

Processing of Diffusion MR Images of the Brain: From Crossing Fibres to Distributed Tractography

Stamatios N. Sotiropoulos

Thesis submitted to the University of Nottingham
for the degree of Doctor of Philosophy

February 2010

Abstract

Diffusion-weighted (DW) magnetic resonance imaging allows the quantification of water diffusion within tissue. Due to the hindrance of water molecules by the various tissue compartments, probing for the diffusive properties of a region can provide information on the underlying structure. This is particularly useful for the human brain, whose anatomy is complex. Diffusion imaging provides currently the only tool to study the brain connectivity and organization non-invasively and in-vivo, through a group of methods, commonly referred to as tractography methods.

This thesis is concerned with brain anatomical connectivity and tractography. The goal is to elucidate problems with existing approaches used to process DW images and propose solutions and methods through new frameworks. These concern data from two popular DW imaging protocols, diffusion tensor imaging (DTI) and high angular resolution diffusion imaging (HARDI), or Q-ball imaging in particular. One of the problems tackled is resolving crossing fibre configurations, a major concern in DW imaging, using data that can be routinely acquired in a clinical setting. The physical constraint of spatial continuity of the diffusion environment is imposed throughout the brain volume, using a multi-tensor model and a regularization method. The new approach is shown to improve tractography results through crossing regions. Quantitative tractography algorithms are also proposed that, apart from reconstructing the white matter tracts, assign relative indices of anatomical connectivity to all regions. A fuzzy algorithm is presented for assessing orientational coherence of neuronal tracts, reflecting the fuzzy nature of medical images. As shown for different tracts, where a-priori anatomical knowledge exists, regions that are coherently connected and possibly belong to the same tract can be differentiated from the background. In a different framework, elements of graph theory are used to develop a new tractography algorithm that can utilize information from multiple image modalities to assess brain connectivity. Both algorithms inherently consider crossing fibre information and are shown to solve problems that affect existing methods.

Acknowledgements

This PhD was part of the Marie Curie collaborative medical image analysis project (CMIAG), funded by the European Union under the FP6 framework.

I would like to thank my supervisor, Dr Christopher Tench, for the excellent guidance and advice throughout the course of my PhD. He has been an endless source of ideas, core research and moral support. I am also thankful to my second supervisor, Dr Li Bai, for her research suggestions and encouragement. Li was the coordinator of the CMIAG research project and I am grateful to her for providing me the opportunity to study in such a great interdisciplinary environment.

I am grateful to Dr Paul Morgan for his invaluable help in acquiring the images used in this thesis. All acquisitions were performed under his ethics protocol. Paul also provided useful feedback and ideas during all the years of my PhD.

I also want to thank Prof. Dorothee Auer for sharing her neuro-radiological expertise in interpreting many of the results of this thesis. Dorothee always provided interesting neuroscience questions and constructive criticism. I am also thankful to Prof. Cris Constantinescu for reviewing most of this work and assessing it from a clinical perspective.

A warm thanks to Prof. Ian Dryden and Dr Alexey Koloydenko for providing feedback during regular meetings. I truly appreciate the time devoted by the other CMIAG faculty members, Prof. Penny Gowland and Dr Alain Pitiot for reviewing the progress of all the CMIAG fellows, including mine. I would also like to thank Prof. Jayaram Udupa for stimulating discussions on his fuzzy segmentation framework, an adaptation of which has been used in this thesis. I am thankful to my examiners, Prof. Daniel Alexander and Prof. Richard Bowtell, for their time and helpful feedback.

I would also like to thank my student mentors, Dr. Peter Steinmetz and Prof. Apostolos Dollas, for helping me obtain the basic qualifications for performing scientific research.

My co-fellows and friends Josef, Jaroslav, Cesar, Antonio, Olivier, Sylvia, Ulas,

Diwei, Yongnan, Amalia and Aaron, deserve a special thanks. They have always been collaborative, willing to help and get scanned.

A huge thanks to Katerina for her support and understanding, especially during the last year of this PhD. Finally, I am grateful to my family that have been unconditionally supportive since the beginning of my scientific career. This thesis is dedicated to them.

Stamatis

Nottingham, 19 Feb. 2010

Publications

Journals

- S. N. Sotiropoulos, L. Bai, P. S. Morgan, C. S. Constantinescu and C. R. Tench. Brain tractography using Q-ball imaging and graph theory: Improved connectivities through fibre crossings via a model-based approach, *NeuroImage*, 49:2444-2456, 2010.
- S. N. Sotiropoulos, L. Bai and C. R. Tench. Fuzzy anatomical connectedness of the brain using single and multiple fibre orientations obtained from diffusion MRI, *Computerized Medical Imaging and Graphics*, in press, 2009.
- S. N. Sotiropoulos, L. Bai, P. S. Morgan, D. P. Auer, C. S. Constantinescu and C. R. Tench. A regularized two-tensor model fit to low angular resolution diffusion images using basis directions, *Journal of Magnetic Resonance Imaging*, 28:199-209, 2008.

Conference Proceedings

- S. N. Sotiropoulos, C. R. Tench, P. S. Morgan and L. Bai. Robust graph-based tracking through crossing fibre configurations, *Proceedings of the IEEE International Symposium on Biomedical Imaging (ISBI)*, pp. 1394-97, Boston, USA, June 2009.
- S. N. Sotiropoulos, L. Bai, P. S. Morgan and C. R. Tench. Graph-based tractography for robust propagation through complex fibre configurations, *Proceedings of the ISMRM Annual Meeting*, p. 852, Honolulu, USA, April 2009.
- S. N. Sotiropoulos, C. R. Tench and L. Bai. In-vivo brain anatomical connectivity using diffusion magnetic resonance imaging and fuzzy connectedness, *Proceedings of the IEEE International Conference on Bioinformatics and Bioengineering (BIBE)*, pp. 1-8, Athens, Greece, October 2008.

-
- S. N. Sotiropoulos, C. R. Tench, P. S. Morgan and L. Bai. Combining Q-ball imaging and fuzzy connectedness: An approach to distributed brain tractography, *Medical Image Understanding and Analysis (MIUA) Proceedings*, pp. 24-28, Dundee, Scotland, July 2008.
 - S. N. Sotiropoulos, C. R. Tench and L. Bai, "Fuzzy anatomical connectedness using diffusion MRI: An approach to tractography of the brain", *Proceedings of the ISMRM Annual Meeting*, p.1846, Toronto, Canada, May 2008.
 - S. N. Sotiropoulos, C. R. Tench, L. Bai, P. S. Morgan, D. P. Auer and C. S. Constantinescu. A regularised two-tensor model fit to low angular resolution diffusion images, *Proceedings of the ISMRM Annual Meeting*, p. 1861, Toronto, Canada, May 2008.

Contents

I	Introduction	1
1	Introduction	2
1.1	Organization of the thesis	5
1.2	Software	8
II	Background	9
2	Diffusion MR Imaging	10
2.1	Self-Diffusion of water molecules	10
2.2	Pulsed gradient spin echo and Q-space	11
2.3	Diffusion imaging and Diffusion tensor imaging	14
	Diffusion Tensor Imaging	16
	Linear estimation of the diffusion tensor	18
	Diffusion tensor properties	19
2.4	Sampling Q-space	22
3	Resolving fibre crossings	24
3.1	DTI and oblate tensors	25
3.2	The multi-tensor model	26
	The "ball and sticks" model	28
	The tensor distribution function	29
3.3	Q-ball imaging and diffusion ODF	30
	ODF for Gaussian diffusion propagator	32

3.4	Spherical harmonics for diffusion ODF estimation	33
	Real spherical harmonics	35
	ODF representation using spherical harmonics	37
3.5	Other HARDI-based approaches	37
	Persistent angular structure and Diffusion orientation transform	37
	Resolving the ADC profile	38
3.6	Spherical deconvolution and fibre ODF	39
	FODF and diffusion ODF	40
3.7	Diffusion spectrum imaging and other q-space approaches	41
	Composite hindered and restricted model of diffusion	42
	Diffusion kurtosis imaging	43
4	White matter tractography	44
4.1	Local approaches	45
	Deterministic streamline tractography	45
	Probabilistic streamline tractography	49
4.2	Global tractography approaches	52
	Front evolution tractography	53
	Graph-based tractography	56
	Global energy minimization tractography	57
4.3	Diffusion simulation tractography	58
III	Original Research	60
5	A regularized two-tensor model for resolving fibre crossings	61
5.1	Introduction	63
5.2	A geometrically constrained two-tensor model	64
5.3	Spatial regularization using relaxation labelling	67

