f_{iso} and ODI. However, the observed decrease in f_{intra} is fairly robust to the difference in diffusivities, and this is inline with the results from BENCH.

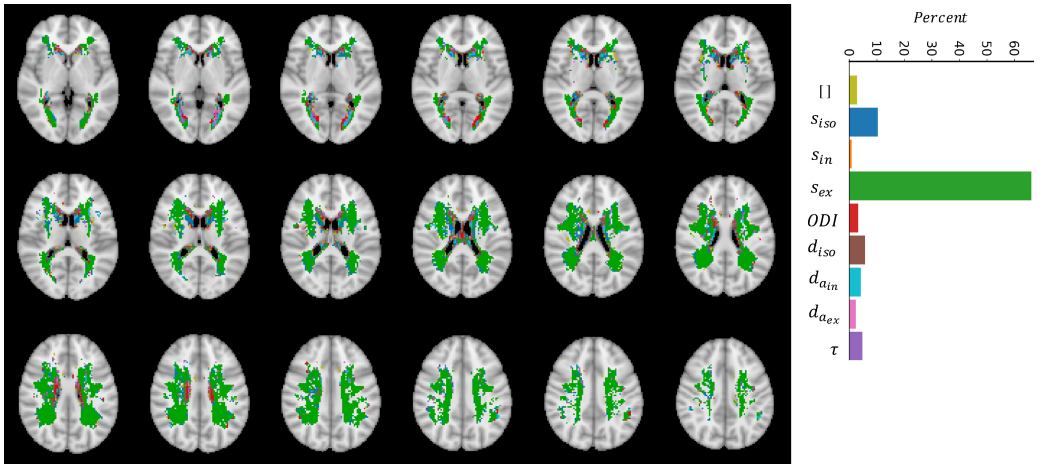


(a) The observed change in WMH and predicted change vectors.

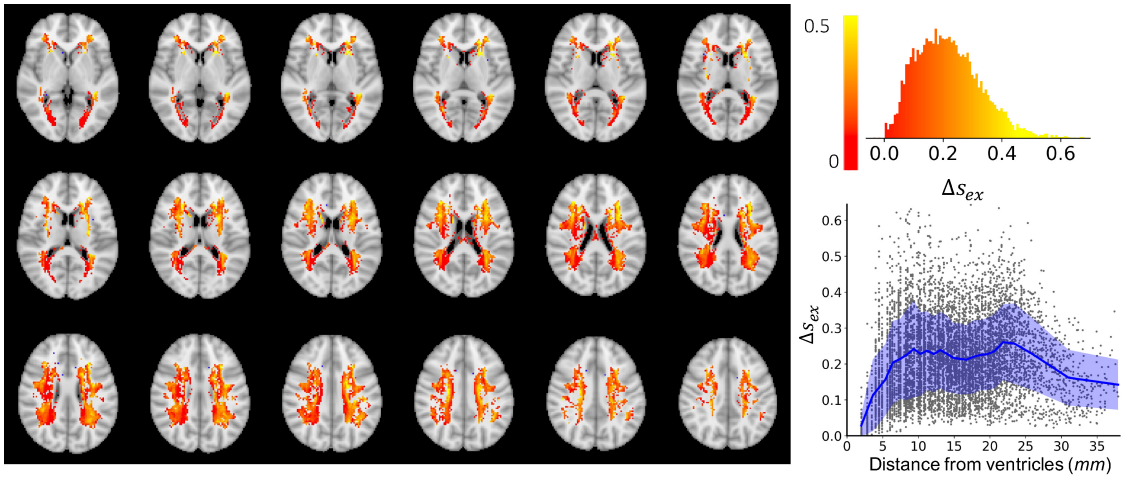


(b) Estimated probability of change in each parameter of the standard model $P(\hat{\Delta\nu} \mid \Delta y)$.

Figure 3.11: a) Each panel shows the estimated amount of change in the measurements if only the corresponding parameter changes, along with the actual observed change in hyperintensities for a patch of voxels in white matter. Each bar is scaled with the best estimated amount of change for that parameter. The observed change in WMH is an increase in the *mean-b0* and, to a lesser extent, and increase in *mean-b1*, and a positive change in the l2 measurements. This is best aligned with the *pattern of change* that an increase in s_{ex} can produce. **b)** Each map shows the estimated probability that change in the corresponding parameter can explain the observed change in the summary measurements between WMH and NAWM at a single axial slice of the brain. The no change model represents the null hypothesis that the change is better explained by noise rather than a change in any one of these parameters. In the majority of the voxels, the change model for s_{ex} has a probability around 1 (yellow) and the remaining parameters are nearly zero(red). This means that a change in s_{ex} is more likely to explain the observed change than any other single parameter change.



(a) Best explaining model of change in each voxel.



(b) Estimated amount of change in s_{ex}.

Figure 3.12: a) The colours indicate which model of change could best explain, i.e. had the highest posterior probability given the observed change in the summary measurements between WMH and NAWM. In the majority of voxels (65%) a change in s_{ex} explained the data better than any other model. However, in the regions very close to the ventricles there is no major winning model. This can be either because of high between subject variability or a different type of change that is not captured by the trained models of change. **b)** The maps show the estimated amount of change in s_{ex} in voxels where s_{ex} was the best model using a maximum a posteriori estimation $\Delta s_{ex} = \underset{\Delta\nu}{\text{argmax}} P(\Delta\nu \mid \Delta y, \Delta\hat{s}_{ex})$. At most of the voxels the estimated amount of change is positive, meaning that an increase in s_{ex} can explain the change in the summary measurements observed in the WMH voxels. The top right panel shows the distribution of estimated amount of change at the voxels where change in s_{ex} was the best model. Most of the estimated changes are between 0 and 0.4. The bottom right panel shows the amount of change vs the distance (in millimeters) from the ventricles.

3.6 Discussion

We presented a Bayesian framework to directly infer changes in parameters of a biophysical model from observed changes in a set of measurements. We applied the method to microstructural modelling of diffusion MRI, where biophysical models usually require many free parameters and are often degenerate.

3.6.1 Comparison with model inversion

The traditional approach to overcome these degeneracies is to constrain some of the parameters to biologically plausible values so that other parameters can be estimated using a conventional measurement (e.g., fixing the diffusivities in NODDI, [10]). Such assumptions reduce the full model parameter space to a restricted subspace, where the model is invertible. This direct inversion approach has the advantage that it gives parameter estimates and that it can model any parameter change in this restricted subspace. However, violation of these assumptions can significantly bias the parameter estimates.

Our proposed approach allows the initial set of parameters to lie anywhere within the full model parameter space (restricted only by broad user-defined priors); and any of these parameters might change. This extra flexibility comes at the price that the parameter changes are assumed to lie along 1D lines in parameter space defined by the user-provided patterns of change $\hat{\Delta\nu}$. For each of these hypothesized 1D change models, we estimate the posterior probability of such a change as well as the most likely amount given the baseline data and the change in it.

To compare this assumption with that made by direct inversion, let us consider a biophysical model with 8 free parameters. Let us further assume that, due to the limited degrees of freedom in our model, we can only fit 3 out of these 8 parameters. In this case direct inversion would require assuming that the microstructural change is limited to a subset of three parameters, i.e., a 3-dimensional subspace of the full 8-dimensional parameter space. In contrast, BENCH assumes by default that the change is caused by one out of the 8 parameters, which corresponds to the microstructural change lying in one of 8 one-dimensional lines in parameter space.

This suggests that if one has prior knowledge of which microstructural parameters are likely to change, it might make sense to use direct inversion with those parameters as free parameters. BENCH would have the advantage in a more exploratory approach, where any of the underlying parameters might have changed. However note that this comparison between approaches is complicated by the fact that using model inversion requires setting a subset of the parameters to some fixed value, which might cause a bias in the free parameters if inaccurately fixed [9, 14].

It is important to note that the user-defined prior distributions for parameters do not directly imply a prior value for the parameters. These priors are used to train the regression models and are required to be wide enough to capture all possible underlying parameter settings. Nevertheless, using broader priors only requires more complex machine learning models that can capture the variation in the relation between the measurements and their derivatives.

In the proposed approach we train the models with simulated data once (without requiring any real data) and use the trained models to estimate the desired probabilities for any real data with the same acquisition protocol. This precomputation saves one from having to integrate over all possible initial parameters when inferring the parameter change in each voxel. Therefore, the inference on real data which only consist of a few 1d integrations for each voxel, runs much faster than the non-linear optimizations in alternative inversion approaches.

The results from simulations suggest that we are able to identify changes in signal fraction accurately for the given brain-like measurement. However, there is a considerable confusion in the diffusivities, meaning that the change in these parameters is not distinguishable from one another. In simulations, we have only accounted for measurement noise, but in real data, particularly in cross-sectional studies, between-subject variability also contributes to noise. Hence, the reported performances and sensitivity to changes in parameters in the simulations section are more reliable when the between-subject variability is less important, for instance, in longitudinal studies. These accuracy values depend on the baseline measurements, underlying parameters, and the nature of how each parameter affects

the measurements. Nevertheless, an important point is even in the case of full confusion in diffusivities, the results from the proposed approach is more reliable compared to the model inversion with fixed parameters. That is because a wrong prior for the fixed parameters can bias the estimates for other parameters, while in the proposed approach we avoid such assumptions. For example, in NODDI any changes in the B0 signal are usually ignored (as a result of the sum constraint on the signal fractions), but in our approach we allow changes in the b0 signal to inform which microstructural parameter might have changed.

In this study we showed that setting a different value for diffusivities in NODDI can result in contradictory inference about changed parameters in white matter hyper intensities. The only consistent change was a decrease in the ratio of intra and extra axonal signal fractions which is in line with the results of BENCH (an increase in s_{ex} with no change in s_{in}). This analysis thus illustrates one of the main benefits of using BENCH: the results do not depend on some prespecified value of a parameter as we integrate over all possible values for the parameters rather than fixing them. Another advantage is that BENCH can provide a more specific explanation for the change, e.g. in this case as opposed to NODDI that only identifies a change in the ratio of the signal fractions, BENCH can specifically tell if it is a change in the extra axonal signal fraction.

The fact that the approach doesn't require the models to be invertible makes it applicable to studying changes in over-parameterised models or models without closed form analytical solution, e.g. simulation-based models. Such simulation-based models provide the opportunity to explore more complex and realistic models of diffusion in a tissue. There is no limitation in the number of parameters as long as they affect the observed data in some way. If several parameters cause the data to change in the same (or very similar way), this approach will give a list of possible parameters underlying the observed change with a probability associated with each. The resulting probability estimates can be used to eliminate unlikely change scenarios.

We utilized the trained models of change for the parameters of the "standard" model for diffusion to investigate which microstructural changes can explain white matter hyperintensities. The results suggest that the change can be associated with an increase in the extracellular signal. This is in line with other findings using more complex diffusion encodings [66], who found an increase in the extracellular T2, which would lead to an increase in the extracellular signal contribution. Comparing with the inversion approach, here we did not assume diffusivities are fixed in various brain regions, but we assumed only one of the parameters has changed as a result of white matter hyperintensity. However, it is possible that simultaneous changes in multiple parameters can better explain the change in the data, which could be tested in the same framework with the extended models of change. For example, a model with combination of the parameters might be able to explain a positive change in *b0_mean* and a negative change in *b2_mean* as it was observed in some voxels. Furthermore, we are limited to detect any changes within the constraints of the "standard" model. Hence, any changes in the signal in the white matter hyperintensities due to phenomena not within the "standard model" (e.g., exchange or non-Gaussian diffusion) would be misinterpreted as changes in the "standard" model parameters.

3.6.2 Summary measures

The choice of summary measurements to train change models is arbitrary, but this choice can affect the performance of the model. It is essential that the summary measurements are able to capture enough information from the data such that they are sensitive to changes in the parameters of interest and insensitive to other changes that are not part of the model parameters. For example, in our simulations we did not include the fibre orientation parameters as part of the free parameters, and therefore we required the summary measures to be rotationally invariant. Hence the choice of decomposing the signals in each shell into spherical harmonics to extract rotationally invariant summary measurements. Of course one can instead use other signal representations, such as measures derived from the

diffusion tensor model, or the kurtosis tensor model, etc, to compute the summary measurements. We chose spherical harmonics over other choices as they are fast to calculate, and the bases are orthogonal which leads to summary measures that capture different aspects of the data.

3.6.3 Future developments

While in the examples shown here these patterns of change only altered a single parameter at a time, in the current framework the *pattern of change* can be any vector in parameter space. In the future we plan to extend this framework to allow for parameter changes in 2D or 3D hyperplanes rather than just along 1D lines (see Appendix 3.8 for the feasibility of this extension). However, the dimensionality of these hyperplanes will always be lower than that of the restricted parameter subspace in which parameters can freely change with the direct inversion approach. Note that computing posterior probabilities in a full Bayesian framework allows for comparison between models of change with different complexities without the need for arbitrary regularisation.

In addition, the model of change can be extended to study continuous changes (e.g. ageing), as opposed to discrete group differences as shown in this work. To do so, one first needs to compute the gradient of change in the measurements with respect to the independent variable, e.g. time, using a regression model. Then one can use the chain rule to relate the rate of change in the measurements to the rate of change in the parameters. Such an approach makes modelling continuous change a straightforward extension of this framework.

Although here we mostly show how our method can be applied to detect changes in parameters given the data, our framework can also be used to optimize data acquisition protocols for detecting changes in particular parameters of interest. For example, in the simulations we show that it is difficult to detect a change in the free-diffusion parameter. Our framework can be used to extend the acquisition (e.g. by adding lower b-values) and, using the output confusion matrices, establish an optimal set of b-shells to enable detection of change in free diffusion.

Finally, while we applied the framework to the specific problem of studying microstructural changes using diffusion MRI in the brain, the framework is general meaning that it can be applied in any field where biophysical models are available. For example, the same approach as described in this study can be applied to dynamical causal models (DCM) [85] for fMRI or MEG/EEG. These are notoriously over-parameterised, but often, are applied in a context where the values of the inferred parameters is of lesser interest than the change in the parameters under different experimental conditions, and its reasonable to assume the change is sparse; the ideal scenario for BENCH.

3.7 Software

BENCH is an open source software implemented in python and available at `https://git.fmrib.ox.ac.uk/hossein/bench`.

3.8 Appendix A. Toy example: Inferring changes in 2D

Consider the forward model

$$f(x) = ax^3 + bx^2 + cx + d \tag{3.23}$$

The model has 4 free parameters (a, b, c, d). Given 3 measurements this model is degenerate, i.e., one cannot estimate all the parameters uniquely. Now consider two instances of this model with parameters (a_1, b_1, c_1, d_1) and (a_2, b_2, c_2, d_2) with 3 measurements for each. Obviously, this system is degenerate and parameter estimation is ill posed. However, if we are only interested in comparing two model instances, we can still infer changes by assuming that the change is sparse. This is the premise of BENCH.

Now we will demonstrate that despite the model degeneracy, we can not only detect changes in a single parameter, but also infer simultaneous changes in pairs of parameters. Consider $(a_1 = 1, b_1 = 1, c_1 = 1, d_1 = 1)$ and $(a_2 = 1.2, b_2 = 0.8, c_2 = 1, d_2 = 1)$, i.e., $\Delta a = +0.2, \Delta b = -0.2, \Delta c = \Delta d = 0$.

When using Monte Carlo simulations to infer parameters for each model given three independent measurements, the posterior distribution is clearly degenerate as shown in Figure 3.13a. In this figure, the blue (resp. red) distribution shows the parameter estimates for (a, b) for the first (resp. second) data set. The intensity of each point encodes the log posterior probability for the estimated parameters. The stars show the true parameter values. The plot demonstrates that parameter estimates are highly correlated (i.e. the model is degenerate).