Regularizing the orientation vector field	68
Fitting basis two-tensor models using 1D non-linear optimization	70
5.4 Regularizing the model selection mask	72
Relaxation labelling on c_p	72
Non-linear filtering of c_p	73
5.5 Data and processing	74
Simulations	74
In-vivo images	75
5.6 Results	76
Simulations	76
Quantification of improvement due to regularization	81
In-vivo human data	82
5.7 Discussion	86
Assumptions and potential improvements	88
Summary	90
Appendix A: Generating uniformly distributed directions on a unit sphere . .	91
Appendix B: Simulating the diffusion-weighted signal	91
Simulating tract geometries	92
6 Fuzzy Anatomical Connectedness of the Brain	94
6.1 Introduction	96
6.2 Fuzzy connectedness framework	97
6.3 Fuzzy connectedness tractography	99
Affinity and connectedness based on DTI orientation estimates	100
Increasing angular resolution	102
Incorporating fibre crossing information	103
6.4 Data and processing	104

Simulations	104
In-vivo images	105
6.5 Results	106
Simulations	106
In-vivo DW images	108
6.6 Discussion	117
Comparison with other tractography algorithms	118
Fuzzy connectedness and Dijkstra's algorithm	120
Assumptions and future work	121
Summary	122
Appendix A: Fuzzy connectedness tractography algorithm	123
7 Graphs and multigraphs for white matter tractography	125
7.1 Introduction	127
7.2 Brain tractography using graph theory	128
Incorporating Q-ball ODFs - Problem outline	129
7.3 Multigraph-based tractography	131
7.4 Estimation of multiple single-peaked ODFs within a voxel	133
7.5 Increasing angular resolution	134
7.6 Q-ball fast marching tractography	135
7.7 Data and processing	136
Simulations	136
In-vivo images	136
Thalamic parcellation	137
7.8 Results	138
Simulations	138
In-vivo human data	141

7.9 Discussion	148
Comparison with other tractography methods	149
Assumptions and future work	151
Summary	153
Appendix A: Multigraph-based tractography algorithm	155
Appendix B: Fast-marching tractography algorithm	156
Appendix C: Multiple prolate tensor model	157
 IV Conclusions and Bibliography	 161
 8 Conclusions	 162
8.1 Summary and conclusions	162
8.2 Future perspectives	164
 Bibliography	 167

List of Figures

1.1	The fibrous structure of brain white matter	3
1.2	Sketch of diffusion in grey and white matter	4
2.1	PGSE experiment	12
2.2	DWI images	16
2.3	DTI Ellipsoids	20
2.4	DTI derived images	21
3.1	Oblate diffusion ellipsoids estimated at crossing regions	26
3.2	Multiple tensor models for resolving fibre crossings	29
3.3	DTI and Q-ball ODFs on the same subject	33
3.4	Spherical Coordinate System	34
3.5	Spherical Harmonics up to 4th order	35
3.6	DTI ellipsoids, ODFs and fODFs on simulated data	41
4.1	Major WM tracts resolved with tractography.	45
4.2	Streamline tractography concept	47
4.3	Streamline tractography in corpus callosum	48
4.4	Streamlines and probabilistic index of connectivity	50
4.5	Fast Marching Tractography	54
5.1	Schematic showing the orientations of the geometrically constrained two-tensor model	66
5.2	Schematic showing the spatial regularization of orientations in crossing regions	68
5.3	Fitting L 1D two-tensor models	71

5.4	Orientation estimates using the 1-parameter hybrid approach	77
5.5	Regularizing the model selection mask	78
5.6	Effect of regularization on planar index and orientation estimates	79
5.7	Evolution of label weights during relaxation labelling	80
5.8	Quantification of estimates improvement due to regularization	81
5.9	Regularization effect against the number of diffusion-sensitizing gradient directions	82
5.10	Regularization effect on the c_p image calculated from in-vivo DW images	83
5.11	Orientations resolved from the regularized two-tensor model	84
5.12	Streamline tractography results for the corpus callosum	85
5.13	Streamline tractography results for the corpus callosum	85
5.14	Revised anisotropy and diffusivity maps	86
6.1	Fuzzy connectedness explanatory example	98
6.2	Affinity values for a pair of neighbouring voxels	101
6.3	Different neighbourhood systems for fuzzy connectedness tractography	103
6.4	FCT results on simulated phantoms	107
6.5	Quantitative simulations for FCT	108
6.6	DTI and Q-ball orientations for the corpus callosum	109
6.7	Fuzzy connectedness maps within the corpus callosum	110
6.8	Histogram-based detection of background FC threshold	111
6.9	Fuzzy connectedness maps within the cingulum	112
6.10	Fuzzy connectedness maps for various WM tracts	113
6.11	Path reconstruction using FCT	114
6.12	Reconstruction of paths in the brainstem using FCT	114
6.13	Comparison of FCT with other quantitative tractography methods	115
6.14	FCT combined with the regularized two-tensor model	117

7.1	Simulations illustrating the limitations of graph-based tractography . . .	130
7.2	Graph and multigraph structure in a crossing region	131
7.3	Neighbourhood systems employed for MGT	135
7.4	Results of different tractography algorithms on simulated crossings . . .	139
7.5	Diffusion ODFs utilized by multigraph-based tractography	142
7.6	Graph and multigraph-based tractography within the corpus callosum .	143
7.7	Most strongly connected paths in the corpus callosum	144
7.8	Connection strengths within the cingulum using multigraph-based track- ing	145
7.9	Thalamic parcellation using multigraph-based tractography	146
7.10	Thalamic clusters over different axial slices	146
7.11	Comparison between GT and MGT across different subjects	147
7.12	Tracking the corpus callosum using ODFs and fODFs	149
7.13	Curvature dependence of arc weights	153

List of Tables

7.1 Mean (st. deviation) difference ΔM between path strengths in horizontal bundle A and crossing bundle B , for various algorithms and crossing angles, across 100 simulations at SNR=40. 140

7.2 Mean (st. deviation) difference ΔM between path strengths in horizontal bundle A and crossing bundle B , for various algorithms and crossing angles, across 100 simulations at SNR=20. 141

List of Algorithms

6.1	Fuzzy Connectedness Tractography with DTI fibre orientations	102
6.2	Update ($\mathbf{j}, k, \mathbf{i}, l, \text{val}$)	123
6.3	Fuzzy Connectedness Tractography	124
7.1	Multigraph-based Tractography	155
7.2	Valid ($\mathbf{i}, n, \mathbf{F}_i^{\text{neigh}}$)	156
7.3	Fast Marching Tractography	157

Part I

Introduction

1

Introduction

The human brain has been described as *"one of the most complex systems in the universe"* (Koch and Laurent, 1999). It is a network of more than 10^{11} individual nerve cells and 10^{16} interconnections (Kandel et al., 2000). Due to the brain's complexity, many aspects of its function remain uncertain. According to the concept of *"structure defining function"* (Buzsaki, 2006), the safest way to start speculating about the brain's function is to inspect its anatomical organization and understand the underlying circuitry. The pattern of structural connections can be used to place specific constraints on brain dynamics and help to identify causal interactions in brain networks (Sporns et al., 2005).

Towards this direction, brain anatomy has been studied for over a century now (Dejerine and Dejerine-Klumpke, 1895; Gray, 1918). Neural tracers and histology, for instance (LaVail and LaVail, 1972), allowed neuroanatomical studies, but their invasive nature has limited their application to animal or post-mortem human brains (Brodmann, 1909; Dejerine and Dejerine-Klumpke, 1895). Dissections (Ludwig and Klingler, 1956) offered another alternative to study anatomy (Haines, 2004; Williams et al., 1997), but still in post-mortem tissue. The development of imaging techniques and of magnetic resonance imaging (MRI) (Lauterbur, 1973; Mansfield and Maudsley, 1977), in

particular, has allowed the study of brain structure in living humans. Due to the limited spatial resolution of these techniques, neuroscience and neurology questions can be tackled at a systems level and on a macroscopic scale. Diffusion-weighted MRI is one of these methods and will be the subject of this thesis.

Diffusion-weighted MRI (DW-MRI) was first applied to living humans in 1986 by Le Bihan et al (Le Bihan et al., 1986). Since then, there has been a lot of research on the field and great potential has been revealed. The underlying concept of the method is that the phenomenon of diffusion of water molecules within tissue is utilized to introduce contrast to images. Diffusion is an intrinsic molecular transport process that induces thermally-driven random motions. But what anatomical information can we infer using DW-MRI?

At the microscopic scale, each neuronal cell comprises of three main parts, the dendrites, the cell body (soma) and the axon (Kandel et al., 2000). These processes grossly correspond to the input, processing and output units of a cell. At a macroscopic scale, cell bodies and dendrites of millions of neurons tend to cluster together to constitute what is known as the *grey matter* of the brain (Kandel et al., 2000). Axons, which can be very long, also cluster together to form axonal bundles and constitute the *white matter*. Axonal bundles, also known as fasciculi or tracts, appear to have a very fibrous and coherent structure in a brain dissection (Fig. 1.1) and therefore axons are also called fibres.



Figure 1.1: The fibrous structure of brain white matter, as shown in post-mortem dissections. Figure adapted from (Williams et al., 1997).

White matter tracts are the wires of the brain. They transmit information between

different brain regions and define the direct interconnections between nodes of the brain network. DW-MRI is currently the only tool that allows study and reconstruction of these tracts, non-invasively and in-vivo, via a process known as tractography (Basser et al., 2000; Mori et al., 1999, 2005). The assumption that is made in DW-MRI is that the observed hindrance of water diffusive motions is connected to the underlying tissue structure. Within tissue, diffusing water molecules are hindered by the various microstructures, such as cell membranes, axons and myelin sheaths that act as barriers to motion. When tissue structural coherence exists, this hindrance is systematic and leads to a measurable preferred diffusion orientation. As shown in the simple sketches of Fig. 1.2, in grey matter, diffusion will be more restricted compared to a barrier-free medium, however no particular orientational preference will be exhibited. On the contrary, white matter exhibits structural coherence and water molecules will diffuse along rather than across the fibres (Beaulieu, 2002). Therefore, the dominant diffusion orientation in white matter is assumed to give an estimate of the orientation of the underlying neuronal bundles. Utilizing these orientations across the whole brain, tractography can be performed (Basser et al., 2000; Mori et al., 1999).




Figure 1.2: Sketch of barrier-free diffusion and diffusion in grey (GM) and white matter (WM). The trajectories of three diffusing molecules are shown. Cell bodies are represented as spheres and axons as cylinders.

Reconstruction of white matter tracts is important for tackling neuroscience/ neuroanatomy questions (Hagmann et al., 2008; van den Heuvel et al., 2009). A map of structural connections, obtained in-vivo, can be used to impose constraints on the analysis of functional connectivity maps. The latter, estimated for instance using electroencephalography (EEG) or functional MRI (fMRI), are only available for living subjects

and DW-MRI allows their integration with structural information of the same subject to provide insight into different aspects of brain's function.

Tractography is also important for studying neurological disorders (Ciccarelli et al., 2008). Norman Geschwind argued in 1965 that interruption of normal white matter connections is likely to lead to a range of impairments (Geschwind, 1965a,b). This disconnection hypothesis suggested that different diseases may affect different tracts. Localization of specific white matter tracts and study of structural changes along the bundles can provide insight to the disease mechanisms and aetiology, as well as offer additional diagnostic information. The potential of tractography to detect pathology-induced changes along specific bundles has been illustrated in many studies (Ciccarelli et al., 2008). Neurosurgical planning has also benefited from DW-MRI (Clark et al., 2003), since mapping intact white matter pathways can assist the surgeon to decide the best strategy and avoid resecting clinically eloquent tissue.

Despite the great potential, white matter tractography is relatively immature. This thesis aims to take steps for improving existing tractography methods. Many of the approaches that will be described in the background sections were developed simultaneously with the work presented here. The goal of this thesis was to elucidate problems with existing algorithms and propose solutions through new frameworks. One of the problems tackled is resolving crossing fibre configurations, a main concern in DW imaging, using, though, data that can be routinely acquired in a clinical setting. Quantitative tractography algorithms are also proposed that, apart from reconstructing the white matter tracts, assign to connected regions relative indices of connectivity. These algorithms inherently consider crossing fibre information and solve problems that existing methods have.

1.1 Organization of the thesis

This thesis is organized in seven chapters. The first three give an overview of existing methods in diffusion-weighted (DW) imaging and techniques to post-process these data in order to reconstruct white matter tracts. The next three chapters present the methodological contributions of this thesis. Each of the proposed methods are applied to human brain DW images and compared with other techniques in the literature. The

last chapter summarizes the results.

Chapter 2 describes diffusion as a molecular transport process and presents the simplest nuclear magnetic resonance (NMR) experiment for weighting the measured signal by the random diffusive motions of water molecules. The *q-space* formalism is explained for estimating the diffusion profile from diffusion-weighted (DW) NMR measurements. The commonly used *diffusion tensor imaging* (DTI) is then derived from the *q-space* relation and presented. Finally, a brief description of acquisition protocols for sampling *q-space* is given. Spherical, Cartesian and hybrid protocols are discussed.

Chapter 3 gives an overview of methods that tackle a major problem in diffusion imaging, which is *resolving crossing fibre configurations* from DW data. What happens when two (or more) distinct fibre tracts cross within an imaging voxel? The diffusion profile becomes complex as there are contributions from different compartments. Model-based and model-free techniques that try to resolve such complex profiles are reviewed. These range from multi-tensor models, methods that resolve a spherical feature of the diffusion profile (orientation distribution functions and persistent angular structure for instance), spherical deconvolution approaches and *q-space* methods.

Chapter 4 illustrates how the resolved diffusion profile can be utilized to reconstruct white matter tracts, a process known as *white matter tractography*. Different methods are presented ranging from curve propagation techniques to front evolution and diffusion simulation approaches. Non-quantitative and quantitative methods are presented, the former providing white matter tracts only and the latter estimating an index of "connectivity", which characterizes how plausible a putative tract is, given the data. A further subdivision is performed between local and global approaches. Local approaches reconstruct tracts by using the diffusion profile of a voxel, in a greedy way, to propagate to subsequent voxels. Global methods consider a feature along the whole path and look for candidate tracts that are optimum according to this feature.

Chapter 5 presents a new method for better resolving crossing fibre configurations from conventional diffusion-weighted images. Diffusion scans performed routinely in a clinical setting must be acquired within a short time frame and therefore provide a very sparse sampling of *q-space* within each voxel. If one is interested in resolving fibre crossings, utilization of advanced techniques is problematic, due to the small number of acquired samples. To overcome this difficulty, a *spatial regularization method* is in-

troduced to counterbalance the sparse intra-voxel sampling with information obtained from neighbouring voxels. The physical constraint of diffusion continuity across neighbouring voxels is imposed using relaxation labelling, an iterative method, on the estimates of a simplified two-tensor model. The regularized approach is shown to reliably resolve crossing configurations in regions where a priori anatomical knowledge exists. Tractography results verify the improved performance compared to other methods.

Chapter 6 presents a new *tractography approach based on a fuzzy framework*. Medical images are by nature fuzzy, due to various factors, for instance noise artefacts, hardware limitations and limited spatial resolution. A fuzzy algorithm is therefore described to reconstruct white matter tracts. Fuzzy transition weights are defined between neighbouring voxels and tracts are reconstructed on the discrete image grid. The proposed method is a global tractography approach that assigns to each voxel a relative connectivity value, representative of the weakest link of the best path connecting it to a predefined seed. It is distributed in the sense that all image voxels are assigned a relative connectivity value. The algorithm considers any number of fibre compartments coexisting in a voxel, thus allowing the incorporation of fibre crossing information. Results are presented for many known white matter tracts that have been successfully reconstructed.

Chapter 7 describes a novel and flexible method to perform *tractography using graph theory*. Building on the framework of the previous chapter, a multigraph representation of the image is generated. To achieve that, a Q-ball reconstruction of the data is performed along with fitting multi-tensor models. Tracts are then reconstructed using information from different imaging modalities; diffusion information derived from DW images and structural information from anatomical images. The multigraph-based tractography is compared with three other methods to show its advantages. It utilizes fibre crossing information and propagates robustly through crossing regions. To illustrate its potential, connectivity-based parcellation of the thalamus is performed.

Chapter 8 summarizes the results presented in chapters 5-7 and discusses future perspectives.

1.2 Software

The new algorithms and methods presented in this thesis were implemented using the C programming language (Kernighan and Ritchie, 1988). For visualization of the results, Matlab 7.4 (The Mathworks, MA, USA) and Paraview 3.6 (<http://www.paraview.org>) were used. UCL's diffusion MRI toolkit - Camino 2 (Cook et al., 2006) was used to perform deterministic and probabilistic streamline tractography, while Oxford's software library - FSL 4.1 (Smith et al., 2004) was utilized to perform brain extraction and image registration.

Part II

Background

2

Diffusion MR Imaging

2.1 Self-Diffusion of water molecules

Diffusion is a molecular transport process that involves thermally-driven random motions (Hobbie, 1997). This process, occurring constantly at non-zero ambient absolute temperatures, was first discovered in 1827 by Robert Brown and is also known as Brownian motion. The random diffusion displacements in an ensemble of molecules can be described using the Einstein relation. In a homogeneous, barrier-free medium, provided that the number of molecules is sufficiently large, the squared diffusion displacements \mathbf{r} , averaged over all the molecules in the ensemble, is directly proportional to the observation time τ . In n dimensional space:

$$\langle \mathbf{r}^2 \rangle = 2nD\tau, \quad n = 1, 2, 3 \quad (2.1)$$

with $\langle \mathbf{r}^2 \rangle$ the mean squared displacement. The above can be seen as a Gaussian process with zero mean displacement and variance $\sigma^2 \propto D\tau$. The scalar proportionality constant D is known as the diffusion coefficient and depends on the medium viscosity,

the particle size and the temperature (Hobbie, 1997). When both the diffusing particles and the medium in which they diffuse are of the same species, the process is called self-diffusion and D the self-diffusion coefficient. At body temperature (37°C), the self-diffusion coefficient of water is $3 \times 10^{-3} \text{ mm}^2/\text{s}$ (Jones, 2009).