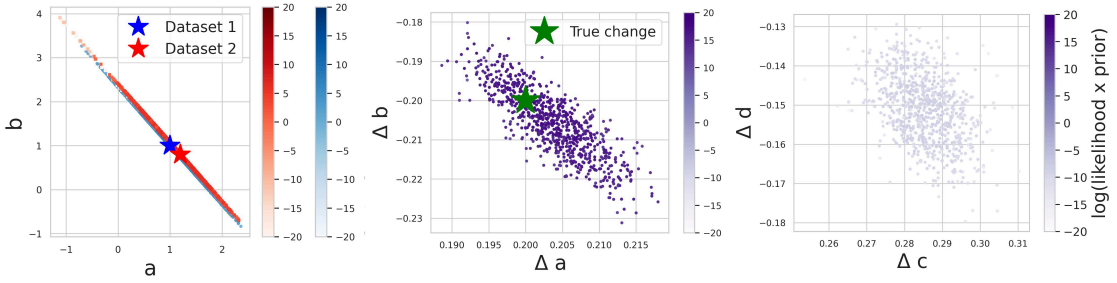
In contrast, figure 3.13b shows Monte Carlo samples for Δa and Δb for the change model. The plot demonstrates that the estimated parameter changes are distributed around the true change value and each sample has a comparatively high posterior probability value. It is therefore possible to infer the true, 2-dimensional change.

We also considered an alternative change model where a and b are fixed and c and d can change. The estimated samples for Δc and Δd are shown in Figure 3.13c. In this case the estimated samples have much lower posterior probabilities (lower intensities) than the a,b change model. Thus, we can use the change model to assess not only the amount of change in 2D, but also which pair of parameters best explains these changes. The changes are still sparse, but not necessarily 1-dimensional.

In BENCH we integrate the approximations of this unnormalized posterior probabilities to compute the the desired probabilities for each model of change in Equ. 3.1. Hence, it this example BENCH (once extended to allow multi-dimensional changes) would correctly infer that it was the parameters a and b that changed, and not the parameters c and d.

3.9 Appendix B. Estimating Quality of Fit

The estimated probability in Equ.3.1 tells how well each model explains the observed change compared to all other defined change models, but it doesn't necessarily reflect to what extent the observed and predicted change are matched. In other words, a model with a poor quality of fit to the data can get a high probability value because its prediction is the closest to data compared to all other models.



(a) Parameter estimation **(b)** Correct change model **(c)** Wrong change model

Figure 3.13: a) Parameter estimation. Each set of dots shows parameter estimates for one instance of the model using MCMC and intensities represent the log posterior probability. The parameter estimates for each data set are highly correlated and all of the points on the lines explain the data equally well, i.e. the models are degenerate and it is not possible to directly compare the parameter estimates. B) Inferred change with the correct model. We ran MCMC with the assumption that change has a particular shape (only a and b changed). The estimated values for Δa and Δb are centred around the correct change (green star) and the unnormalized posterior probabilities are comparatively high. C) Inferred change using a wrong model. We run a similar MCMC but this time assuming c and d can change. In this case the estimated posterior probabilities are much smaller compared to the previous change model, i.e. this model of change cannot explain the change in the measurements as well as the model in (b).

Also, it is possible that more than one change model predict the data accurately and hence all get low probabilities in Equ.3.1.

To estimate how well a change model can explain the change in data one can look at the chi-squared distance between the predictions of the change model and the measured change:

$$d = (\Delta y - \mu)^T \Sigma^{-1} (\Delta y - \mu) \tag{3.24}$$

In the above expression, Δy is the observed change in the data, and μ and Σ are the mean and covariance of change in the measurements predicted by the best model. This statistic follows a chi-squared distribution and a higher d means more discrepancy between the observed change and the predicted change.

Figure 3.14 shows the distribution of d for the case of one parameter change that is explained by the correct model (blue) and the case of two parameter change that is mistakenly identified as a single parameter change (orange). Accordingly, our recommendation when the discrepancy is high is to consider revising the change models, as the winning model is poorly explaining the observed change. For

example, one can define biologically feasible linear combinations of the parameters as change directions.

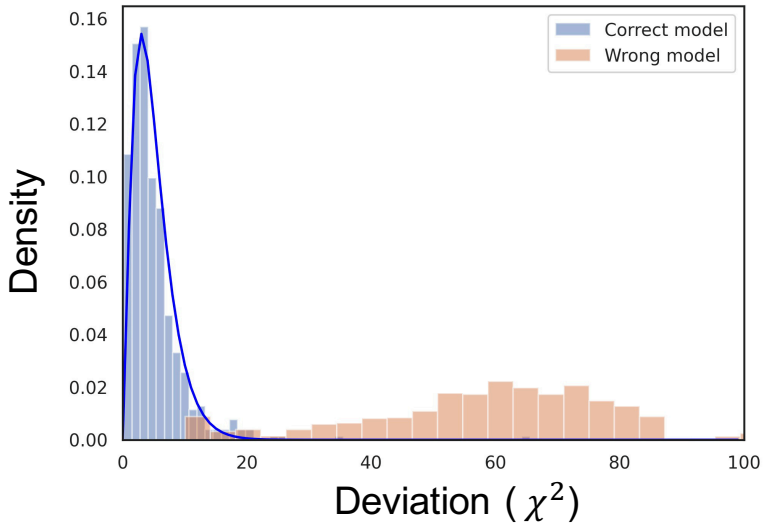


Figure 3.14: Distribution of distance for the correct change model (blue) and a wrong model(orange). Given a baseline measurement (y) and a change (Δy), we estimate the most likely change in the parameters as well as the most likely amount of change in that direction using our trained change models. These estimates can then be used to predict the distribution of expected change in the measurements. Using the discrepancy between this prediction and the actual observed change, we can determine the quality of the change model in explaining the data. The histograms are showing the Mahalanobis distance (i.e., the offset normalised by the covariance matrix as defined in 3.24) between the actual and the predicted change in the measurement when the correct change model is used (blue) and when a wrong change model is used (orange) for several instances of simulated data. The blue curve shows the pdf of χ^2 distribution with $df =$ the number of measurements.

"Two heads are better than one."

—John Heywood

4

Hierarchical Modelling of Crossing Fibres in the White Matter

Contents

4.1 Abstract 74
4.2 Introduction 74
4.3 Theory 76
4.3.1 Crossing fibre model 76
4.3.2 Hierarchical model 78
4.3.3 Parameter Inference 80
4.3.3.1 Determining the number of fibres in each voxel 82
4.3.4 Fixel based threshold free cluster enhancement 82
4.4 Methods 85
4.4.1 Simulation settings 85
4.4.2 Data sampling and voxel correspondence 85
4.4.3 Fixel template creation 86
4.4.4 GLM analysis 87
4.5 Results 87
4.5.1 Simulations 87
4.5.2 White matter fixel template 88
4.5.3 Fixel based GLM 93
4.6 Software implementation 98
4.7 Discussion 98
4.7.1 Fixel template creation and usage 99
4.7.2 Future Directions 101

4.1 Abstract

While diffusion MRI is typically used to estimate microstructural properties of tissue in volumetric elements (voxels), more specificity can be obtained by separately modelling the properties of individual fibre populations within a voxel. In the context of cross-subjects modelling, these so-called fixel-based analyses require identifying equivalent fibre populations. This is usually done post-hoc, after estimating fibre orientations for individual subjects independently and subsequently matching the fixels between subjects. This approach can fail due to individual differences in fibre orientation distributions.

Here, we introduce a hierarchical framework for fitting crossing fibre models to diffusion MRI data in a population of subjects. This hierarchical setup guarantees that the crossing fibres are consistent by construction and, therefore, comparable across subjects. We propose an expectation-maximisation approach to fit the model, which can scale to large numbers of subjects. This approach produces a crossing-fibre white matter fibre template, which can be used to estimate fibre-specific parameters that are consistent across subjects and, hence, can be used in fixel-based statistical analyses.

4.2 Introduction

Diffusion MRI is widely used to study between-subjects variations in white matter, both in healthy cohorts and in patient populations [7, 8, 36]. Like in voxel-based morphometry for studying grey matter [86], comparing white matter between subjects requires care in the between-subjects spatial alignment. This is particularly crucial when the objective of the analysis is the comparison of microstructural effects (from diffusion MRI models), rather than macroscopic variations in gross anatomy.[53]

Recently, fixel-based analysis of white matter has emerged as a technique for cross-subject analysis of white matter not only within each voxel, but also within each fibre population in the voxel[60, 87] In fixel-based analysis, a white matter orientation model is used to assign properties to fibre populations crossing within

each voxel (e.g., the density of fibres aligned with a given direction). These metrics are then compared across subjects, allowing investigation of fibre-specific changes. This approach is in contrast to more traditional analyses that use voxelwise metrics, such as fractional anisotropy (FA) or mean diffusivity (MD), as markers of white matter. While these voxelwise metrics have been shown to be highly sensitive to changes in white matter tissue (e.g.[88]), they cannot easily be interpreted in terms of specific changes to the underlying fibre populations within the voxel (see Figure 4.1).

Fixel-based approaches attempt to resolve this fibre-specificity issue. By assigning measures to a discrete set of fibre orientations within a voxel, between-subjects differences can then be associated with specific fibre orientations, which in turn may correspond to fibre pathways associated with specific functions. However, not only do fixel-based analyses need to resolve between-subject spatial alignment (just like voxelwise methods do), they also need to resolve between-subject assignment of fibres. In order to "compare like with like", the discrete set of fibres within a voxel must all be in correspondence across subjects [89].

So far, this between-subject fibre assignment has been conducted as a post-processing step, following independent processing of each subject separately. In [89], fibre orientations are estimated on a subject by subject basis using the ball and sticks model [90]. Then, to maximise correspondence both spatially and between subjects, fibre labels are swapped in multiple stages. The method was restricted to the white matter skeleton [53] in order to ensure validity of the spatial assignment. More recently, fixel-based analyses were conducted using continuous fibre ODF models [59, 60] fitted to each subject separately. Peaks from the fODF were subsequently matched and relabelled across subjects. This approach is, in fact, what is currently referred to as fixel-based analysis.

However, even after accurate spatial normalisation with fibre reorientation, post-hoc between-subjects fibre reassignment can often fail due to variability in white matter orientation estimates. This can either be natural variability in anatomy between subjects, or variability due to noise in the parameter estimation process. In either case, fibre reassignment can be ill-posed.

Here, we propose a solution to this problem that takes part earlier in the processing pipeline, at the stage of voxelwise white matter orientation modelling, rendering post-hoc fibre reassignment unnecessary. We build a cross-subject hierarchical model of crossing fibres where, in each voxel, the parameters of the model (including fibre orientations) are drawn from a population distribution. This means that the assignment of the fibre orientations is matched across subjects by construction, rather than through a post-processing step. We fit both the individual subjects and the group parameters simultaneously to all the data and obtain group and individual estimates of orientations and associated fixel parameters. After a template has been created, it can be used to compare signal fractions in different fibres across a larger sample of individuals by fitting the ball-and-sticks model with the template acting as the prior distribution.

In addition, we propose an expectation-maximisation algorithm that enables us to fit the model in a way that is scalable to a large cohort of subjects. We also extend the voxel-based Threshold-Free Cluster Enhancement (TFCE) approach [91] to enable statistical analysis of fixels with proper correction for multiple comparisons across space. The method and implementation are validated through simulations. Finally, we create a group-averaged fixel template for white matter using data from the UK Biobank [92], and we used this template in an example analysis of fixel changes with ageing.

4.3 Theory

4.3.1 Crossing fibre model

While our framework can be adapted to work on different crossing fibre models, we choose the ball-and-sticks model because it is simple, widely used, and it lends itself naturally to a hierarchical extension. According to this model, the diffusion signal in a voxel originates from the combination of a ball compartment that represents water with isotropic diffusion and one or more stick compartments that represent

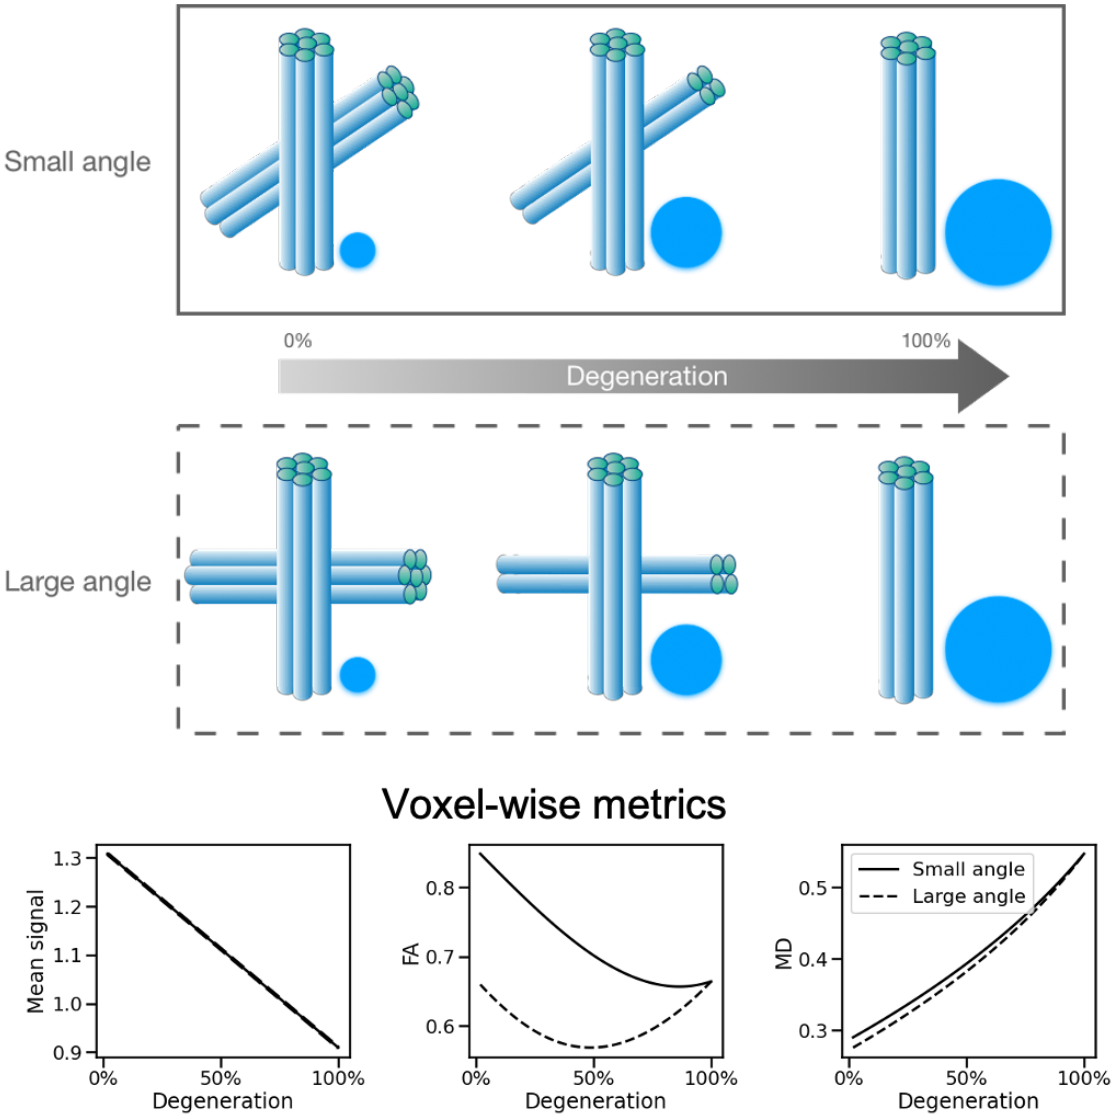


Figure 4.1: Illustration of the effect of a fibre-specific biological process on voxel-wise measurements. The top panels show a scenario of axonal degeneration in a crossing fibre region, where one of the fibre populations (shown by cylinders) gradually degenerates and is replaced by water with isotropic diffusion (shown by circles). Two possible scenarios for the same process are depicted in the two panels, one in which the angle between the two fibre populations is small, and another in which the angle is large. The bottom row shows the change in voxel-wise measurements of mean signal, fractional anisotropy, and mean diffusivity in both scenarios (solid line for the small angle and dashed line for the large angle). The angle between the fibres has no effect on the mean signal and a small effect on MD, but it substantially affects the changes seen in FA.

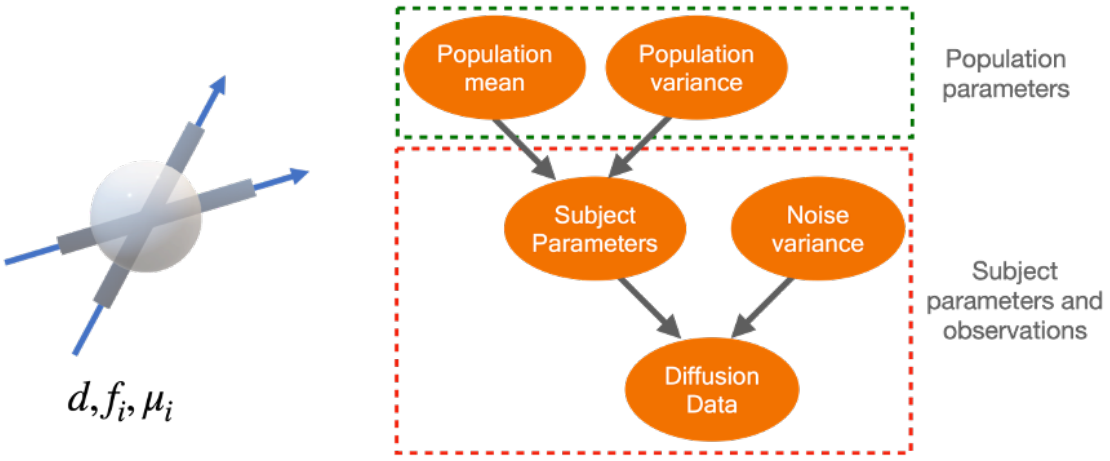


Figure 4.2: The generative forward model of diffusion MRI signals for the whole population. Left: A schematic representation of the ball and sticks model for diffusion MRI signal in white matter voxels. According to this model, water only diffuses axially inside each population of fibres and isotropically outside the fibres. Right: The hierarchical graphical model for the diffusion MRI signal in a white matter voxel for a population of subjects. In this model, we assume that, within a single subject, each parameter of the ball and stick model is drawn from a population distribution. The red box shows subject-specific parameters and observations, and the green box shows population-level parameters. This hierarchical structure will introduce a dependency between the subjects, which has the advantage of ensuring that the fibre labels are consistent across subjects, therefore allowing for the comparison of fibre-specific parameters between subjects.

water that diffuses along a fibre population. The diffusion signal according to the ball and sticks model is generated by

$$S = S_0[\sum_i f_i e^{-bd(\mu_i^T g)^2} + (1 - \sum_i f_i)e^{-bd}] \tag{4.1}$$

where b, g are the b-value and the gradient orientation, d is the diffusion coefficient, μ_i, f_i are the orientation and signal fraction of the i'th stick and S_0 is the baseline signal with no diffusion encoding. Hence, the set of free parameters in this model are $\{S_0, d, f_i, \mu_i\}$ per subject per voxel.

4.3.2 Hierarchical model

In a hierarchical framework, we assume that, in each voxel, each parameter in each subject (ν_s) represents a sample drawn from a population distribution for that parameter. This hierarchical structure ensures that all parameters, including the labelling of the fibres, are consistent across subjects.

Table 4.1: Description of notations

Notation	Description	Dimension
M	Generative biophysical model	$R^p \rightarrow R^d$
D_s	Data for subject s *	d
ν_s	Parameters of the model for subject s	p
σ_s^2	Noise variance for subject s	d
ν_g	Mean of the population parameters **	p
σ_g^2	Variance of the population parameters	p

* subscript s denotes subject specific parameters.
** subscript g denotes group parameters.

The population distributions of all scalar parameters are assumed to be Gaussian with unknown means and variances, while orientation parameters are assumed to follow a Watson distribution with unknown mean and concentration. Finally, each subject's diffusion MRI signal is derived from the subject parameters following the forward model (M) in 4.1 with additive noise ($D_s = M(\nu_s) + \epsilon$). The noise distribution is assumed to be a zero-centred Gaussian with a variance (σ_s^2). This variance also has an exponential prior distribution with hyperparameter α_n. The diagram in Figure 4.2 shows the full generative forward model. Using the dependency structure that is shown in the graph, the joint distribution of all free parameters factors into:

$$P(\{\nu_s\}, \{\sigma_s^2\}, \nu_g, \sigma_g^2, \{D_s\}) = \prod_s \left[P(D_s|\nu_s, \sigma_s^2)P(\sigma_s^2)\right] \times \prod_s P(\nu_s|\nu_g, \sigma_g^2) \times P(\nu_g)P(\sigma_g^2) \tag{4.2}$$

The curly brackets indicate sets of subject-specific parameters. The conditional probabilities in the graphical model are defined as:

$$\begin{array}{lcll} P(D_s|\nu_s, \sigma_s^2) & = & Normal(M(\nu_s), \sigma_s^2 I) & \\ P(\nu_s|\nu_g, \sigma_g^2) & = & Normal(\nu_g, \sigma_g^2) & \text{for scalar } \nu_s \\ P(\nu_s|\nu_g, \sigma_g^2) & = & Watson(\nu_g, \sigma_g^2) & \text{for directional } \nu_s \end{array} \tag{4.3}$$

For the Watson distribution, ν_g is the mean orientation, and σ_g^2 is the orientation dispersion index [10]. We assume uniform priors within plausible ranges for all the group mean parameters (ν_g) and a power-law prior for the group variances (σ_g^2) with hyperparameter. The power law prior is used to prevent the group variance

from going to zero. The prior distributions of the parameters are given by

$$\begin{array}{lcl} P(\nu_g) & \propto & 1 \\ P(\sigma_g^2) & \propto & e^{\alpha_g \sigma_g^2} \\ P(\sigma_s^2) & \propto & e^{-\alpha_n \sigma_s^2} \end{array} \tag{4.4}$$

α_g and α_n(positive real numbers) are hyperparameters that are adjusted by the user. Setting a high value for α_g implies a stronger constraint for group variances to converge to zero, thereby favouring individual differences. On the other hand, a high α_n increases the penalty associated with estimating a high noise variance.

4.3.3 Parameter Inference

We employ a maximum a-posteriori approach to fit the parameters of the hierarchical model to the data, that is to maximise the conditional probability of all the parameters of interest given the data ($P(\nu_g, \sigma_g^2, \{\nu_s\}, \{\sigma_s^2\}|\{D_s\})$). The free parameters of this model are the group means and variances for each model parameter, subject parameters and subject specific noise variance. The full posterior probability of the model according to the hierarchical structure factorises into:

$$P(\{\nu_s\}, \{\sigma_s^2\}, \nu_g, \sigma_g^2 \Big| \{D_s\}) \propto \prod_s P\left(D_s | M\left(\nu_s\right), \sigma_s^2\right) P(\nu_s|\nu_g, \sigma_g^2) \times P(\nu_g) P(\sigma_g^2) \tag{4.5}$$

Fitting this entire model requires simultaneously optimising all of these parameters, which creates a space with a very high number of dimensions that grows linearly with the number of subjects. This makes optimization infeasible for datasets with more than a few tens of subjects.

To make the optimisation feasible for large datasets, we employ an expectation-maximisation approach: In the expectation step, the group parameters are estimated while the subject parameters are fixed; in the maximisation step, the group parameters are held fixed while the subject parameters are estimated. These stages are repeated until all of the parameters have converged. Convergence is achieved when all group and subject parameters remain unchanged between subsequent iterations.

Expectation

In the expectation step, the subject parameters are set to the best guess made so far so that the population parameters can be estimated. Following Equ. 4.5 this leads to optimising the following conditional posterior:

$$P(\nu_g, \sigma_g^2 | \{\nu_s\}, \{\sigma_s^2\}, \{D_s\}) \propto \prod_s P(\nu_s \Big| \nu_g, \sigma_g^2) \times P(\nu_g) P(\sigma_g^2) \tag{4.6}$$

Since we assume that there is no correlation between parameters a-priori, each parameter is evaluated on its own at this stage. For the scalar parameters, plugging the probabilities according to Equ. 4.6 and optimising Equ. 4.6 results in:

$$\nu_g = \frac{1}{n} \sum_s \nu_s, \sigma_g^2 = \frac{\frac{1}{2} \sum_s (\nu_g - \nu_s)^2}{\alpha_g + 1 + \frac{n}{2}} \tag{4.7}$$

Numerical optimization is used for the maximisation of the posterior to estimate group orientation parameters, including average orientation and dispersion for each fibre population. For all the model fittings (both for individual fits and in the hierarchical model) we have used non-linear model fitting as implemented in the SciPy package Nelder-Mead method. [93, 94]

Maximisation

In this stage, the population mean and variance parameters are fixed, and the subject parameters are estimated by maximising the conditional posterior probability. According to Equation 4.5:

$$P(\{\nu_s\} \Big| \nu_g, \sigma_g^2, \{D_s\}) = \prod_s P(D_s | \nu_s, \sigma_s^2) P(\nu_s | \nu_g, \sigma_g^2) P(\sigma_s^2) \tag{4.8}$$

Since the group parameters are fixed in this stage, the likelihood is independent for each subject and factorises into subject specific parts. This makes it possible to optimise each subject's parameters separately. We therefore independently estimate subject parameters ν_s by optimising the logarithm of the posterior function through numerical optimisation.

4.3.3.1 Determining the number of fibres in each voxel

To determine the number of fibres in each voxel, we first fit the hierarchical model using one, two, and three fibres. We then set a threshold on the difference in the quality of the fit achieved by adding a fibre population. The log-likelihood of the data $\log\left(P(D_s|\nu_s, \sigma_s^2)\right)$ summed over all subjects is used as the quality of fit metric.

Figure 4.3 shows the full pipeline for fitting the hierarchical model, including the creation of a white matter fixel template representing the group average orientations and population volume fractions in the model.

4.3.4 Fixel based threshold free cluster enhancement

The natural next step following fitting fixels is to fit a cross-subject model of the parameters, e.g., the fibre signal fractions, using a general linear model (GLM). As this is done in each voxel, statistical techniques for controlling false positives are required due to the large number of tests involved. One popular technique is Threshold-Free Cluster Enhancement (TFCE) [91], which is a method for multiple comparison correction developed for voxelwise cross-subject analyses. TFCE works in an analogous way to cluster-based thresholding methods, but without requiring a cluster-forming threshold. Here, we show how we can extend TFCE to the context of fixel-based analyses.

TFCE works by forming contiguous clusters of voxels that are above a given statistical threshold, then calculating a statistic that combines the cluster height and extent, and finally integrating this statistic across different thresholds. To apply this method to fixel data, we need to modify the definition of spatial neighbourhood used for forming clusters.

In our definition, two fixels are regarded as neighbours if they are located in adjacent voxels (27-neighbourhood) and if the angle between them is below a predefined threshold (see Figure 4.4). This notion of adjacency has three advantages: it is symmetric; it maintains adjacency between fixels that are potentially in the same tract; and, finally, it prevents information from being mixed between crossing fibres within the same voxel. We construct a graph of connections between fixels,

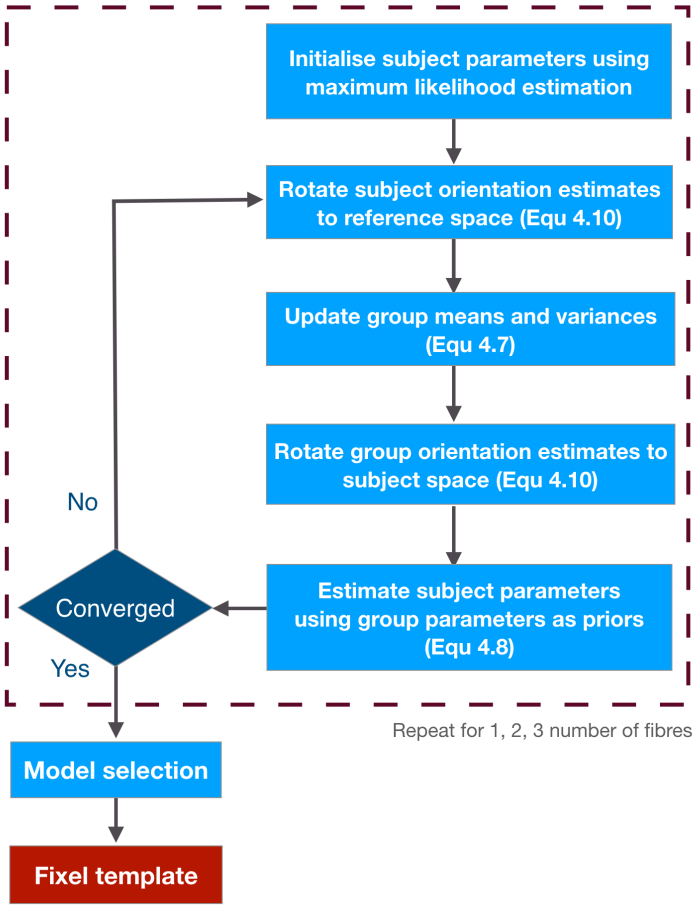


Figure 4.3: Template creation pipeline. Firstly, the diffusion data is fitted with models for each subject independently using maximum likelihood estimation. We then use the parameters from all subjects to estimate the population distribution for each parameter. We then rotate the orientations back into the space of each subject. The population parameters (mean and variance) are then utilised as prior distributions to re-fit the subject parameter to the diffusion data. This process is repeated until convergence, and for models with 1, 2, and 3 fibres separately. To then determine the number of fixels in a given voxel, we evaluate the improvement in the likelihood function across different models for that voxel. This gives us a fixel template that we can use as a prior distribution to fit fixels to each subject. The inputs of the pipeline are diffusion data for each subject that is resampled into a structural template space, and transformations from each subject's diffusion data to that template. Note that all the computations happen for each voxel in the reference space and all the diffusion data is only fetched once needed, so there is no need to register the diffusion data into the template space. To accomplish this, for every voxel in the reference space we assign a matching voxel from each subject; that is the closest voxel in the subject's space after transforming the coordinates. The relevant equations are mentioned in each box.

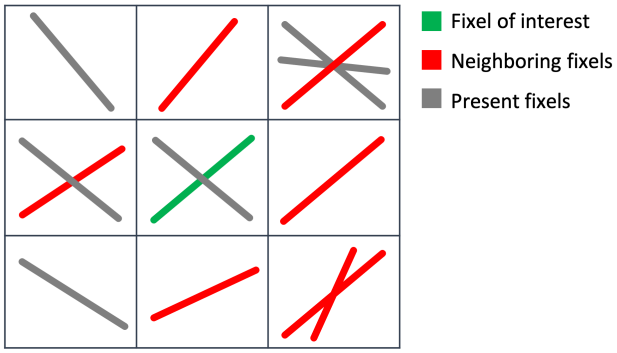


Figure 4.4: 2D section of fixel neighbourhood for TFCE. The neighbours of a fixel (green fixel) are all the fixels that are located in an adjacent voxel (at least one corner is shared, in total 27 voxels) and make angles smaller than a threshold (red fixels). Multiple fixels in a neighbouring voxel may be taken into account as neighbours of the same fixel.

where each node in the graph is a fixel, and two fixels that meet the above criteria for being neighbours share an edge.

For the fixel of interest, the cluster extent $e(h)$ at threshold h is defined as the number of nodes in the connected component containing the fixel, only considering the sub-graph that includes fixels with statistics greater than h. As a result, we assign a TFCE score to each fixel as:

$$TFCE = \int_0^t e(h)^E h^H dh \tag{4.9}$$

Here E and H are parameters that weigh the importance of cluster extent and height, respectively, t is the statistics for the fixel of interest.

Since the TFCE scores may not have a Gaussian distribution, we use standard nonparametric permutation testing to estimate p-values for each fixel. In order to achieve this, we permute the original data labelling (for example, subjects between two groups for examining group differences) and re-estimate the TFCE score multiple times to produce a null distribution. The p-value is then calculated as the likelihood of observing the actual TFCE score or higher in that null distribution.

4.4 Methods

4.4.1 Simulation settings

To determine the effectiveness of the hierarchical framework, we evaluated it with simulated data. Diffusion data for 30 participants in a single crossing fibre voxel were simulated using a ball-and-sticks model with two fibres. Each subject's parameters are randomly sampled from a population distribution. The signal fractions are sampled from Gaussian distributions with means $0.5, 0.3$ and standard deviation of $0.1, 0.1$ and then truncated to range $[0, 1]$ and normalised to sum to 0.8. Therefore, on average the first fibre population is stronger, but that may not be the case for every single subject. The stick orientations are samples from a Watson distribution with low dispersion (ODI=0.1), so that each fibre population forms a cluster across subjects. Following that, using an acquisition protocol that includes 5 b0 images and 50 diffusion-encoded images with $b = 1, 2um^2/ms$ (identical to that used by the UK biobank dataset [92], diffusion data for each subject is simulated using a custom python script. All diffusion signals are then distorted by adding Gaussian noise with SNR=50.