The diffusion process can be also described by Fick's laws (Hobbie, 1997). Fick's first law states that diffusion appears macroscopically as a net flux from regions of high particle concentration C to regions of low concentration. The second law is of particular interest as it describes the spatio-temporal dynamics of diffusion. Also known as the Heat equation or the Diffusion equation, it states that in a homogeneous medium:

$$\frac{dC}{dt} = D \cdot \nabla^2 C . \quad (2.2)$$

For self-diffusion, the concentration C can be replaced by the conditional probability density function $P(\mathbf{r}, \tau)$ (Callaghan, 1991) that describes the average probability for any particle in an ensemble having a relative diffusion displacement \mathbf{r} after time τ . Assuming a barrier-free medium and that all particles are concentrated at the same location at time $\tau = 0$, the solution of the diffusion equation gives a trivariate Gaussian distribution (Callaghan, 1991) with zero mean and variance $\sigma^2 = 2D\tau$ in each direction, similar to the predictions of Einstein's relation:

$$P(\mathbf{r}, \tau) = \frac{1}{\sqrt{(4\pi D\tau)^3}} \cdot \exp\left(\frac{-\mathbf{r}^2}{4D\tau}\right) = \frac{1}{\sqrt{(4\pi D\tau)^3}} \cdot \exp\left(\frac{-\mathbf{r}^T \cdot \mathbf{r} \cdot D^{-1}}{4\tau}\right) . \quad (2.3)$$

2.2 Pulsed gradient spin echo and Q-space

Diffusion-weighted (DW) MR imaging is sensitive to the self-diffusion of water molecules within tissues. It can indirectly provide structural tissue information, by allowing estimation of the water diffusion profile in different tissue regions. A way to introduce diffusion-weighting to MR sequences along with the q-space formalism (Callaghan, 1991) are explained in this section.

Let a spin (hydrogen proton), at time $t=0$, be at location \mathbf{x}_0 . The diffusion profile of this spin can be described by a conditional probability density function (pdf) $P_s(\mathbf{x}|\mathbf{x}_0, \tau)$. For a single spin, this pdf gives the probability of diffusion from \mathbf{x}_0 to

\mathbf{x} in time τ . As mentioned before, for an ensemble of spins, we can define a pdf $P(\mathbf{r}, \tau) = P(-\mathbf{r}, \tau) = \int P_s(\mathbf{x}_0 + \mathbf{r}|\mathbf{x}_0, \tau) \cdot \rho(\mathbf{x}_0) d\mathbf{x}_0$, with ρ giving the spin density at a specific location and \mathbf{r} being the relative spin displacement (Callaghan, 1991). The function P gives the average probability for any particle in the ensemble to have a diffusion displacement \mathbf{r} and is encountered in the literature under various names, such as the diffusion propagator, the diffusion spectrum or the diffusion scatter pattern.




Figure 2.1: Pulsed gradient spin echo experiment used to introduce diffusion-weighting to MR measurements.

In the q-space formalism (Callaghan, 1991), P can be calculated as the Fourier transform of diffusion-weighted MR measurements. Diffusion-weighting along a given direction is commonly achieved using the pulsed gradient spin echo (PGSE) sequence introduced by Stejskal and Tanner (Stejskal and Tanner, 1965). PGSE is a modified spin echo experiment that utilizes diffusion-sensitizing magnetic field gradients applied along the direction of interest (Fig. (2.1)). Let $\mathbf{G}(t)=[G_x(t) \ G_y(t) \ G_z(t)]^T$ such a gradient and γ the gyromagnetic ratio (for protons $\gamma=42.576$ MHz/T). Then, a spin located at position \mathbf{x} along the direction of \mathbf{G} will experience a Larmor frequency offset $\omega_G(\mathbf{x}, t) = \gamma \cdot \mathbf{x} \cdot \mathbf{G}(t)$ (Haacke et al., 1999). Thus, the accumulated phase during the gradient application will be:

$$\phi_G(\mathbf{x}) = - \int_0^\infty \omega_G(\mathbf{x}, t) dt = -\gamma \cdot \mathbf{x} \cdot \int_0^\infty \mathbf{G}(t) dt . \quad (2.4)$$

Let \mathbf{q} the vector

$$\mathbf{q} = \gamma \int_0^{\infty} \mathbf{G}(t) dt, \quad (2.5)$$

which has the same orientation as the gradient vector \mathbf{G} (in the case of a constant gradient of duration δ , the \mathbf{q} vector reduces to $\mathbf{q} = \gamma \mathbf{G} \delta$). Then, the phase acquired by a spin at location \mathbf{x} will be $-\mathbf{q} \cdot \mathbf{x}$, according to Eq. (2.4) and (2.5). During the PGSE experiment, two identical gradients are applied on either side of a 180° RF pulse, as shown in Fig. 2.1. If the gradients are short enough ($\delta \ll \Delta$) so that no diffusion displacements take place during their duration, a spin at location \mathbf{x}_0 will acquire a phase $-\mathbf{q} \cdot \mathbf{x}_0$ during the application of the first gradient. If the spin is at location \mathbf{x}_1 during the application of the second gradient, it will acquire a phase $-\mathbf{q} \cdot \mathbf{x}_1$. The 180° pulse reverses the phase change that occurred prior to its application, i.e. the one induced by the first gradient. Therefore, the net phase for this spin will be $-\mathbf{q} \cdot (\mathbf{x}_1 - \mathbf{x}_0) = -\mathbf{q} \cdot \mathbf{r}$, \mathbf{r} being the displacement of the spin position between times 0 and Δ . In the case of a stationary spin, the two phases cancel each other and complete rephasing occurs. If, however, a particle moves during the inter-gradient interval, then a phase will be acquired. Different random displacements \mathbf{r} by different particles will lead to phase dispersion and an attenuation of the observed signal $S(\mathbf{q}, \tau)$ will be observed, compared to the signal $S(0)$ that would be obtained without the application of diffusion-sensitizing gradients.

The measured signal attenuation $E(\mathbf{q}, \tau) = S(\mathbf{q}, \tau) / S(0)$ can be obtained by integrating over the induced phases obtained by all possible diffusion displacements across an ensemble of spins and weighting the displacements by the diffusion propagator P . A Fourier transform between E and P results (Callaghan, 1991):

$$E(\mathbf{q}, \tau) = \int P(\mathbf{r}, \tau) \cdot e^{-i\mathbf{q} \cdot \mathbf{r}} d\mathbf{r} = \mathcal{F}[P(\mathbf{r}, \tau)]. \quad (2.6)$$

Using this Fourier relationship the pdf P can be estimated from q-space measurements. The wavevector \mathbf{q} and the diffusion displacement \mathbf{r} define a Fourier pair, just as the \mathbf{k} vector and the spatial location \mathbf{R} in MR imaging. Remember that for an imaging gradient $\mathbf{G}_i(t)$, the \mathbf{k} vector and the signal S are (Haacke et al., 1999):

$$\mathbf{k} = \frac{1}{2\pi} \gamma \int_0^{\infty} \mathbf{G}_i(t) dt \quad \text{and} \quad S(\mathbf{k}) = \int \rho(\mathbf{R}) \cdot e^{-i2\pi \mathbf{k} \cdot \mathbf{R}} d\mathbf{R}.$$

Therefore, \mathbf{q} and \mathbf{r} are reciprocal in the sense that measurements at high "frequencies" \mathbf{q} will be more sensitive to subtle diffusion displacements, while low \mathbf{q} -value measurements capture larger and coarser displacements (Basser, 2002). In (Wedeen et al., 2005) it is further proved that the diffusion MR signal is real and positive. This means that $\mathcal{F}[E(\mathbf{q})] = \mathcal{F}[|E(\mathbf{q})|]$, i.e. only the modulus of the measurements is necessary to compute the diffusion propagator:

$$|E(\mathbf{q}, \tau)| = \int P(\mathbf{r}, \tau) \cdot e^{-i\mathbf{q} \cdot \mathbf{r}} d\mathbf{r} = \mathcal{F}[P(\mathbf{r}, \tau)] . \quad (2.7)$$

The derivation of the Fourier transform of Eq. (2.6) assumes that the width of the diffusion-sensitizing gradients is negligible compared to the diffusion time of interest (i.e. $\delta \ll \tau$ and $\Delta \approx \tau$). In practice, this narrow pulse approximation is not valid, as δ and Δ are of a similar magnitude. Under certain assumptions, the effects of the non-negligible δ can be modelled (a discussion is included in (Alexander, 2006)). For example, assuming a Gaussian diffusion propagator and rectangular gradients, the non-negligible δ reduces the effective diffusion time to $\tau = \Delta - \delta/3$ (Stejskal and Tanner, 1965). In (Mitra and Halperin, 1995), it is further shown that the Fourier transform between the measured signal and the diffusion propagator still holds for a modified interpretation of the displacement vector \mathbf{r} . The vector \mathbf{r} can be considered as the displacement of the mean spin position within the period $[0, \delta]$ relative to the mean spin position within $[\Delta, \delta + \Delta]$.