4.4.2 Data sampling and voxel correspondence

For the hierarchical framework to work, there needs to be spatial correspondence between subjects. This is achieved through image registration of each subject's data to a common reference space. Registration was performed using FSL's MMORF [95], a multimodal registration tool that uses both scalar structural (T1-weighted) images and diffusion tensor images (DTIs) to drive the registration. We used the Oxford-MultiModal-1 (OMM1) template [96] as our standard reference space for voxel correspondence. OMM1 is preferred over MNI152 since it was constructed using simultaneous alignment of T1-weighted, T2-FLAIR and DTI images, and therefore includes high quality, unbiased, scalar and tensor data in a common space. The inclusion of DTI data during registration ensures that matching of white matter orientation is taken into account during warp optimization. This is important for

obtaining both correct fixel assignment across subjects, and correct estimates of the group mean and variance of fixel orientation in template space.

We chose not to resample the diffusion data into a shared reference space after registration due to memory and space constraints. Instead, we built correspondence by extracting the necessary data at each stage of inference using the warps between individuals and a reference space, with nearest neighbour interpolation. To move across individual and common spaces, we are required to map the voxel position as well as the registration-induced local re-orientation and apply them to the orientation parameters. Because registration was driven by directional information within the white matter, this step also contributes to reducing cross-subject variability in fibre orientation estimates. In order to rotate all orientation parameters, we used the warpfield's local Jacobian matrix. We chose not to isolate the rotation component from the Jacobian in order to also account for non-rotational types of distortion, such as shearing. Instead, we used the complete Jacobian matrix to transform the orientation vector and then renormalized the resized orientations to unit length. For example, any given direction parameter in subject native space is resampled to the reference space using:

$$\mu_{ref} = \frac{J\mu_{nat}}{|J\mu_{nat}|} \tag{4.10}$$

Where μ_{nat} is the orientation parameter in subject native space, J is the Jacobian matrix of the warp field at that voxel, and μ_{ref} is the orientation after resampling to the reference space.

4.4.3 Fixel template creation

To create a fixel template, we used the diffusion data of 100 unrelated subjects from the HCP dataset for young adults [28]. The diffusion data includes three shells (b-values: 1, 2, 3 with 90 orientations each) and 18 b0 images (in total 288 volumes). This data has been preprocessed using the standard preprocessing pipeline in FSL for the HCP dataset [29].

Each subject was registered to OMM1 space using MMORF. Registration was driven using brain-extracted T1-weighted and DTI modalities. As fixel analysis is only valid within white matter, we created a white matter mask by running FSL's FAST [6, 97] on the T1-weighted image of the OMM1 template. The template creation process then proceeded as described above. Following the creation of this template, its parameters were fixed and it was used as a prior distribution in order to fit the diffusion model to new, unseen subjects as detailed in the following section.

4.4.4 GLM analysis

In order to illustrate an example fixel-based analysis utilising the fixel template, we ran a GLM to model changes in fibre signal fractions with age. The diffusion data from 400 randomly selected subjects (200 male, age 50+/-10) in the UK biobank were used for the GLM analysis. This dataset contains two shells (b-values= $1, 2us/(mm^2)$ each in 50 directions) and five b0 images (105 volumes in total). The standard preprocessing pipeline for UK biobank data was used to preprocess the diffusion data [98]. Once again, MMORF registration was used to estimate warps for transforming parameters between individuals' diffusion space and the OMM1 template space. In this instance T2-FLAIR data were used in addition to T1-weighted and DTI data to drive the registration.

We used fixel signal fractions as dependent variables, and age as the independent variable in our GLM analysis, with sex and head size as confounding factors. The t-statistics for the regression against age were estimated using a linear fit for each fixel independently. Then, we estimated voxel-wise p-values from permutation tests (5000 permutations) using the fixel-based TFCE scores calculated with the proposed cluster enhancement method.

4.5 Results

4.5.1 Simulations

The individual (independent) fitting approach (maximum likelihood estimates) and hierarchical model are employed to estimate individual subject parameters and

population-level distribution parameters from simulated diffusion data with known ground truth. The estimated parameters from each approach versus the actual parameters for each subject is shown in Figure 4.5.

The results indicate that the hierarchical fit is more accurate for estimating the model parameters. In particular, the individual fitting approach failed to detect the weaker secondary fibre population in the majority of subjects (most of the second fibre signal fractions are zeros). The line plots show the estimated orientation of the sticks for each subject in both approaches. The individual fit has no clear clustered pattern whereas the hierarchical fit was able to assign the same label to corresponding sticks.

4.5.2 White matter fixel template

The white matter fixel template contains the mean and variance of the ball-and-sticks model parameters for each white matter voxel, including the diffusivity, direction, and signal fraction of the fibres in that voxel. This provides a basis for comparing the signal fractions in different fibres over a broader sample of participants by fitting the ball-and-sticks model using the template as prior.

Figure 4.6 displays two sections of the template's fixel orientations and two magnified areas of crossing fibre regions for better visual inspection. These orientations are calculated using the group average across subjects. Although there are no between-voxel spatial constraints, the orientations of the fixels exhibit good spatial continuity and conform to well-established anatomical features.

Figure 4.7 compares the estimated fibre orientations for a sample white matter crossing fibre voxel with and without the hierarchical model. As a result of the subject-specific nature of the individual fit, existing fibre populations may not be oriented consistently across subjects, and some individuals may have a different number of fixels. Also, despite the fact that there is a high degree of consistency in the orientation of the fibre population between subjects, in some of them the relative strength of the fibre populations are swapped.

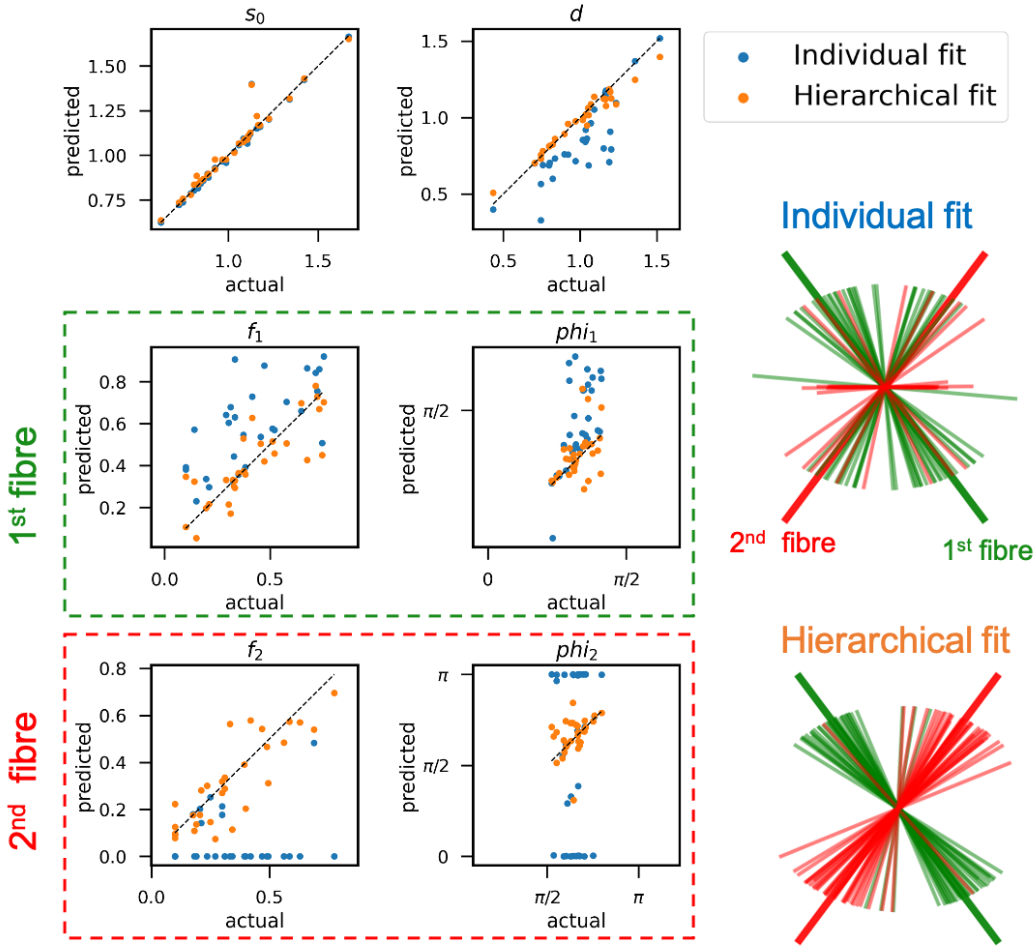


Figure 4.5: Scatter plots of estimated vs. actual parameters for the 2-fibre ball and stick model in a single voxel for 30 subjects using hierarchical and individual fits. For easier visualisation of the orientation parameters, we fixed the values of their spherical polar angle for all subjects to be 90 degrees and only showed estimates for the azimuthal angle phi. Each dot represents a subject, and the closer it is to the identity line (represented by the dashed line), the better the fit. The fibres in the individual fits are sorted by strength within each subject, while the fibres in the hierarchical fits are labelled according to the strength of the group-average fibre. For some subjects, the individual fit has swapped the values of f1 and f2 as well as the associated angles. In the lower-right corner, the estimated orientations of the sticks for individual fit and hierarchical fit are shown. The bold, elongated lines display the actual group average orientation. Despite the fact that the strength may have been in reverse order for some individuals, the hierarchical fit was able to label the sticks matching the correct labelling. In contrast, we observe a large amount of group variance in the individual fit method.

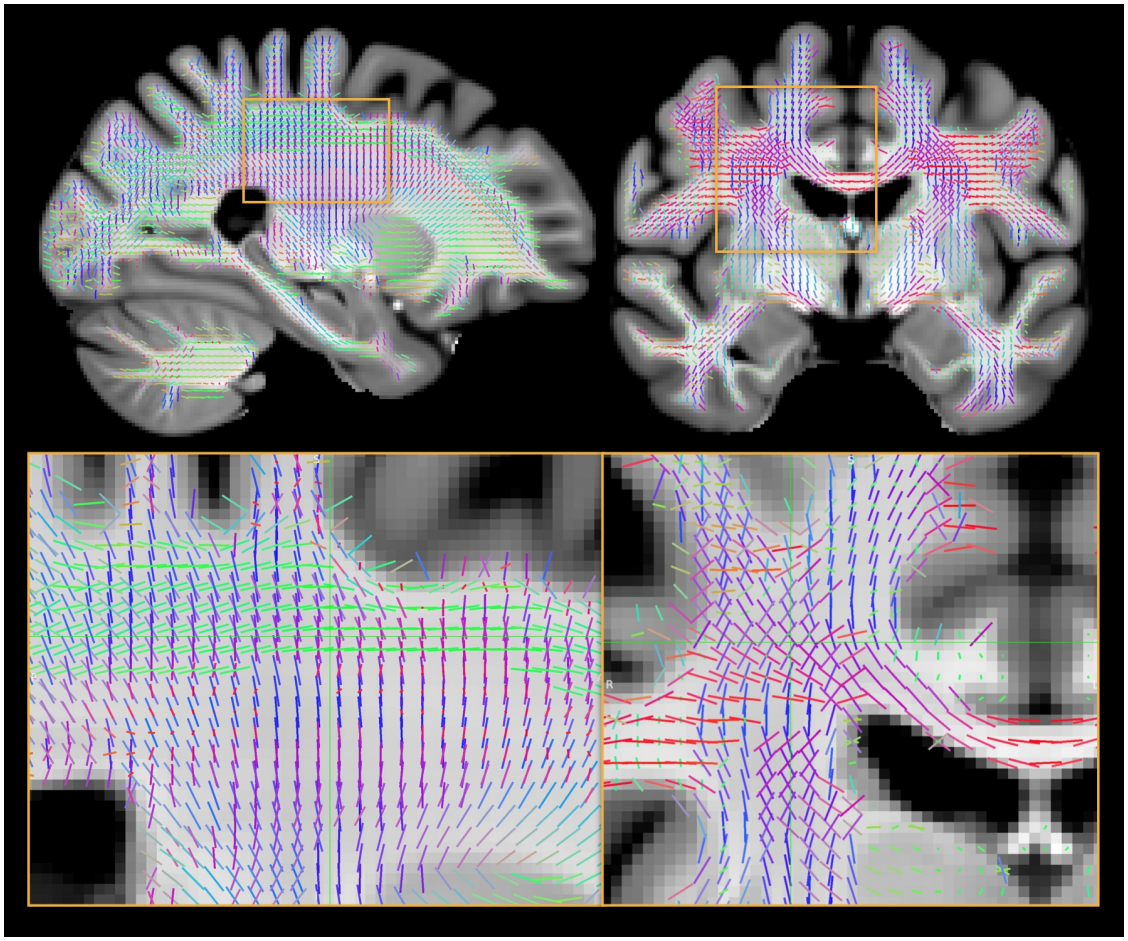


Figure 4.6: A sagittal and coronal slice of the white matter fixel template created using HCP diffusion data overlaid on the OMM1 T1-weighted image. The fixels that are present in each voxel are depicted by lines that are colour-coded according to their orientation (L-R: red, I-S: blue, A-P: green). The number of fixels is decided by thresholding the likelihood function improvement that is achieved by adding more sticks. Two sample patches of crossing fibre areas are magnified for better assessment. Although the technique does not involve regularisation across voxels, the fixels exhibit a high degree of spatial continuity. At the boundaries, we observe some fixels entering the grey matter radially.

The same fit is displayed in the lower panel, but this time the hierarchical model with 2 fibre populations was used. The results demonstrate that the fibres are aligned in two distinct groups and are matched between subjects. Furthermore, note that just the labelling has changed, and that the estimated orientations for each subject are extremely close to those from the individual fits. In other words, the hierarchical fit does not bias the orientations towards the group average and mostly only re-labels them.

Maps and a histogram of the cross-subject variation in orientation of the first

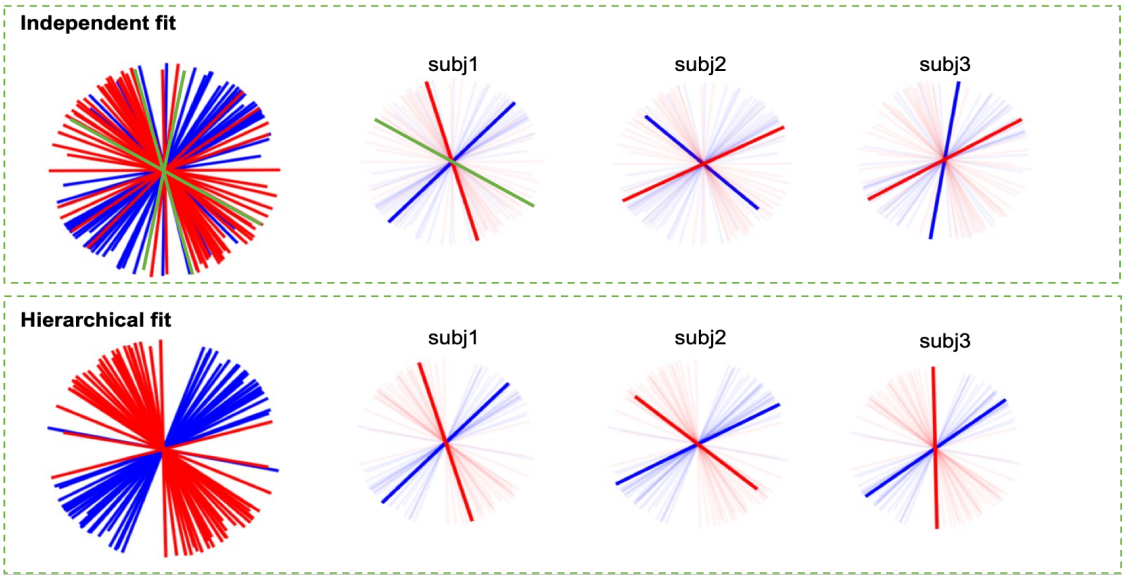


Figure 4.7: Estimated fixels in a sample voxel with crossing fibres for 30 subjects. The top panel shows aggregation of all subject fixels, along with three representative subjects. The first, second, and third fibres are represented by the colours blue, red, and green, respectively, in order of signal fraction. Various types of variability are observed across subjects: In this voxel in Subject 1, there are three distinct populations of fibres, in Subject 2, the first and second fixels are reversed compared to the other subjects, and in Subject 3, the fibres are oriented at an acute angle to one another. Any post hoc analysis to match fixels between subjects becomes challenging as a result of these differences. The hierarchical fit for the same data is shown in the lower panel. The labelling is based on the strength of the group average. In the aggregate plot, we can see that the fibres have organised into two distinct clusters. Similarly, the sample subjects demonstrate a strong correspondence with the template and between them.

fibre population (all white matter voxels) are shown in Figure 4.8 for both individual and hierarchical fits. The across subjects variability of the major fibre orientation in the individual fit is substantially higher than that for the hierarchical fit. This is due, in part, to the fact that certain subjects' first and second (or third) fibre populations have switched places, and also because the hierarchical fit tends to pull orientations closer to the group average when data is not strongly supportive of a particular orientation (e.g., in outlier or noisy subjects). The maps depict a slice of the cross-subject dispersion estimated in the white matter voxels. In the hierarchical fit, any high variance resides close to the grey matter boundaries, where there is no strong and consistent fibre orientation structure.