2.3 Diffusion imaging and Diffusion tensor imaging

The simplest diffusion imaging experiment assumes free diffusion and a homogeneous diffusion coefficient D in each voxel. Then, using the \mathbf{q} -space relationship and the Gaussian diffusion propagator described in Eq. (2.3), the signal obtained from a PGSE experiment follows the equation:

$$S(\mathbf{q}, \tau) = S(0) \cdot \exp(-|\mathbf{q}|^2 D \tau) . \quad (2.8)$$

This expression was obtained by Stejskal and Tanner (Stejskal and Tanner, 1965),

who found that $S(\mathbf{q}, \tau) = S(0) \cdot \exp(-bD)$. The value $b = \tau |\mathbf{q}|^2$ effectively determines the amount of diffusion weighting applied during acquisition. For rectangular gradients of finite duration δ and under the assumption of Gaussian diffusion, the effective diffusion time reduces to $\tau = \Delta - \delta/3$, as mentioned in the previous section. Then $\mathbf{q} = \gamma \mathbf{G} \delta$ gives $b = \gamma^2 G^2 \delta^2 (\Delta - \delta/3)$. Le Bihan et al (Le Bihan et al., 1986) incorporated the PGSE experiment within an MRI sequence to obtain for the first time in 1986 diffusion-weighted images of the human brain. A single diffusion coefficient was then estimated in each image voxel using Eq. (2.8). As pointed out in (Le Bihan et al., 1986), these coefficients do not purely reflect free diffusion and are thus smaller than the theoretical self-diffusion coefficient of water. Diffusion in the brain is hindered by the various tissue compartments. If enough diffusion time is allowed, particles will hit the boundaries and diffuse less than expected in a free medium. Furthermore, any type of incoherent motion that can cause spin dephasing during the application of the gradients will contribute to the measured echo attenuation. These include, apart from diffusion, microcirculation in the capillary network and fluids flowing with several velocities within tissue. For all these reasons, the estimated quantities were termed *apparent diffusion coefficients* (ADCs).

Bulk motions, such as cardiac cycle pulsations and involuntary head movements, can also be mistakenly perceived as diffusion and contribute to the measured signal. To reduce sensitivity to bulk motions, DW echo-planar imaging (EPI) was introduced by Turner et al (Turner et al., 1990). With single-shot EPI an entire image is captured in less than 0.1 seconds and motion artefacts are reduced.

The assumption of isotropic behaviour of the diffusion coefficient is utilized in DW imaging. However, this is not always the case for brain tissue. Fig. 2.2 shows DW images of the same brain region when gradients with identical magnitude but different direction are applied. The signal attenuation in certain regions varies significantly with gradient direction and this anisotropic behaviour will be reflected in the estimated coefficients. More specifically, anisotropic diffusion is encountered in regions with coherent tissue structure, such as brain white matter (WM) (Beaulieu, 2002). In these regions, diffusion displacements are on average larger along a certain "preferred" orientation. Water molecules, hindered by the axonal membranes and the myelin sheaths, diffuse along rather than across neuronal axons. For regions with no tissue structure, such as

CSF-filled areas, or regions with random structure, such as brain grey matter (GM) the assumption of isotropy seems more appropriate, i.e. displacements are expected to be of the same magnitude along any orientation.




Figure 2.2: Single-shot EPI diffusion-weighted images acquired with gradients applied along three different orientations (anterior-posterior, left-right, superior-inferior). An EPI image without diffusion-weighting is shown on the left. To aid visualization, the dynamic range of the non-DW image is different, therefore the signal attenuation is not obvious in all regions.

Diffusion Tensor Imaging

To capture potential anisotropic diffusion, diffusion tensor imaging (DTI) has been introduced (Basser et al., 1994a). In DTI, the diffusion propagator is assumed to be a zero-mean tri-variate Gaussian function (Basser et al., 1994b):

$$P(\mathbf{r}, \tau) = \frac{1}{\sqrt{(4\pi\tau)^3 |\mathcal{D}|}} \cdot \exp\left(\frac{-\mathbf{r}^T \cdot \mathcal{D}^{-1} \cdot \mathbf{r}}{4\tau}\right), \quad (2.9)$$

with a covariance matrix $\Sigma = 2\tau\mathcal{D}$ and \mathcal{D} called the diffusion tensor. This is a generalization of the diffusion propagator of Eq. (2.3) obtained for a barrier-free medium, however with the assumption of medium homogeneity removed. Thus, instead of a single diffusion coefficient D in all dimensions, different coefficients are allowed along different directions. This is reflected in the diffusion tensor \mathcal{D} , which is a 3x3 symmetric, positive definite matrix:

$$\mathcal{D} = \begin{bmatrix} D_{xx} & D_{xy} & D_{xz} \\ D_{xy} & D_{yy} & D_{yz} \\ D_{xz} & D_{yz} & D_{zz} \end{bmatrix}.$$

Using the q-space relationship of Eq. (2.6), the DTI model for the signal attenuation (Basser et al., 1994a) can be obtained. We need to use the following Fourier transform pair, according to which the Fourier of a multivariate Gaussian function with zero mean and covariance Σ is also a multivariate Gaussian with zero mean and covariance Σ^{-1} (Pollard, 2001):

$$\frac{1}{\sqrt{(2\pi)^3 |\Sigma|}} \cdot \exp\left(\frac{-\mathbf{r}^T \cdot \Sigma^{-1} \cdot \mathbf{r}}{2}\right) \xleftrightarrow{\mathcal{F}} \exp\left(\frac{-\mathbf{q}^T \cdot \Sigma \cdot \mathbf{q}}{2}\right).$$

Then, if \mathbf{g} is a unit vector such that $\mathbf{q} = |\mathbf{q}| \mathbf{g}$ and $b = \tau |\mathbf{q}|^2$ as before, the diffusion MR signal will be:

$$\begin{aligned} E(\mathbf{q}, \tau) = \exp(-\tau \cdot \mathbf{q}^T \cdot \mathcal{D} \cdot \mathbf{q}) &\implies E(\mathbf{g}, \tau) = \exp(-b \cdot \mathbf{g}^T \cdot \mathcal{D} \cdot \mathbf{g}) \implies \\ S(\mathbf{g}, \tau) &= S(0) \cdot \exp(-b \cdot \mathbf{g}^T \cdot \mathcal{D} \cdot \mathbf{g}). \end{aligned} \quad (2.10)$$

As in DW imaging, for rectangular magnetic field gradients \mathbf{G} of duration δ , the effective diffusion time is $\tau = \Delta - \delta/3$ (Basser et al., 1994a) and the b value is $b = \gamma^2 G^2 \delta^2 (\Delta - \delta/3)$.

Using the above equation, the tensor \mathcal{D} can be estimated from the diffusion measurements S . Having 6 unknowns, at least 6 unique DW measurements are needed to determine \mathcal{D} , provided that $S(0)$ is measured. Keeping the scalar b constant at a high value (around 1000 s/mm²) (Jones et al., 1999) and changing the direction of the diffusion gradients (i.e. sampling a sphere with constant radius in q-space) will give the required measurements. To improve immunity against noise, more than 6 non-collinear directions (usually 15-30) are utilized (Jones, 2004). To reduce any noise bias and keep the statistical properties of the calculated parameters as rotationally invariant as possible, the sampling directions should be evenly distributed on the sphere (Hasan et al., 2001; Jones, 2004). Having a set of DW images, the diffusion tensor can be estimated by performing multivariate regression on Eq. (2.10) (Basser et al., 1994a).

Linear estimation of the diffusion tensor

Given M DW measurements per voxel, we need to estimate the 6 unknown elements of the diffusion tensor \mathcal{D} . According to the DTI model, for $k = 1 \dots M$:

$$\begin{aligned} S(\mathbf{g}_k, \tau) &\equiv S_k = S(0) \cdot \exp(-b \cdot \mathbf{g}_k^T \cdot \mathcal{D} \cdot \mathbf{g}_k) \implies \\ C_k &\equiv -\ln\left(\frac{S_k}{S(0)}\right) = b \mathbf{g}_k^T \cdot \mathcal{D} \cdot \mathbf{g}_k \implies \\ C_k &= b g_{k1}^2 D_{xx} + b g_{k2}^2 D_{yy} + b g_{k3}^2 D_{zz} + 2b g_{k1} g_{k2} D_{xy} + 2b g_{k1} g_{k3} D_{xz} + 2b g_{k2} g_{k3} D_{yz} . \end{aligned}$$

This is a system of M linear equations with 6 unknowns, the elements of the symmetric diffusion tensor. In matrix form, this can be written as:

$$\begin{bmatrix} C_1 \\ C_2 \\ \vdots \\ C_M \end{bmatrix} = \begin{bmatrix} b g_{11}^2 & b g_{12}^2 & b g_{13}^2 & 2b g_{11} g_{12} & 2b g_{11} g_{13} & 2b g_{12} g_{13} \\ b g_{21}^2 & b g_{22}^2 & b g_{23}^2 & 2b g_{21} g_{22} & 2b g_{21} g_{23} & 2b g_{22} g_{23} \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ b g_{M1}^2 & b g_{M2}^2 & b g_{M3}^2 & 2b g_{M1} g_{M2} & 2b g_{M1} g_{M3} & 2b g_{M2} g_{M3} \end{bmatrix} \cdot \begin{bmatrix} D_{xx} \\ D_{yy} \\ D_{zz} \\ D_{xy} \\ D_{xz} \\ D_{yz} \end{bmatrix} \\ \implies \mathbf{C} = \mathbf{A} \cdot \mathbf{X} .$$

The solution to this system for $M > 6$ is (Press et al., 1992):

$$\mathbf{X} = (\mathbf{A}^T \cdot \mathbf{A})^{-1} \mathbf{A}^T \cdot \mathbf{C} \quad (2.11)$$

and is known as the ordinary least squares (OLS) estimate. The assumption of this approach is that noise is zero-mean Gaussian and that all measurements have the same noise variance. For fitting to the diffusion-weighted signal S_k , this is a reasonable assumption, unless the signal to noise ratio (SNR) is very low (Pajevic and Basser, 2003). However, the model is fitted to the logarithm of the signal C_k . The logarithmic transformation increases the variance of small measurements (note that as $x \rightarrow 0$, $\ln x \rightarrow -\infty$ rapidly) relative to the variance of large measurements. As suggested by (Salvador et al., 2005), noise for the logarithm of the measurements follows approximately a zero-mean Gaussian distribution, for $\text{SNR} > 3$. However, the variance across measurements

differs, with the logarithm of the S_k measurement having a variance inversely proportional to S_k^2 .

Therefore, a more appropriate approach would be weighted least squares (WLS) (Basser et al., 1994a; Kingsley, 2006b). Each equation of the above linear system is weighted by a number representative of the reciprocal variance of the corresponding measurement. Thus, measurements with a large variance will have a lower impact in the estimation process. The WLS solution is:

$$\mathbf{W} \cdot \mathbf{C} = \mathbf{W} \cdot \mathbf{A} \cdot \mathbf{X} \implies \mathbf{X} = (\mathbf{A}^T \cdot \mathbf{W} \cdot \mathbf{A})^{-1} \mathbf{A}^T \cdot \mathbf{W} \cdot \mathbf{C} . \quad (2.12)$$

\mathbf{W} is a $M \times M$ diagonal matrix with element w_{kk} the weight corresponding to the k^{th} measurement. Due to the logarithmic transformation, a good choice for the weights is $w_{kk} \propto S_k^2$ (Salvador et al., 2005). It should be pointed out that scaling all the weights with the same number will not make any difference in the estimates.

Diffusion tensor properties

Once \mathcal{D} is determined, its diagonal elements represent ADCs along the three laboratory coordinate directions (slice selection, readout, phase encoding direction), while the off-diagonal elements represent the correlation of displacements along these coordinate directions (Pierpaoli et al., 1996). However, the ADCs along the directions of a local coordinate system, determined by the tissue anatomy and structure, are of interest. Therefore, the diffusion tensor is commonly diagonalized (Press et al., 1992) and its eigenvalues $\lambda_1, \lambda_2, \lambda_3$ and eigenvectors $\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3$ are calculated (Kingsley, 2006b). The eigen-decomposition of the diffusion tensor is:

$$\mathcal{D} = [\mathbf{e}_1 | \mathbf{e}_2 | \mathbf{e}_3] \begin{bmatrix} \lambda_1 & 0 & 0 \\ 0 & \lambda_2 & 0 \\ 0 & 0 & \lambda_3 \end{bmatrix} [\mathbf{e}_1 | \mathbf{e}_2 | \mathbf{e}_3]^T = \lambda_1 \mathbf{e}_1 \mathbf{e}_1^T + \lambda_2 \mathbf{e}_2 \mathbf{e}_2^T + \lambda_3 \mathbf{e}_3 \mathbf{e}_3^T . \quad (2.13)$$

The eigenvalues of the diffusion tensor provide diffusion coefficients along the orientations defined by its respective eigenvectors (Pierpaoli et al., 1996). A common representation of iso-probability surfaces of the Gaussian diffusion propagator estimated

by DTI is the diffusion ellipsoid (Basser, 1995), with axes along the three eigenvectors \mathbf{e}_i and magnitude equal to the respective root mean square (rms) diffusion displacements $\sqrt{2\tau\lambda_i}$. The shape of the ellipsoid reveals the anisotropy degree of the tensor, as shown in Fig. 2.3. For anisotropic ellipsoids found in WM, the principal eigenvector \mathbf{e}_1 , i.e. the one associated with the largest eigenvalue λ_1 , has been found to coincide with the principal fibre orientation of the underlying tissue structure (Pierpaoli et al., 1996).




Figure 2.3: DTI ellipsoids for an isotropic (left) and an anisotropic (right) tensor.

While the diffusion ellipsoid effectively plots the rms diffusion displacements, the ADC profile (Basser, 1995) is an angular plot of the diffusion coefficients that one would measure along different directions, given a tensor \mathcal{D} . The ADC value along any arbitrary direction \mathbf{g} is:

$$ADC_{\mathbf{g}} = \mathbf{g}^T \cdot \mathcal{D} \cdot \mathbf{g} = \lambda_1(\mathbf{g} \cdot \mathbf{e}_1)^2 + \lambda_2(\mathbf{g} \cdot \mathbf{e}_2)^2 + \lambda_3(\mathbf{g} \cdot \mathbf{e}_3)^2. \quad (2.14)$$

Note that the above expression is equal to $-\ln[S(\mathbf{g}, \tau)/S(0)]/b$. Plotting these values along many directions on the sphere will give the ADC profile, also known as an ADC peanut, due to its shape in highly anisotropic regions.

A number of scalar rotationally invariant parameters can be extracted from the diffusion tensor. The Mean Diffusivity (MD) can be obtained as the mean of the tensor's eigenvalues or in general as the average sum of three ADCs measured in three

orthogonal directions (Basser, 1995):

$$MD = \frac{D_{xx} + D_{yy} + D_{zz}}{3} = \frac{\lambda_1 + \lambda_2 + \lambda_3}{3}. \quad (2.15)$$

The Fractional Anisotropy (FA) is the most commonly used anisotropy measure and is a normalized expression of the variance of the tensor eigenvalues (Basser, 1995). It is 0 for perfectly isotropic ($\lambda_1 = \lambda_2 = \lambda_3$) and 1 for perfectly anisotropic tensors ($\lambda_1 \neq 0, \lambda_2 = \lambda_3 = 0$):

$$FA = \sqrt{\frac{3 \sum_{i=1}^3 (\lambda_i - \bar{\lambda})^2}{2 \sum_{i=1}^3 \lambda_i^2}}. \quad (2.16)$$




Figure 2.4: Diffusion tensor derived images, including the mean diffusivity, the fractional anisotropy and the principal tensor eigenvector. The latter is colour-coded by orientation. As shown by the coloured sphere, superior-inferior (S-I) orientation corresponds to blue, anterior-posterior (A-P) to green and medial-lateral (M-L) to red.

Axial slices of the MD and FA maps are shown in Fig. 2.4. The mean diffusivity is relatively constant within the brain parenchyma at $0.7 \times 10^{-3} \text{ mm}^2/\text{s}$. Interestingly, this value remains roughly the same across human subjects and also across a range of other mammalian brains (Basser and Jones, 2002). The FA values are high for white matter and low for grey matter and CSF-filled regions. A colour-coded image of the principal

fibre orientation (i.e. principal tensor eigenvector) (Pajevic and Pierpaoli, 1999) is also presented. Colours are scaled by the FA, so that low anisotropy regions are suppressed. We can observe familiar orientations for some known tracts, such as inferior-superior orientation within the corona radiata and medial-lateral orientation for the corpus callosum.

2.4 Sampling Q-space

As explained in the previous section, DTI assumes a Gaussian model for the diffusion propagator. Therefore, only a small number of q-space measurements are needed to estimate the diffusion tensor. A spherical acquisition scheme is commonly applied in DTI, typically consisting of 6-30 samples on a q-space sphere. The magnitude $|\mathbf{q}|$ is kept constant, while the direction of the applied diffusion-sensitizing gradients is varied.