The number of fibre populations in each white matter voxel is depicted in Figure

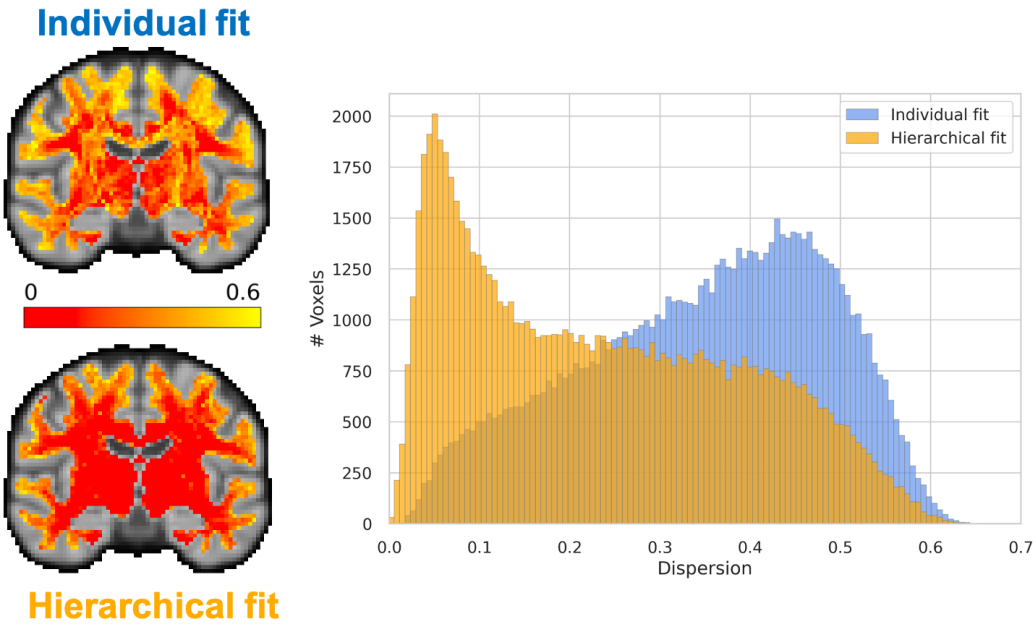


Figure 4.8: Variability of the primary fiber population across subjects with and without hierarchical fit: The histogram shows the distribution of across subject dispersion in all white matter voxels for both individual fits and hierarchical fit. A noticeable reduction in cross-subject variability is observed, indicating a higher degree of alignment among the primary fiber population across subjects for the majority of voxels. This reduction can be attributed to both the hierarchical framework drawing subject parameters towards the group average and also better relabelling of fibres. The accompanying maps display dispersions for a single slice of the standard brain. In the hierarchical fit, it is evident that the majority of high dispersion voxels are located at the grey matter boundaries, while fibres in the deep white matter exhibit increased alignment across subjects. In contrast, this pattern is not observed in the individual fitting approach, especially at the crossing fibre regions.

4.9 in example slices. The improvement in the likelihood function by adding an extra fibre population is the criterion for selecting the number of fibre populations in each voxel. The left histogram depicts the distribution of the difference in log-likelihood between models with two and one fibre population, summed across all participants. We apply a threshold to this histogram in order to identify voxels containing a single fibre population. The threshold is chosen by visually inspecting the maps to ensure a certain degree of symmetry across hemispheres and reasonable spatial homogeneity. The histograms on the top right displays the distribution of the log-likelihood difference between the 2 and 3 fibre models, but only for the voxels

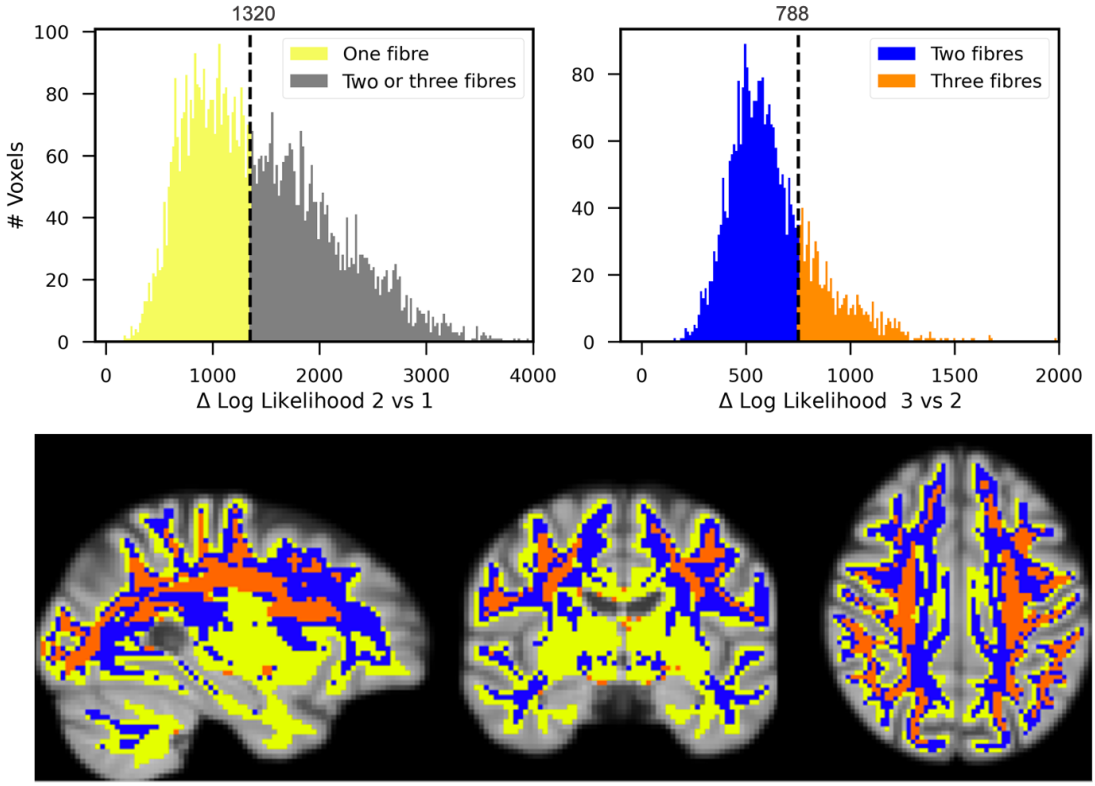


Figure 4.9: Determining the number of fixels present in each voxel. The histogram on the left displays the distribution (across the voxels in the white matter mask) of the increase in subject-averaged log likelihood from one to two fibres. We expect a better fit to data with more free parameters; hence all numbers are positive. To determine which voxels, contain only one fibre we threshold these differences. The threshold is chosen by visually inspecting the maps and taking into account across hemispheres symmetry and spatial homogeneity. The right histogram illustrates the improvement between two to three fibres exclusively for voxels that pass the criteria for one versus two models. This histogram is thresholded to indicate which voxels have two fixels and which have three fixels. The criteria for threshold is once more spatial homogeneity and cross-hemisphere symmetry. Maps depict a slicing through the template, with each voxel coloured according to the number of fixels it contains (yellow for 1, blue for 2, and orange for 3).

that are over the 1 fibre threshold. We apply a threshold to this histogram in order to identify voxels with two or three fibre populations.

4.5.3 Fixel based GLM

To illustrate how the fixel-based framework can be put to use for statistical analysis, we conducted a GLM analysis to assess the changes in white matter fibre tracts through ageing. For each subject, the HCP white matter fixel template was used as a prior to estimate fixel strength and orientation within each voxel. Then, we

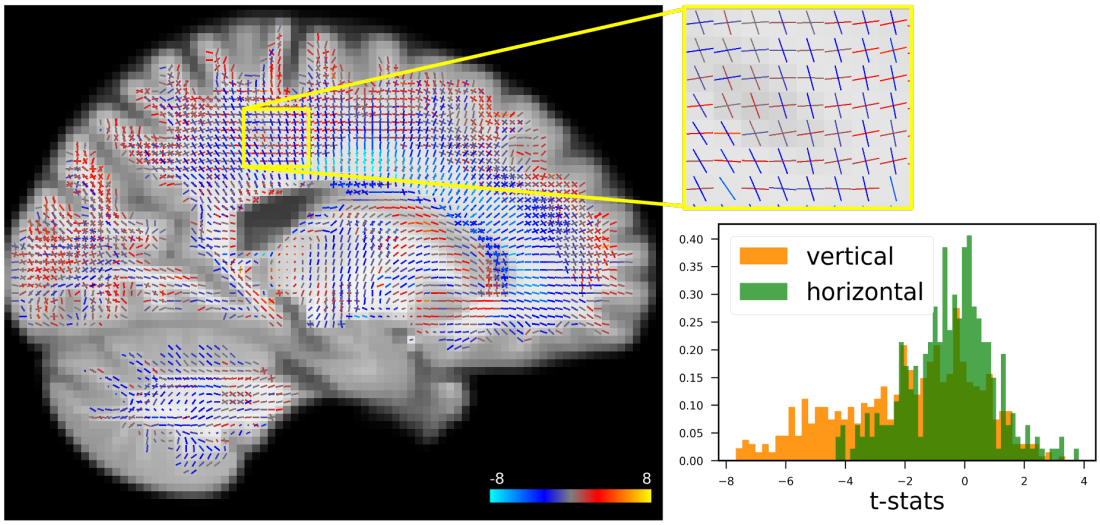


Figure 4.10: (Left) A sample slice of t-statistics for fixels. Colour of lines indicate correlation with age (blue for negative, red for positive correlations). (Right). Histogram of the t-statistics for horizontal and vertical fixels in the magnified region and 5 slices parallel to it. In this illustration, the distribution for vertical fibres is shown to be highly biased towards negative values, whereas the horizontal fibre distribution is centred around zero. In other words, on average, vertical (projection and commissural) tracts in this region degrade with age, whilst horizontal (association) tracts do not alter consistently.

performed a GLM with age as the independent variable (controlled for sex and head size) and fixel signal fraction as the dependent variable. For each fixel, we calculate a t-value from the GLM, which is then fed into the fixel-based TFCE to estimate the p-values via 5000 iterations of non-parametric permutation testing.

Figure 4.10 depicts a section of white matter with fixels coloured according to the t-statistics associated with the age regressor from the GLM; With red denoting an increase in fibre signal fraction and blue denoting a reduction. To facilitate a better visual evaluation, a patch of crossing fibre areas is magnified. We observe both positive and negative changes in the region, with the majority of horizontal (anterior-posterior) fibres displaying positive changes and the majority of vertical (superior-inferior) fibres displaying negative changes. The histogram depicts the distribution of t-stats separated for horizontal and vertical fibres inside the patch. The t-stats for horizontal fibres are centred around zero, whereas the vertical population is skewed to the left.

Figure 4.11 displays the changes in signal fractions and fractional anisotropy for a representative deep white matter voxel as a function of age. The FA scatter

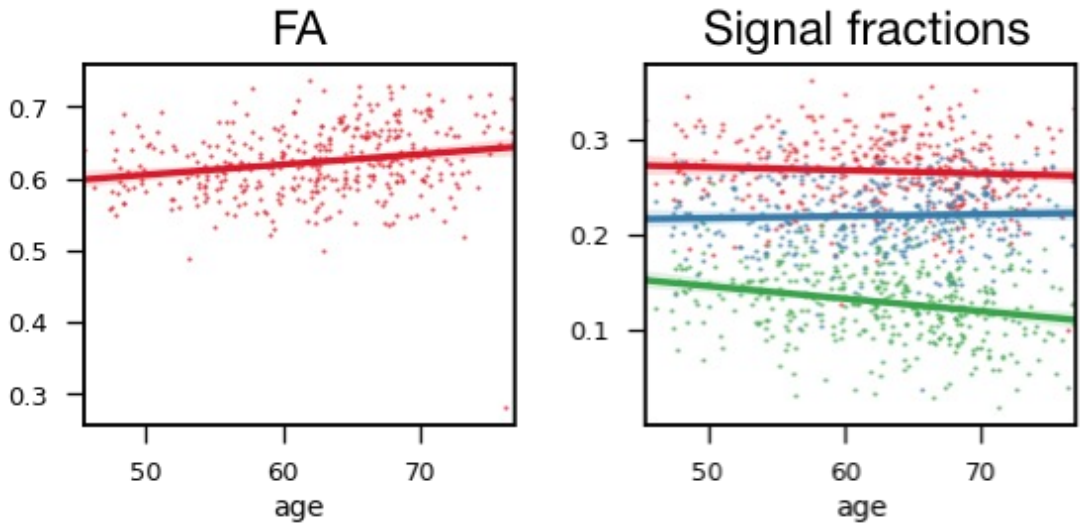


Figure 4.11: Comparison of FA and signal fractions in a sample voxel. Fractional anisotropy versus age for 400 subjects is shown in the scatter plot on the left (each dot represents a subject). In this voxel, the trend line (thick red line) suggests a positive FA-age correlation, which is conventionally attributed to an increase in fibre strength. The scatter plot on the right shows the signal proportions with trend lines for the three fibre populations present in this voxel (Red for first, Blue for second, and green for third fibre). While the first two fibre populations in this voxel changed very little with age, the third population decreased substantially with age. As a result, anisotropy increases because diffusion decreases along the direction of the changed fibre population relative to the other two.

plot trend indicates a significant rise in FA, which is sometimes interpreted as an increase in fibre "integrity", when we would usually expect a decrease with ageing. Examining the fixel signal fractions helps us interpret this counter-intuitive result. The signal fractions for the three fibre populations in this voxel as a function of age are displayed in the scatter plot on the right. Rather than an increase in overall strength, the population of fibres with the lowest signal fraction is decreasing at a faster rate compared with the other two. This can explain the apparent increase in FA since anisotropy should increase, particularly when there is a large angle between this fibre population and the other two (see Figure 4.1).

Figure 4.12 depicts a map of fixels with significant changes ($p<0.05$) following multiple comparison correction with fixel-based threshold free cluster enhancement. Only negative changes in the vicinity of the corpus callosum survived the significance test.