For certain applications, more measurements are needed on a q-space sphere (at least 60, for instance). Such spherical schemes were termed, by Tuch et al (Tuch et al., 2002), high angular resolution diffusion imaging (HARDI) protocols. HARDI is used, for example, in Q-ball imaging (Tuch, 2004) and PAS-MRI (Jansons and Alexander, 2003), where features of the diffusion propagator P are of interest rather than P itself. HARDI-based methods will be described in detail in the next chapter.

When the diffusion propagator P needs to be estimated in a model-free way, without imposing any assumptions on it, the q-space relation can be utilized (Eq. (2.6)). This is the underlying idea of q-space imaging (Assaf et al., 2002) and diffusion spectrum imaging (Wedeen et al., 2005) that sample the q-space and estimate the diffusion propagator via a Fourier transform. Assaf et al have performed one-dimensional q-space imaging in healthy subjects and patients (Assaf et al., 2002) to obtain the diffusion propagator along a given direction. Measurements were acquired along a line in q-space by varying the gradient strength $|\mathbf{G}|$. Generalizing to 3D, diffusion spectrum imaging (DSI) (Wedeen et al., 2005) samples a Cartesian q-space grid by varying both the gradient strengths and gradients directions. Despite being demanding in terms of acquisition time, DSI provides directly an estimate of the propagator P .

Combining spherical and Cartesian acquisition schemes, a hybrid protocol has been recently introduced (Wu and Alexander, 2007). Hybrid Diffusion Imaging (HYDI)

acquires measurements on concentric q -space shells. Angular resolution is gradually increased with $|\mathbf{q}|$ value, giving relatively low angular sampling at low $|\mathbf{q}|$ and HARDI measurements at high $|\mathbf{q}|$ values. Using these data, various types of analysis can be followed; for instance, DTI analysis can be performed using low $|\mathbf{q}|$ value shells, while DSI analysis can be performed by regridding the measurements from all acquired shells on a Cartesian lattice.

3

Resolving fibre crossings

Diffusion tensor imaging (DTI) has been the most commonly used DW-MRI method, both in research and clinical environments. This is mainly due to: a) the low sampling requirements that allow fast scanning protocols and b) the processing simplicity associated with the required linear model regression. Despite its potential, DTI suffers from the assumption that, within a voxel, the diffusion scatter pattern exhibits a single directional maxima. This accommodates regions with isotropic diffusion profile, where no coherent structure is present and also regions with highly anisotropic profile, where white matter axons run in parallel.

However, given the low resolution of the DW images (1-2 mm), one can easily imagine more complex scenarios within a voxel in brain white matter, such as crossing, bending and fanning fibres. For all these cases, DTI will reveal a diffusion profile with a single directional maxima, making some fibre configurations indistinguishable from a bundle of unidirectional axons (Seunarine and Alexander, 2009). To resolve such complex scenarios, other model-based or model-free methods have been developed. A few of the existing methods have been introduced in the previous chapter (e.g. DSI), however a more thorough review is presented here, focusing on crossing

fibre configurations that have a diffusion profile with multiple maxima.

3.1 DTI and oblate tensors

A special case of crossing fibres can be picked up by DTI. For close to perpendicular crossings of (geometrically) identical tracts, the estimated diffusion tensor has a disk-like shape, commonly known as planar or oblate (Wiegell et al., 2000). An oblate tensor has approximately equal principal and secondary eigenvalues, both being much larger than the third eigenvalue ($\lambda_1 \approx \lambda_2 \gg \lambda_3$). Its principal eigenvector \mathbf{e}_1 alone does not provide any meaningful structural information, however the plane spanned by \mathbf{e}_1 and \mathbf{e}_2 (and is normal to \mathbf{e}_3) has been empirically found to coincide with the plane in which the fibres cross (Wiegell et al., 2000). In the case of two identical populations with the same volume fraction, crossing at right angles, \mathbf{e}_1 and \mathbf{e}_2 can point anywhere within this plane (with the restriction that they are orthogonal), while \mathbf{e}_3 has a more direct interpretation. Examples of oblate tensors estimated using DTI on voxels with crossing fibres are presented in Fig. 3.1. Different crossing angles were simulated and the estimated diffusion ellipsoids are presented. Note how the FA and the ellipsoid directionality increase as the crossing angle deviates from 90 degrees. This is due to an increase in the difference between the estimated eigenvalues (λ_1 increases and λ_2 decreases). However, even for the case of 60° crossing, \mathbf{e}_3 is normal to the crossing plane.

To discriminate between the three tensor shapes, spherical, oblate and prolate, Westin et al introduced three scalar indices (Westin et al., 2002). The spherical, planar and linear index, respectively, are defined as:

$$c_s = \frac{3\lambda_3}{\lambda_1 + \lambda_2 + \lambda_3}, \quad c_p = \frac{2(\lambda_2 - \lambda_3)}{\lambda_1 + \lambda_2 + \lambda_3}, \quad c_l = \frac{\lambda_1 - \lambda_2}{\lambda_1 + \lambda_2 + \lambda_3}. \quad (3.1)$$

Westin's indices range from zero to one and their sum equals to one. Each of them quantifies similarity to a perfect sphere, plane or line, while $1 - c_{s,p,l}$ quantifies deviation from these shapes. Fig. 3.1 shows how the value of the planar index changes with the crossing angle, when a diffusion tensor is estimated within a voxel with crossing fibres.




Figure 3.1: Oblate diffusion ellipsoids estimated at crossing regions. A mixture of two prolate tensors (FA=0.8) with different crossing angles (90° - 60°) was used to simulate the data. The diffusion tensor (DTI) model was then fitted to the simulated data. For each case, the FA and the planar index c_p of the tensor are shown. The black arrows correspond to the DTI tensor eigenvectors and the blue dashed lines to the crossing orientations.

3.2 The multi-tensor model

The multi-tensor model is a direct extension of the DTI model that allows estimation of $N \geq 1$ fibre orientations (diffusion profile maxima). It stems from the assumption that the diffusing medium can be described by N homogeneous compartments. Within each of them, Gaussian diffusion takes place without exchange to other compartments (Tuch et al., 2002). Then, the diffusion propagator is described as a mixture of N trivariate zero-mean Gaussians with covariance matrices $2\tau\mathcal{D}_i, i = 1 \dots N$ and volume fractions f_i ; each tensor \mathcal{D}_i corresponding to a different compartment. First introduced by Tuch et al (Tuch et al., 2002), the model is:

$$S(\mathbf{g}) = S(0) \cdot \sum_{i=1}^N f_i \exp(-b\mathbf{g}^T \mathcal{D}_i \mathbf{g}), \quad \text{with} \quad \sum_{i=1}^N f_i = 1, \quad (3.2)$$

b and \mathbf{g} characterizing the applied diffusion-sensitizing magnetic field gradient (\mathbf{g} being

a unit direction vector). The principal eigenvectors of the estimated tensors provide fibre orientation estimates. Due to the increased number of parameters (e.g. 13 for $N=2$), this model is usually fitted to HARDI data (Tuch et al., 2002).

The multi-tensor framework has a number of drawbacks. The number N of compartments coexisting in a voxel should be decided in advance, creating thus a model selection problem. Furthermore, the model is non-linear and so can be difficult to fit. Many studies estimate the parameters by minimizing the sum of squared residuals (Alexander and Barker, 2005; Kreher et al., 2005; Peled et al., 2006; Tuch et al., 2002). The estimates may depend on the values used to initialize the fitting procedure, as the objective function exhibits multiple local minima (Tuch et al., 2002). This makes the solution highly unstable, especially for $N>2$. To increase the fitting robustness, a number of geometrical constraints have been employed. A two-tensor model with the eigenvalues of each tensor being fixed to predefined values was utilized in (Tuch et al., 2002), while cylindrically symmetric tensors ($\lambda_2 = \lambda_3$) were fitted in (Alexander and Barker, 2005). Peled et al estimated two cylindrically symmetric tensors that were also restricted to lie within the crossing plane revealed by DTI (see Fig. 3.1), in regions where an oblate DTI tensor was estimated. To overcome partial volume, a two anisotropic and one isotropic tensor model was introduced in (Kreher et al., 2005) for the estimation of two fibre orientations. This captured scenarios where two crossing WM populations coexist in a voxel with fluid-filled regions or GM structures.