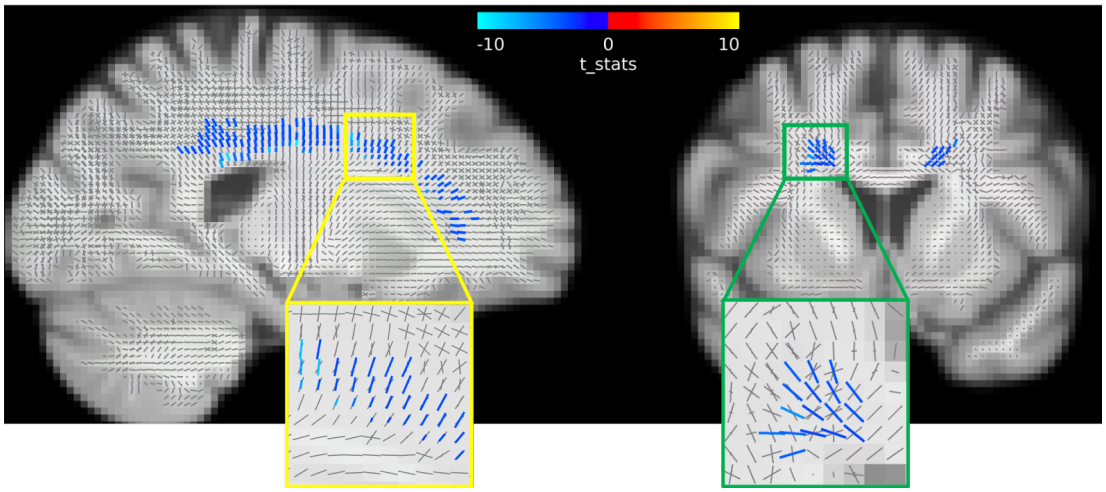


Figure 4.12: Statistically significant changes. Fixel-specific threshold-free cluster enhancement is applied to the t-stats derived from the GLM analysis on fixel signal fractions versus age (E = 1, H = 3). Permutation tests with 5000 repetitions on the TFCE values are used to estimate p-values for each fixel. The maps depict the significant fixels (colour-coded based on the t-statistics) with corrected $p < 0.05$. Only negative changes in signal fractions pass the statistical significance test, and they are all situated in the same tract (body of corpus callosum) in both hemispheres.

In Figure 4.13, we examine the changes with age in pairs of fixels forming crossing fibres in all voxels of the white matter. The figure displays the t-statistics of the first (x-axis) and second (y-axis) population of crossing fibres. Changes that occur in both populations of fibres and in the same direction (around the identity line in the scatter plots) represent voxel-wise changes, as opposed to fibre/fixel specific changes that occur only in one of the fibres (away from the identity line). High voxel-wise change regions (absolute t-stat in both fibres being more than 3) are shown by blue/green dots in the scatter plot and displayed in the same colour on the maps. Voxels that exhibit substantial changes (t-stats difference is greater than 3) specific to a fibre are identified by the colour red and displayed as such on the maps. The maps indicate that most voxel-based changes happen along the boundaries of grey matter, implying alterations in the partial volume of grey/white matter with age. In contrast, the majority of fixel-specific modifications arise within the white matter, likely mirroring changes in tract-specific organization of white matter. However, it is worth mentioning that alterations near the boundaries could partly be attributed to suboptimal registration, which can be tested independently.

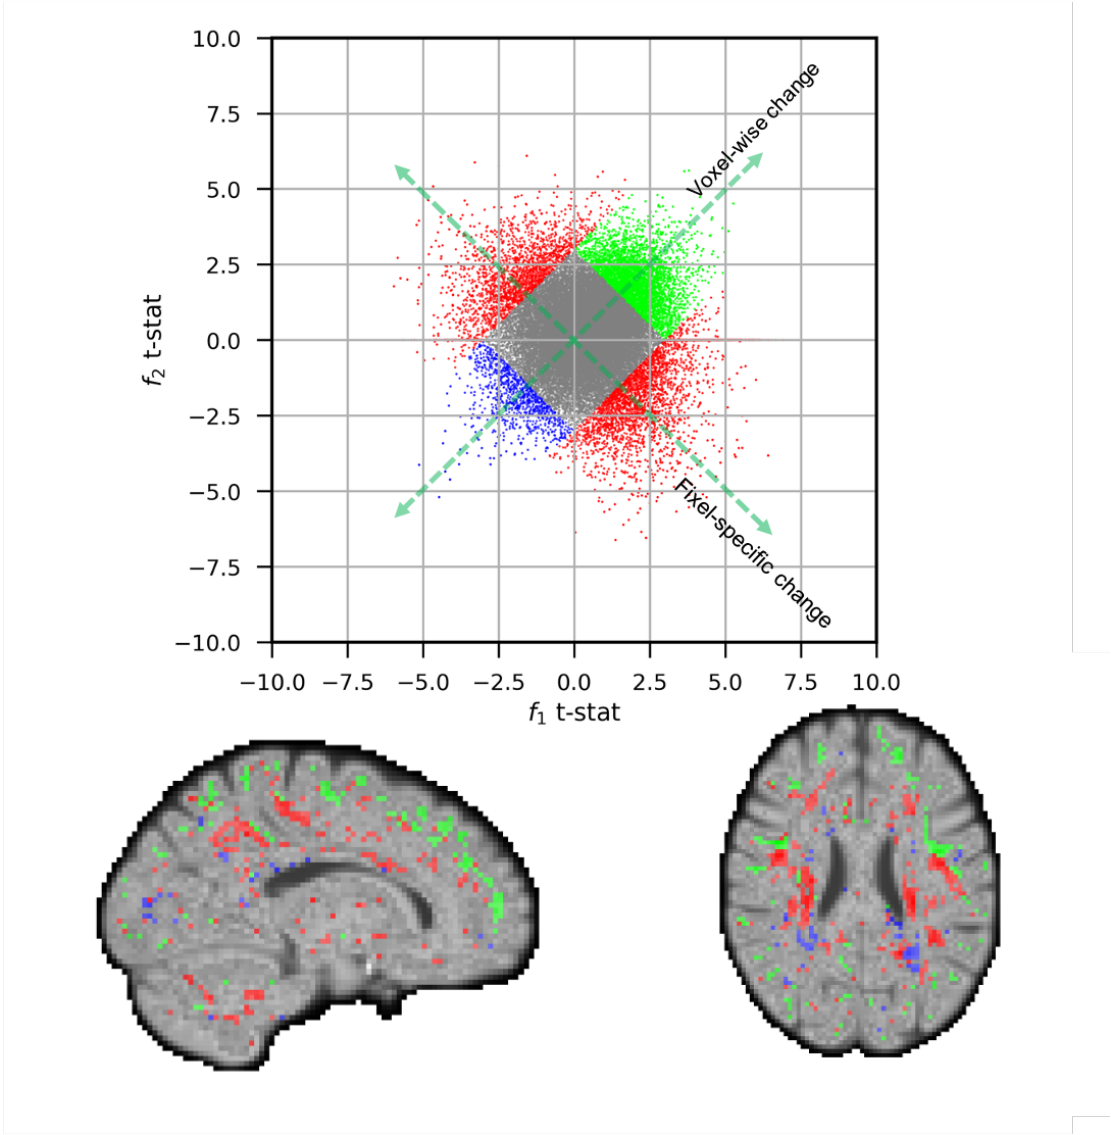


Figure 4.13: Voxelwise vs. fixel-specific changes. The scatter plot depicts the t-statistics for the signal fraction vs age for the first and second fibre populations in all white matter voxels (each point represents one voxel). If two fixels in the same voxel both change in the same direction (points that are near the identity line), this indicates that the change is occurring in the whole voxel in the same manner, i.e., it is a voxel-wise change. Examples of voxels with this type of change are displayed in blue for negative changes and green for positive changes. In contrast, if the fixels within a voxel change in opposite directions with ageing (points away from the identity line) this indicates that the change is fibre-specific. Examples of these voxels are highlighted in red. The maps display the areas where voxel- and fixel-specific changes are observed with the same colours as the scatter plot. The majority of the positive changes observed at the voxel level occur near the grey matter boundaries, whereas the majority of the fixel-specific changes or negative changes in voxel-level occur in deep white matter.

4.6 Software implementation

WHIM (White matter HIerarchical Modelling) is an open-source software package https://git.fmrib.ox.ac.uk/hossein/whim that can be used to run the full hierarchical inference, to generate new templates, or to run analysis on new data using existing templates. WHIM will be integrated with future versions of FSL.

4.7 Discussion

In this study, we introduce a new approach for extracting fibre-specific measurements (fixels based metrics) from diffusion MRI data. We used a hierarchical framework to fit the ball-and-sticks model to data from multiple subjects that offers some advantages over an individual fitting approach. First, the cross-subject fibre labelling issue is inherently resolved by the hierarchical structure of the model. This eliminates the need for post hoc fibre matching, allowing us to directly compare tract-specific parameters across individuals. The hierarchical structure also helps with robust parameter estimation for noisy subjects as it makes use of pooled data from multiple subjects. Natural expansions are viable in light of the hierarchical structure, such as hierarchically modelling subpopulations of individuals or employing a high-level model (such as age variables) as part of the inference process.

The hierarchical structure has some disadvantages as well. First and foremost, it is computationally more expensive than fitting the model to individuals separately. We addressed this issue by employing an EM algorithm, which converges within a few iterations and allows for both memory efficiency and convenient parallelization across subjects. However, we note that once a template has been created, it can be utilised as a prior for fitting new subjects and would have the same computational costs as the individual fitting approach. Individual variations in the number of fixels within a voxel constitute another issue. Our method assumes that all subjects have the same number of fixels and is inflexible in this regard, but at least the hierarchical model enforces comparison of fixels with similar characteristics. Moreover, the hierarchical model requires accurate registration of white matter both in terms of location and

local fibre orientation.Therefore, the interpretation of the results must always take into account potential contributing factors, including the effects of misregistration.

While the proposed framework is one option to extract fibre specific measurements, it is not the only one; the MRTRIX software [59] also has a very solid pipeline implemented that has been used in many studies. There are two key differences compared to our approach. Firstly, we estimate individual fibre orientations using the ball and sticks model while the other approach uses spherical deconvolution. Secondly, they solve the fibre assignment across subjects post hoc, after model fitting, whereas we perform it as part of the modelling process. Both methods provide proxy measures for fibre strength.

As of now, the diffusion tensor metrics such as FA and MD are the most frequently used measurements to examine the integrity of brain tissue. A GLM analysis for FA vs age reveals that a number of brain areas exhibit significant positive correlation with age. While a decrease of FA is often interpreted as reduced tract integrity, increases in FA are more difficult to interpret. In Figure 11, we provided a concrete example of real data demonstrating that a drop in the signal fraction of a crossing fibre in an example voxel in deep white matter is contributing to a rise in FA. Also, after applying the cluster corrections, the results of the fixel based GLM analysis show that all significant changes in the fibre strengths are negative. This would imply that the FA increases are confounded by the geometrical structure of the tracts and their relative alterations within each voxel.

4.7.1 Fixel template creation and usage

We have applied the proposed methods on a few subjects from the HCP young adult dataset to produce a white matter fixel template. This template contains the probability distribution for the orientation of fibres present in each white matter voxel in the Oxford Multimodal Template 1 (OMM-1), which can be used in the future for tract-specific statistical analysis. It may be essential to create a new template if a cohort-specific template is required (e.g., for disease or a different

age range), if using a different base diffusion model (other than ball-and-sticks), or if using a structural template other than OMM-1.

It is worth re-emphasising that whilst the hierarchical model can be fitted to a cohort of subjects using our EM algorithm, the resulting group parameters constitute a template that can be employed in its own right as a prior to fit additional subjects that are not part of the same cohort. This is not only useful for computational reasons. It can also benefit in fitting lower quality data than those used to create the template. For instance, it is well established that in diffusion modelling, there is a tension between acquiring multiple shells to have a better estimate of the diffusion parameters (e.g., ADC and kurtosis etc.) and acquiring multiple directions to have more accurate crossing fibre modelling. Most datasets compromise on one or the other. The template can be created using high-quality, multi-shell and multi-directions data, such as the HCP data, and then used as a prior to constrain the analysis of lower quality data. We have used exactly this approach here, with a template created using HCP, and deployed to analyse the UK Biobank data.

In this study, the hierarchical structure was applied to all parameters of the ball and sticks model, excluding the baseline signal (S0) parameter. This decision was made due to the potential for artefacts, such as bias field, to induce intrasubject variability in S0 with a complicated cross-subject distribution. Incorporating a parameter into the hierarchy can help to regularise its estimates especially in subjects with noisy data, but it can also induce biases. It is not required to include all parameters in the hierarchy; for example, signal fractions might be left out to avoid potential biases. However, stick orientations must always be maintained in the hierarchy to enable consistent labelling between subjects.

In order to determine how many fixels are present in each voxel, we have relied on the increase in log-likelihood as a measure of goodness of fit. In determining an appropriate cut-off for this metric, we take into account the spatial coherence in the produced maps and whether they are consistent with the information we currently have on the total number of tracts in each region. To determine the thresholds, one may also use metrics such as AIC and BIC. However, there is always an arbitrary

trade-off between the number of free parameters in the model and its goodness of fit. In any case, the thresholds are fairly flexible and can be re-adjusted with minimal effort if there is a compelling reason to do so.

In addition to the free parameters that are fitted using data, the hierarchical model also includes user-set hyperparameters (α_g, α_n) that constrain the prior on the group variance and the noise variance. Having a high α_g imposes a greater penalty on group variances to approach zero, hence increasing the chance to favour individual differences. A high α_n increases the penalty for estimating high noise variance. Changing either of these parameters within a suitable range, however, should not significantly alter the model's behaviour, although it may be desirable in some circumstances.

4.7.2 Future Directions

The hierarchical structure can be enhanced further in a number of ways. The template can be created in a nested hierarchical structure for multiple groups of subjects, and comparisons can then be made between group parameters. Another modification is to employ a regression model to incorporate independent variables (such as age) into the hierarchical structure. Moreover, in order to better handle noisy subjects, one can additionally include the uncertainty in subject parameter estimates into the model and convert it to a mixed-effect analysis.

The fixel template is created by employing the ball and sticks model to extract the orientations of the fibre tracts within each voxel. The ball and stick model were chosen as it is a well-tested model designed to robustly extract crossing fibres in the brain's white matter. However, this does not imply that the tract specific metrics are restricted to these model parameters. The presented template can be used as a starting point for developing more complex and specific biophysical models to fit the data. For instance, the template can be used to fit a multi-compartment model (e.g., NODDI [10]) with crossing fibres by utilising the information on the number and orientation of fibres in each voxel. The parameters of those

biophysical models can subsequently be employed as fixel-specific measurements that are compared across subjects.

As an illustration of the proposed method, we have employed a GLM framework to conduct fixel-specific statistical analysis in this study. The steps for performing such an analysis and the results to anticipate are laid out in detail here. Age has been used as an independent variable to clarify usage, but we hope that the tool will be useful to explore tract-specific alterations as a function of disease or more intricate biological processes. Additionally, one can investigate tract-specific alterations in the brain by incorporating these measurements with more data-driven methods such as Independent Component Analysis or Canonical Correlation Analysis. Moreover, tractography approaches can be applied to the template to reconstruct the major white matter tracts, enabling the across subjects' comparison of brain structure along the tracts.

In order to increase the power of statistical tests for fixel-based data, here we present a modified version of threshold free cluster enhancement. Similar to the original TFCE technique, this method utilises two hyperparameters (E and H) to adjust how heavily cluster extent and height play towards promoting significant clusters. In this study we have used the original method's (E=3, H=2) values for these parameters, which have been shown through simulations to be appropriate for voxel-wise data. Future development may involve adjusting them to accommodate for the different neighbourhood structure of fixel-based data compared to that of voxel-based data.

"This draft comes to an end here, but the story remains."

—Saadi, Persian poet

5 Conclusion and Future Directions

Contents

5.1 Bayesian estimation of change 105
5.2 Hierarchical modelling of crossing fibres in the white matter 107

In this thesis, we have proposed new methods to enhance the specificity of information regarding brain microstructure, derived from diffusion MRI (dMRI) data. These methods hold significant relevance in the field of microstructural modelling of dMRI data, a field that can be substantially enriched with increased specificity. This enhancement is crucial as, despite the data's high sensitivity to microstructural alterations and pathologies, it is remarkably unspecific. This lack of specificity manifests itself in two primary areas.

Firstly, even in the simplest scenarios, the models necessitate numerous parameters to approximate reality and generate meaningful results that can inform further research or clinical applications effectively. Yet, the available data is often scarce, leading to an obstacle in fitting the models accurately. This scarcity thus challenges reliable extraction of useful insights and necessitates numerous non-validated assumptions.

Secondly, specificity concerns arise from the spatial domain. This issue is particularly pronounced within white matter, which is a complex network of millions of fibres intricately intertwined. The collected diffusion data, being fundamentally volumetric, doesn't readily provide fibre-specific information. To this end, we design and utilise crossing fibre models capable of extracting fibre-specific information from the general volumetric data. A key challenge here is to account for cross-subject variability - the natural differences observed in fibre location and orientation between subjects. Respecting these differences while ensuring consistent comparability of fibre-specific information across different subjects is a delicate task, requiring precision, an in-depth understanding of the data, and proper modelling techniques, all of which are the focus of a part if this thesis.

Although dMRI serves as an invaluable tool for non-invasive brain examination, it has inherent limitations when studying brain microstructure. Notably, the acquisition of dMRI data is indirect and susceptible to noise and various artefacts. Also, brain tissue microstructure is notably complex. Furthermore, most models are primarily based on normal brain tissue. However, in clinical applications, the primary interest lies in examining diseased tissue, which may substantially differ from normal tissue. This difference implies that existing models and assumptions may not fully apply to diseased tissue, and results should be interpreted cautiously. As such, numerous assumptions and simplifications are necessary to extract meaningful microstructural information from the data, making it vital to employ complementary data sources and alternative approaches for validating these assumptions or modifying them to reflect reality more accurately. MRI-histology validation studies present a significant area of research in this regard. [99–103]

Our objective has been to bridge the gap between indirect data and the complex characteristics of brain tissue, thus improving the applicability of current data and models, particularly in clinical settings. We have explored the use of machine learning models and collective population data to gain more accurate information about brain microstructure. By leveraging these fronts, we have augmented the precision and reliability of microstructural information derived from dMRI.

5.1 Bayesian estimation of change

The primary objective of this work was to address the degeneracy problem in the field of microstructural modelling of diffusion MRI. In this regard, instead of fitting models to the data in order to estimate all potentially relevant parameters, we developed a tool to answer a more apt question: Which microstructural changes can explain the observed differences between a diseased group and a control group?

To do so, we make an assumption that the change is sparse, meaning that only a limited number of parameters systematically change between conditions. This enables the inclusion of more parameters without concerning ourselves with the fitting problem or the need for more unvalidated prior assumptions.

Indeed, the key assumption underpinning our approach is the premise of sparse change, an assumption that may not be universally valid across different conditions. Nevertheless, we believe it generally offers a more reliable premise compared to assuming parameter value sparsity, particularly when addressing biological processes. In well-constructed models that faithfully reflect reality, we anticipate that changes in different model parameters, each representing an independent biophysical factor, should be largely independent, and rarely occur together. It's worth noting, however, that our framework is Bayesian, designed to yield probabilities that reflect comparisons across all evaluated change models. If a change cannot be uniquely attributed to a parameter, the framework responds by providing a significant probability for all possibly changed parameters. Furthermore, we've devised a metric to assess the proximity between the modelled change and the observed one, which offers a valuable measure for verification and sanity check. This added layer of reassurance enhances confidence in the approach and its outcomes, reinforcing its potential for wider application.

Also, while advancements in data acquisition techniques enables the accumulation of supplementary data and facilitates the fitting of current state-of-the-art models [104], the integration of new data within the change framework can also stimulate further enhancements of these models. For instance, water exchange, which impacts the diffusion signal, can disclose additional details about the microstructure of

the tissue. [105] This can be quantified by modulating the diffusion time, but it also increases the number of free parameters required for modelling the data. Under such circumstances BENCH proves beneficial as it facilitates the exploration of alterations of more free parameters.

Additionally, BENCH's distinctive feature—its sole reliance on the forward model—offers an opportunity to get beyond limits set by models having analytical solutions. For instance, Monte Carlo simulations of diffusion can be implemented on synthetic substrates that closely resemble tissue microstructure, thus training change models that mimic realistic changes like alterations in axonal packing. We believe that this application of change framework could be a powerful instrument in the field for addressing problems that were once viewed as intractable.

In this study, the tool has been validated and applied to an example dataset, enabling us to examine microstructural changes in white matter hyperintensities. Concurrently, an application is in progress where histology data is combined with MRI to investigate the shifts in MRI signals attributed to histological changes. This endeavor employs change models trained on an expansive model for white matter, incorporating additional compartments for soma and glial cells. Subsequently, these change models are utilized to interpret regional variations in histology data derived from a single macaque brain [106]. This process necessitated modifications to the BENCH framework, enabling it to handle continuous data, rather than just differences between two groups, and incorporating a linear model for the correlation with histology. This strategy has the potential to forge a significant link between the parameters of the biophysical models of diffusion and histological measurements such as Myelin and Nissl staining. Our aspiration is to continue refining this tool, ultimately applying it to a broader spectrum of fundamental and clinical applications with multi-modal data, such as genetics.

Lastly, the proposed framework's potential extends beyond diffusion MRI modeling, providing the capacity for adaptation to other modalities or frameworks in the future. For instance, it could be used to model the BOLD MRI signal within a Dynamic Causal Modeling (DCM) framework [85]. This is especially

beneficial for DCM, given its inherently overparameterized nature, due to the multi-layered modeling of the BOLD signal. We are hopeful that researchers across various fields will find this tool advantageous and will incorporate it into their respective models and datasets.

5.2 Hierarchical modelling of crossing fibres in the white matter

The aim of this work was to develop a tool to extract fibre-specific information in a manner that allows for comparability across subjects. To this end, we employed hierarchical Bayesian modelling, which enable nested modelling of data from multiple subjects. This hierarchical framework offers two main advantages: it establishes correspondence in model parameters across subjects that governs comparability, and it provides more robust parameter estimation for individual subjects, especially those with noisy data.

In this work, we've implemented a pipeline based on the 'ball and sticks' model to create a white matter fibre template. The value of this template lies in its ability to extract a fibre-specific metric, which can then be utilized for statistical analysis within a GLM framework. We designed this template with the capacity for further refinement, allowing for the integration of more sophisticated models and thereby enabling the extraction of additional fibre-specific metrics. For example, after estimating the fixel orientation, it's possible to incorporate dispersion models to evaluate fibre cross-sections. Furthermore, the template has the potential to enhance single fibre models, such as NODDI. This enhancement could be accomplished by integrating fibre orientation density derived from fixel orientation, thereby facilitating the modelling of biophysical models of crossing fibres.

Furthermore, the template's basis on abstract model parameters allows for adaptability across diverse datasets with varying acquisition settings. This flexibility is particularly valuable in clinical contexts where data quality is often poor. For example, models could be trained on high-quality research datasets and subsequently used to estimate parameters in clinical images.

In the implementation of this hierarchical framework, we have established our work on certain assumptions such as near-perfect registration and minimal dislocation or reorientation of tracts. While these assumptions generally hold true for a homogeneous group of subjects, they might not extend appropriately to severely diseased tissues or across age ranges. To counteract this potential limitation, we suggest adapting the hierarchical framework into a nested structure where each sub-population has its own population distribution. Through this adaptation, we can even render a more robust comparison across populations by directly modelling the populations and their differences.

In general, this approach holds considerable potential to enhance the robustness and utility of clinical imaging data. We anticipate that this will be particularly useful for studying diseases involving neurodegeneration in specific brain tract, offering a more coherent and accessible approach to fibre-specific analysis in diffusion MRI data.

References

[1] P Mansfield. "Multi-planar image formation using NMR spin echoes". In: *Journal of Physics C: Solid State Physics* 10.3 (Feb. 14, 1977), pp. L55–L58. URL: https://iopscience.iop.org/article/10.1088/0022-3719/10/3/004 (visited on 05/15/2023).

[2] P. C. Lauterbur. "Image Formation by Induced Local Interactions: Examples Employing Nuclear Magnetic Resonance". In: *Nature* 242.5394 (Mar. 1973), pp. 190–191. URL: https://www.nature.com/articles/242190a0 (visited on 05/15/2023).

[3] Denis Le Bihan et al. "Diffusion tensor imaging: Concepts and applications". In: *Journal of Magnetic Resonance Imaging* 13.4 (Apr. 2001), pp. 534–546. URL: https://onlinelibrary.wiley.com/doi/10.1002/jmri.1076 (visited on 05/15/2023).

[4] Denis Le Bihan. "Looking into the functional architecture of the brain with diffusion MRI". In: *Nature Reviews Neuroscience* 4.6 (June 2003), pp. 469–480. URL: http://www.nature.com/articles/nrn1119 (visited on 05/15/2023).

[5] Nikos K. Logothetis. "The Underpinnings of the BOLD Functional Magnetic Resonance Imaging Signal". In: *The Journal of Neuroscience* 23.10 (May 15, 2003), pp. 3963–3971. URL: https://www.jneurosci.org/lookup/doi/10.1523/JNEUROSCI.23-10-03963.2003 (visited on 05/15/2023).

[6] Mark Jenkinson et al. "FSL". In: *NeuroImage* 62.2 (Aug. 2012). Publisher: Elsevier BV, pp. 782–790.

[7] Derek K. Jones, ed. *Diffusion MRI: theory, methods, and application*. Oxford ; New York: Oxford University Press, 2010. 767 pp.

[8] Heidi Johansen-Berg and Timothy E. J. Behrens. *Diffusion MRI: from quantitative measurement to in-vivo neuroanatomy*. 2nd ed. OCLC: 899573164. Amsterdam: Elsevier Science, 2014.

[9] Dmitry S Novikov et al. "Quantifying brain microstructure with diffusion MRI: Theory and parameter estimation". In: *NMR in Biomedicine* 32.4 (2019), e3998.

[10] Hui Zhang et al. "NODDI: Practical in vivo neurite orientation dispersion and density imaging of the human brain". en. In: *NeuroImage* 61.4 (July 2012), pp. 1000–1016. URL: http://www.sciencedirect.com/science/article/pii/S1053811912003539 (visited on 09/09/2020).

[11] Yaniv Assaf et al. "AxCaliber: a method for measuring axon diameter distribution from diffusion MRI". In: *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine* 59.6 (2008), pp. 1347–1354.

[12] Dmitry S. Novikova et al. "Rotationally-invariant mapping of scalar and orientational metrics of neuronal microstructure with diffusion MRI". en. In: *NeuroImage* 174 (July 2018), pp. 518–538. URL: https://www.sciencedirect.com/science/article/pii/S1053811918301915 (visited on 09/09/2020).

[13] Yaniv Assaf and Peter J. Basser. "Composite hindered and restricted model of diffusion (CHARMED) MR imaging of the human brain". In: *NeuroImage* 27.1 (Aug. 2005), pp. 48–58. URL: https://linkinghub.elsevier.com/retrieve/pii/S1053811905002259 (visited on 05/15/2023).

[14] Ileana O. Jelescu et al. "Degeneracy in model parameter estimation for multi-compartmental diffusion in neuronal tissue". en. In: *NMR in biomedicine* 29.1 (Jan. 2016), p. 33. URL: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC4920129/ (visited on 09/09/2020).

[15] Stamatios N Sotiropoulos, Timothy EJ Behrens, and Saad Jbabdi. "Ball and rackets: inferring fiber fanning from diffusion-weighted MRI". In: *Neuroimage* 60.2 (2012), pp. 1412–1425.

[16] T. E. J. Behrens et al. "Characterization and propagation of uncertainty in diffusion-weighted MR imaging". eng. In: *Magnetic Resonance in Medicine* 50.5 (Nov. 2003), pp. 1077–1088.

[17] J-Donald Tournier et al. "Direct estimation of the fiber orientation density function from diffusion-weighted MRI data using spherical deconvolution". In: *Neuroimage* 23.3 (2004), pp. 1176–1185.

[18] Stéphanie Debette and HS Markus. "The clinical importance of white matter hyperintensities on brain magnetic resonance imaging: systematic review and meta-analysis". In: *Bmj* 341 (2010).

[19] Marek Kubicki et al. "The application of DTI to investigate white matter abnormalities in schizophrenia". In: *Annals of the New York Academy of Sciences* 1064.1 (2005), pp. 134–148.

[20] Nobutsugu Hirono et al. "Impact of white matter changes on clinical manifestation of Alzheimer's disease: a quantitative study". In: *Stroke* 31.9 (2000), pp. 2182–2188.

[21] Mark A Horsfield and Derek K Jones. "Applications of diffusion-weighted and diffusion tensor MRI to white matter diseases–a review". In: *NMR in Biomedicine: An International Journal Devoted to the Development and Application of Magnetic Resonance In Vivo* 15.7-8 (2002), pp. 570–577.

[22] Nicolae Sarbu et al. "White matter diseases with radiologic-pathologic correlation". In: *Radiographics* 36.5 (2016), pp. 1426–1447.

[23] Douglas C Dean et al. “Characterizing longitudinal white matter development during early childhood”. In: *Brain Structure and Function* 220 (2015), pp. 1921–1933.

[24] Naama Barnea-Goraly et al. “White matter development during childhood and adolescence: a cross-sectional diffusion tensor imaging study”. In: *Cerebral cortex* 15.12 (2005), pp. 1848–1854.

[25] F. Bloch. “Nuclear Induction”. In: *Phys. Rev.* 70 (7-8 Oct. 1946), pp. 460–474. URL: https://link.aps.org/doi/10.1103/PhysRev.70.460.

[26] John E Tanner and Edward O Stejskal. “Restricted self-diffusion of protons in colloidal systems by the pulsed-gradient, spin-echo method”. In: *The Journal of Chemical Physics* 49.4 (1968), pp. 1768–1777.

[27] Edward O Stejskal and John E Tanner. “Spin diffusion measurements: spin echoes in the presence of a time-dependent field gradient”. In: *The journal of chemical physics* 42.1 (1965), pp. 288–292.

[28] David C. Van Essen et al. “The WU-Minn Human Connectome Project: An overview”. In: *NeuroImage* 80 (Oct. 2013). Publisher: Elsevier BV, pp. 62–79.

[29] Stamatios N. Sotiropoulos et al. “Advances in diffusion MRI acquisition and processing in the Human Connectome Project”. en. In: *NeuroImage*. Mapping the Connectome 80 (Oct. 2013), pp. 125–143. URL: https://www.sciencedirect.com/science/article/pii/S105381191300551X (visited on 06/13/2021).

[30] Jesper L.R. Andersson, Stefan Skare, and John Ashburner. “How to correct susceptibility distortions in spin-echo echo-planar images: application to diffusion tensor imaging”. In: *NeuroImage* 20.2 (2003), pp. 870–888. URL: https://www.sciencedirect.com/science/article/pii/S1053811903003367.

[31] Stephen M. Smith et al. “Advances in functional and structural MR image analysis and implementation as FSL”. In: *NeuroImage* 23 (Jan. 2004), S208–S219. URL: https://linkinghub.elsevier.com/retrieve/pii/S1053811904003933 (visited on 05/15/2023).

[32] Jesper L.R. Andersson and Stamatios N. Sotiropoulos. “An integrated approach to correction for off-resonance effects and subject movement in diffusion MR imaging”. In: *NeuroImage* 125 (2016), pp. 1063–1078. URL: https://www.sciencedirect.com/science/article/pii/S1053811915009209.

[33] Anastasia Yendiki et al. “Spurious group differences due to head motion in a diffusion MRI study”. In: *Neuroimage* 88 (2014), pp. 79–90.

[34] Mark S Graham et al. “Quantitative assessment of the susceptibility artefact and its interaction with motion in diffusion MRI”. In: *PloS one* 12.10 (2017), e0185647.

[35] P.J. Basser, J. Mattiello, and D. LeBihan. “MR diffusion tensor spectroscopy and imaging”. In: *Biophysical Journal* 66.1 (Jan. 1994), pp. 259–267. URL: https://linkinghub.elsevier.com/retrieve/pii/S0006349594807751 (visited on 05/15/2023).

[36] Peter J Basser, James Mattiello, and Denis LeBihan. “Estimation of the effective self-diffusion tensor from the NMR spin echo”. In: *Journal of Magnetic Resonance, Series B* 103.3 (1994), pp. 247–254.

[37] O Ciccarelli et al. “Investigation of MS normal-appearing brain using diffusion tensor MRI with clinical correlations”. In: *Neurology* 56.7 (2001), pp. 926–933.

[38] Ahmed T Toosy et al. “Characterizing function–structure relationships in the human visual system with functional MRI and diffusion tensor imaging”. In: *Neuroimage* 21.4 (2004), pp. 1452–1463.

[39] Andrew L Alexander et al. “Diffusion tensor imaging of the brain”. In: *Neurotherapeutics* 4.3 (2007), pp. 316–329.

[40] Olga Ciccarelli et al. “From diffusion tractography to quantitative white matter tract measures: a reproducibility study”. In: *Neuroimage* 18.2 (2003), pp. 348–359.

[41] Heidi Johansen-Berg et al. “Functional–anatomical validation and individual variation of diffusion tractography-based segmentation of the human thalamus”. In: *Cerebral cortex* 15.1 (2005), pp. 31–39.

[42] Ping Li, Jennifer Legault, and Kaitlyn A Litcofsky. “Neuroplasticity as a function of second language learning: Anatomical changes in the human brain”. In: *Cortex* 58 (2014), pp. 301–324.

[43] Armin Schwartzman, Robert F Dougherty, and Jonathan E Taylor. “Cross-subject comparison of principal diffusion direction maps”. In: *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine* 53.6 (2005), pp. 1423–1431.

[44] Pawel J Winklewski et al. “Understanding the physiopathology behind axial and radial diffusivity changes—what do we know?” In: *Frontiers in neurology* 9 (2018), p. 92.

[45] Balin Balinov et al. “The NMR self-diffusion method applied to restricted diffusion. Simulation of echo attenuation from molecules in spheres and between planes”. In: *Journal of Magnetic Resonance, Series A* 104.1 (1993), pp. 17–25.

[46] CH Neuman. “Spin echo of spins diffusing in a bounded medium”. In: *The Journal of Chemical Physics* 60.11 (1974), pp. 4508–4511.

[47] Marco Palombo et al. “SANDI: A compartment-based model for non-invasive apparent soma and neurite imaging by diffusion MRI”. In: *NeuroImage* 215 (2020), p. 116835. URL: https://www.sciencedirect.com/science/article/pii/S1053811920303220.

[48] Marco Palombo et al. “Joint estimation of relaxation and diffusion tissue parameters for prostate cancer with relaxation-VERDICT MRI”. In: *Scientific Reports* 13.1 (2023), p. 2999.

[49] Aaron Szafer, Jianhui Zhong, and John C Gore. “Theoretical model for water diffusion in tissues”. In: *Magnetic resonance in medicine* 33.5 (1995), pp. 697–712.

[50] Derek K Jones and Mara Cercignani. “Twenty-five pitfalls in the analysis of diffusion MRI data”. In: *NMR in Biomedicine* 23.7 (2010), pp. 803–820.

[51] Lindsay Snook, Chris Plewes, and Christian Beaulieu. “Voxel based versus region of interest analysis in diffusion tensor imaging of neurodevelopment”. In: *Neuroimage* 34.1 (2007), pp. 243–252.

[52] Gwenaëlle Douaud et al. “DTI measures in crossing-fibre areas: increased diffusion anisotropy reveals early white matter alteration in MCI and mild Alzheimer’s disease”. In: *Neuroimage* 55.3 (2011), pp. 880–890.

[53] Stephen M. Smith et al. “Tract-based spatial statistics: Voxelwise analysis of multi-subject diffusion data”. In: *NeuroImage* 31.4 (July 2006). Publisher: Elsevier BV, pp. 1487–1505.

[54] Lars T. Westlye et al. “Life-Span Changes of the Human Brain White Matter: Diffusion Tensor Imaging (DTI) and Volumetry”. In: *Cerebral Cortex* 20.9 (Dec. 2009), pp. 2055–2068. eprint: `https://academic.oup.com/cercor/article-pdf/20/9/2055/17303951/bhp280.pdf`. URL: `https://doi.org/10.1093/cercor/bhp280`.

[55] Jan Scholz et al. “Training induces changes in white-matter architecture”. In: *Nature neuroscience* 12.11 (2009), pp. 1370–1371.

[56] Kirsi Maria Kinnunen et al. “White matter damage and cognitive impairment after traumatic brain injury”. In: *Brain* 134.2 (2011), pp. 449–463.

[57] Michael Bach et al. “Methodological considerations on tract-based spatial statistics (TBSS)”. In: *Neuroimage* 100 (2014), pp. 358–369.

[58] Matthew F Glasser and James K Rilling. “DTI tractography of the human brain’s language pathways”. In: *Cerebral cortex* 18.11 (2008), pp. 2471–2482.

[59] J-Donald Tournier et al. *MRtrix3: A fast, flexible and open software framework for medical image processing and visualisation*. Cold Spring Harbor Laboratory, Feb. 15, 2019.

[60] David A. Raffelt et al. “Investigating white matter fibre density and morphology using fixel-based analysis”. In: *NeuroImage* 144.0 (Jan. 2017). Publisher: Elsevier BV, pp. 58–73.

[61] Ben Jeurissen et al. “Multi-tissue constrained spherical deconvolution for improved analysis of multi-shell diffusion MRI data”. In: *NeuroImage* 103 (2014), pp. 411–426.

[62] Danilo Rezende and Shakir Mohamed. “Variational inference with normalizing flows”. In: *International conference on machine learning*. PMLR. 2015, pp. 1530–1538.

[63] Jens H Jensen et al. “Diffusional kurtosis imaging: the quantification of non-gaussian water diffusion by means of magnetic resonance imaging”. In: *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine* 53.6 (2005), pp. 1432–1440.

[64] Dmitry S. Novikov et al. “Quantifying brain microstructure with diffusion MRI: Theory and parameter estimation”. en. In: *NMR in Biomedicine* 32.4 (Apr. 2019). Publisher: John Wiley & Sons, Ltd. URL: `https://onlinelibrary.wiley.com/doi/10.1002/nbm.3998` (visited on 09/09/2020).

[65] Marco Reisert, Elias Kellner, and Valerij G. Kiselev. “Disentangling micro from mesostructure by diffusion MRI: A Bayesian approach”. en. In: *NeuroImage* 147 (Feb. 2017). Publisher: Academic Press, pp. 964–975. URL: `https://www.sciencedirect.com/science/article/pii/S1053811916305353` (visited on 09/09/2020).

[66] Björn Lampinen et al. “Searching for the neurite density with diffusion MRI: Challenges for biophysical modeling”. In: *Human brain mapping* 40.8 (2019), pp. 2529–2545.

[67] Marco Reisert, Valerij G. Kiselev, and Bibek Dhital. “A unique analytical solution of the white matter standard model using linear and planar encodings”. In: *Magnetic Resonance in Medicine* 81.6 (2019), pp. 3819–3825. eprint: `https://onlinelibrary.wiley.com/doi/pdf/10.1002/mrm.27685`. URL: `https://onlinelibrary.wiley.com/doi/abs/10.1002/mrm.27685`.

[68] Santiago Coelho et al. “Resolving degeneracy in diffusion MRI biophysical model parameter estimation using double diffusion encoding”. In: *Magnetic resonance in medicine* 82.1 (2019), pp. 395–410.

[69] Björn Lampinen et al. “Towards unconstrained compartment modeling in white matter using diffusion-relaxation MRI with tensor-valued diffusion encoding”. In: *Magnetic resonance in medicine* 84.3 (2020), pp. 1605–1623.

[70] Niels D Prins and Philip Scheltens. “White matter hyperintensities, cognitive impairment and dementia: an update”. In: *Nature Reviews Neurology* 11.3 (2015), pp. 157–165.

[71] Joanna M Wardlaw et al. “Neuroimaging standards for research into small vessel disease and its contribution to ageing and neurodegeneration”. In: *The Lancet Neurology* 12.8 (2013), pp. 822–838.

[72] Alida A Gouw et al. “Heterogeneity of small vessel disease: a systematic review of MRI and histopathology correlations”. In: *Journal of Neurology, Neurosurgery & Psychiatry* 82.2 (2011), pp. 126–135. eprint: `https://jnnp.bmj.com/content/82/2/126.full.pdf`. URL: `https://jnnp.bmj.com/content/82/2/126`.

[73] Pauli Virtanen et al. “SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python”. In: *Nature Methods* 17 (2020), pp. 261–272.

[74] Michael Kazhdan, Thomas Funkhouser, and Szymon Rusinkiewicz. “Rotation invariant spherical harmonic representation of 3 d shape descriptors”. In: *Symposium on geometry processing*. Vol. 6. 2003, pp. 156–164.

[75] Mauro Zucchelli, Samuel Deslauriers-Gauthier, and Rachid Deriche. “A computational Framework for generating rotation invariant features and its application in diffusion MRI”. en. In: *Medical Image Analysis* 60 (Feb. 2020), p. 101597. URL: `http://www.sciencedirect.com/science/article/pii/S1361841519301379` (visited on 09/09/2020).

[76] Karla L. Miller et al. “Multimodal population brain imaging in the UK Biobank prospective epidemiological study”. eng. In: *Nature Neuroscience* 19.11 (2016), pp. 1523–1536.

[77] Fidel Alfaro-Almagro et al. “Image processing and Quality Control for the first 10,000 brain imaging datasets from UK Biobank”. In: *Neuroimage* 166 (2018), pp. 400–424.

[78] Yongyue Zhang, J Michael Brady, and Stephen Smith. "Hidden Markov random field model for segmentation of brain MR image". In: *Medical Imaging 2000: Image Processing*. Vol. 3979. International Society for Optics and Photonics. 2000, pp. 1126–1137.

[79] Mark W Woolrich et al. "Bayesian analysis of neuroimaging data in FSL". In: *Neuroimage* 45.1 (2009), S173–S186.

[80] Jesper LR Andersson, Mark Jenkinson, and Stephen Smith. "High resolution nonlinear registration with simultaneous modelling of intensities". In: *BioRxiv* (2019), p. 646802.

[81] Ludovica Griffanti et al. "BIANCA (Brain Intensity AbNormality Classification Algorithm): A new tool for automated segmentation of white matter hyperintensities". en. In: *NeuroImage* 141 (Nov. 2016), pp. 191–205. URL: http://www.sciencedirect.com/science/article/pii/S1053811916303251 (visited on 12/14/2020).

[82] Rutger H. J. Fick, Demian Wassermann, and Rachid Deriche. "The Dmipy Toolbox: Diffusion MRI Multi-Compartment Modeling and Microstructure Recovery Made Easy". In: *Frontiers in Neuroinformatics* 13 (2019), p. 64. URL: https://www.frontiersin.org/article/10.3389/fninf.2019.00064.

[83] Amy FD Howard et al. "Estimating intra-axonal axial diffusivity in the presence of fibre orientation dispersion". In: *bioRxiv* (2020). eprint: https://www.biorxiv.org/content/early/2020/10/10/2020.10.09.332700.full.pdf. URL: https://www.biorxiv.org/content/early/2020/10/10/2020.10.09.332700.

[84] Nicolas Kunz, Analina R da Silva, and Ileana O Jelescu. "Intra-and extra-axonal axial diffusivities in the white matter: Which one is faster?" In: *Neuroimage* 181 (2018), pp. 314–322.

[85] Karl J Friston, Lee Harrison, and Will Penny. "Dynamic causal modelling". In: *Neuroimage* 19.4 (2003), pp. 1273–1302.

[86] John Ashburner and Karl J. Friston. "Voxel-Based Morphometry—The Methods". In: *NeuroImage* 11.6 (June 2000), pp. 805–821. URL: https://linkinghub.elsevier.com/retrieve/pii/S1053811900905822 (visited on 05/03/2023).

[87] Thijs Dhollander et al. "Fixel-based Analysis of Diffusion MRI: Methods, Applications, Challenges and Opportunities". In: *NeuroImage* 241 (Nov. 2021), p. 118417. URL: https://linkinghub.elsevier.com/retrieve/pii/S1053811921006923 (visited on 05/03/2023).

[88] Jan Scholz et al. "Training induces changes in white-matter architecture". In: *Nature Neuroscience* 12.11 (Nov. 2009), pp. 1370–1371. URL: http://www.nature.com/articles/nn.2412 (visited on 05/03/2023).

[89] Saad Jbabdi, Timothy E.J. Behrens, and Stephen M. Smith. "Crossing fibres in tract-based spatial statistics". In: *NeuroImage* 49.1 (Jan. 2010), pp. 249–256. URL: https://linkinghub.elsevier.com/retrieve/pii/S105381190900946X (visited on 05/03/2023).

[90] T. E. J. Behrens et al. "Probabilistic diffusion tractography with multiple fibre orientations: What can we gain?" en. In: *NeuroImage* 34.1 (Jan. 2007), pp. 144–155. URL: http://www.sciencedirect.com/science/article/pii/S1053811906009360 (visited on 09/09/2020).

[91] S Smith and T Nichols. "Threshold-free cluster enhancement: Addressing problems of smoothing, threshold dependence and localisation in cluster inference". In: *NeuroImage* 44.1 (Jan. 1, 2009). Publisher: Elsevier BV, pp. 83–98.

[92] Karla L Miller et al. "Multimodal population brain imaging in the UK Biobank prospective epidemiological study". In: *Nature neuroscience* 19.11 (2016), pp. 1523–1536.

[93] J. A. Nelder and R. Mead. "A simplex method for function minimization". In: *The Computer Journal* 7.4 (Jan. 1, 1965). Publisher: Oxford University Press (OUP), pp. 308–313.

[94] Pauli Virtanen et al. "SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python". In: *Nature Methods* 17 (2020), pp. 261–272.

[95] Frederik J. Lange et al. "A Symmetric Prior for the Regularisation of Elastic Deformations: Improved anatomical plausibility in nonlinear image registration". In: *NeuroImage* 219 (Oct. 2020), p. 116962. URL: https://linkinghub.elsevier.com/retrieve/pii/S1053811920304481 (visited on 05/03/2023).

[96] Christopher Arthofer et al. "Multimodal MRI template construction from UK Biobank: Oxford-MM-0". In: *Organisation for Human Brain Mapping (OHBM)* (2021).

[97] Y. Zhang, M. Brady, and S. Smith. "Segmentation of brain MR images through a hidden Markov random field model and the expectation-maximization algorithm". In: *IEEE Transactions on Medical Imaging* 20.1 (2001). Publisher: Institute of Electrical and Electronics Engineers (IEEE), pp. 45–57.

[98] Fidel Alfaro-Almagro et al. "Image processing and Quality Control for the first 10,000 brain imaging datasets from UK Biobank". In: *NeuroImage* 166 (Feb. 2018), pp. 400–424. URL: https://linkinghub.elsevier.com/retrieve/pii/S1053811917308613 (visited on 05/03/2023).

[99] Jeroen Mollink et al. "Evaluating fibre orientation dispersion in white matter: Comparison of diffusion MRI, histology and polarized light imaging". In: *NeuroImage* 157 (2017), pp. 561–574. URL: https://www.sciencedirect.com/science/article/pii/S1053811917304706.

[100] Amy FD. Howard et al. "Joint modelling of diffusion MRI and microscopy". In: *NeuroImage* 201 (2019), p. 116014. URL: https://www.sciencedirect.com/science/article/pii/S1053811919305956.

[101] Jason P Lerch et al. "Studying neuroanatomy using MRI". In: *Nature neuroscience* 20.3 (2017), pp. 314–326.

[102] Nadja Van Camp et al. "A complementary diffusion tensor imaging (DTI)-histological study in a model of Huntington's disease". In: *Neurobiology of aging* 33.5 (2012), pp. 945–959.

[103] Luis Concha et al. “In vivo diffusion tensor imaging and histopathology of the fimbria-fornix in temporal lobe epilepsy”. In: *Journal of Neuroscience* 30.3 (2010), pp. 996–1002.

[104] Santiago Coelho et al. “Reproducibility of the Standard Model of diffusion in white matter on clinical MRI systems”. In: *NeuroImage* 257 (2022), p. 119290. URL: https://www.sciencedirect.com/science/article/pii/S1053811922004104.

[105] Ileana O. Jelescu et al. “Neurite Exchange Imaging (NEXI): A minimal model of diffusion in gray matter with inter-compartment water exchange”. In: *NeuroImage* 256 (2022), p. 119277. URL: https://www.sciencedirect.com/science/article/pii/S1053811922003986.

[106] Amy FD Howard et al. “The BigMac dataset: an open resource combining multi-contrast MRI and microscopy in the macaque brain”. In: *bioRxiv* (2022). eprint: https://www.biorxiv.org/content/early/2022/09/10/2022.09.08.506363.full.pdf. URL: https://www.biorxiv.org/content/early/2022/09/10/2022.09.08.506363.