Regarding the model selection problem various methods have been employed. In (Tuch et al., 2002), the two-tensor model was fitted in voxels with a low correlation between the measured signal and the one predicted by the DTI model. In (Kreher et al., 2005), both the multi-tensor and the DTI models were applied to a voxel and an F-test assessed the hypothesis that the multi-tensor model improved the fit significantly over the DTI model. Westin's planar index was used in (Peled et al., 2006) to identify oblate DTI tensors, considered representative of crossing regions. In (Alexander et al., 2002), the structural complexity was characterized using a spherical harmonic expansion of the ADC profile ($-\ln[S(\mathbf{g})/S(0)]/b$). For a single fibre population, the ADC obtained from the DTI model is quadratic in terms of the elements of the applied direction \mathbf{g} (see Eq. (2.14)); or equivalently in terms of the sinusoids of the measurement angles θ and ϕ (Eq. (3.12)). It can thus be described by spherical harmonics up to 2nd order (Frank,

2002). By testing the goodness of fit of spherical harmonic series, truncated at 2nd and higher orders, to the measured ADCs, a voxel was classified as containing one or more fibre populations. This model selection technique was further employed in (Parker and Alexander, 2003) to fit one or two tensor models.

The "ball and sticks" model

A special case of the multi-tensor model was introduced by Behrens et al (Behrens et al., 2003, 2007). The tensor mixture comprises of a perfectly isotropic (the "ball") and N perfectly anisotropic compartments (the "sticks"), one per fibre population. The model reduces to:

$$S(\mathbf{g}) = S(0) \left[\left(1 - \sum_{i=1}^N f_i\right) \exp(-bd) + \sum_{i=1}^N f_i \exp(-bd(\mathbf{g} \cdot \mathbf{v}_i)^2) \right], \quad (3.3)$$

with d being the diffusivity and \mathbf{v}_i a vector describing the i^{th} fibre orientation. According to this model, diffusion is allowed along each fibre orientation, while the isotropic compartment captures all other diffusing motions. A modification of this model is presented in (Hosey et al., 2005), with diffusion perpendicular to fibre orientations being explicitly modelled. A graphical representation of the multi-tensor and ball and sticks model is presented in Fig. 3.2.

A common feature of the above studies (Behrens et al., 2003, 2007; Hosey et al., 2005) is that they use a Bayesian framework to perform the model regression. The model parameters ω are treated as random variables and the joint posterior probability density function, given the data Y , is defined using Bayes' theorem: $\Pi(\omega|Y) = \Pi(Y|\omega)\Pi(\omega) / \int_{\omega} \Pi(Y|\omega)\Pi(\omega)d\omega$. A random sample of values for each parameter can then be drawn from the posterior $\Pi(\omega|Y)$, using e.g. Markov-Chain Monte-Carlo (MCMC) algorithms (Behrens et al., 2003). Estimating a distribution rather than a single value per model parameter is particularly useful for probabilistic tractography methods, as it will be shown in the next chapter.




Figure 3.2: Graphical representation of the multi-tensor and the ball and sticks model used to resolve fibre crossings.

The tensor distribution function

A direct extension of the multi-tensor model has been recently introduced (Jian et al., 2007; Leow et al., 2009). Rather than estimating a finite number of tensors, a tensor distribution function (TDF) is resolved. The TDF is defined on the space of second order tensors \mathcal{D} and assigns weights to the tensors that best explains the measured signal. The underlying model is (Leow et al., 2009):

$$S(\mathbf{g}) = S(0) \int_{\mathcal{D}} \text{TDF}(\mathcal{D}) \exp(-b\mathbf{g}^T \cdot \mathcal{D} \cdot \mathbf{g}) d\mathcal{D} . \quad (3.4)$$

According to Eq. 3.4, the diffusion propagator is modelled as an infinite sum of zero-mean trivariate Gaussians, each having a covariance matrix $2\tau\mathcal{D}_i$ and a volume fraction $\text{TDF}(\mathcal{D}_i)$. Once the TDF is estimated, a tensor orientation distribution (TOD) can be obtained by integrating over the eigenvalues:

$$\text{TOD}(\mathbf{u}) = \int_{\lambda} \text{TDF}(\mathcal{D}(\lambda, \mathbf{u})) d\lambda , \quad (3.5)$$

with $\lambda = (\lambda_1, \lambda_2, \lambda_3)$ and \mathbf{u} the principal direction of tensor \mathcal{D} . The TOD has a very sharp angular profile and can provide multiple peaks in crossing regions, without the

need of model selection.

A similar model as in Eq. 3.4 is followed in (Jian et al., 2007). However, the TDF is parameterized by a mixture of Wishart distributions, the latter being distributions defined for non-negative matrix-valued random variables. The shape of the distributions is predefined and common for all image voxels, assuming that all fibres have the same anisotropy profile.

3.3 Q-ball imaging and diffusion ODF

Most multiple tensor approaches utilize HARDI data, due to the increased number of model parameters. Other model-free or model-based approaches also utilize spherical sampling schemes with relatively high angular resolution. Q-ball imaging is a model-free approach that is presented here, while other HARDI-based methods are reviewed in the next section.

Q-ball imaging (QBI) (Tuch, 2004) is amongst the most popular HARDI-based methods used to resolve fibre crossings. It reconstructs the angular profile of the diffusion propagator P , commonly known as the orientation distribution function (ODF). While P represents the probability of displacing $r \mu m$ along a direction \mathbf{u} , the ODF represents how likely is to observe any diffusion displacement along \mathbf{u} . It is the radial projection of P and can be obtained by integrating P over all possible displacement distances along a unit direction \mathbf{u} (Tuch, 2004):

$$ODF(\mathbf{u}) = \frac{1}{Z} \int_0^{\infty} P(r\mathbf{u}) dr, \quad (3.6)$$

with Z being a constant that ensures ODF normalization to unit mass.

The Q-ball imaging ODF is estimated using the Funk-Radon transform (\mathcal{FRT}) of the attenuated signal $E(\mathbf{q}) = S(\mathbf{q})/S(0)$. The \mathcal{FRT} of the sampled $E(\mathbf{q})$ along a direction \mathbf{u} is the integral of $E(\mathbf{q})$ values along the equator perpendicular to \mathbf{u} . In (Tuch, 2004), it is shown that the \mathcal{FRT} of the signal attenuation is approximately equal to the ODF:

$$ODF(\mathbf{u}) \approx \mathcal{FRT}[E(\mathbf{q})](\mathbf{u}) \equiv \int \delta_D(\mathbf{u} \cdot \mathbf{q}) \cdot E(\mathbf{q}) d\mathbf{q}, \quad (3.7)$$

with δ_D being the Dirac function. The approximation improves for larger sampling radii $|\mathbf{q}|$. For relatively lower $|\mathbf{q}|$ the \mathcal{FRT} provides a smoothed version of the true ODF.

To implement Eq. (3.7) numerically, a large number of measurements E are needed along the equator perpendicular to \mathbf{u} . Then $ODF(\mathbf{u}) \approx (1/\mathcal{Z}) \sum_i E(\tilde{\mathbf{q}}_i)$, with $\tilde{\mathbf{q}}_i$ being points on the respective equator and \mathcal{Z} a constant ensuring normalization to unit mass. Since the points $\tilde{\mathbf{q}}_i$ will not necessarily coincide with the measurement points, an interpolation scheme is utilized. The measured signal can be expressed as a linear combination of some basis functions $\psi(\cdot)$, such that (Alexander, 2005a):

$$E(\mathbf{q}) = \sum_{j=1}^R c_j \psi_j(\mathbf{q}) . \quad (3.8)$$

As will be shown in section 3.4, the coefficients c_j can be estimated from the measurements and Eq. (3.8) can then provide an estimated measurement for any point $\tilde{\mathbf{q}}_i$. In (Tuch, 2004), spherical radial basis functions were employed to interpolate the signal. A large number of diffusion-sensitizing gradients were also used (>252) to robustly estimate the ODFs. In more recent studies (Descoteaux et al., 2007; Hess et al., 2006), spherical harmonics were utilized, offering a more compact set of basis functions that reduce scanning requirements to 50-60 directions. Spherical harmonics also exhibit a property that allows the analytic computation of the Funk-Radon transform, without the need for numerical integration.

The ODF exhibits multiple local maxima in crossing regions, which are routinely used as fibre orientation estimates; even if these estimates are biased for small crossing angles (Zhan and Yang, 2006). A generalized anisotropy measure (GFA) can be also defined, using the standard deviation of the ODF values on the sphere (Tuch, 2004):

$$GFA = \sqrt{\frac{n \sum_{i=1}^n (ODF(\mathbf{u}_i) - \overline{ODF})^2}{(n-1) \sum_{i=1}^n ODF(\mathbf{u}_i)^2}} , \quad (3.9)$$

with n evaluations of the ODF needed at points \mathbf{u}_i (evenly distributed on the sphere) and \overline{ODF} being the mean ODF value.

ODF for Gaussian diffusion propagator

A simple analytic expression can be obtained for the diffusion ODF, if the Gaussian diffusion propagator of DTI is utilized. Using Eqs. (2.9) and (3.6) we obtain:

$$ODF(\mathbf{u})_{DTI} = \frac{1}{\sqrt{(4\pi\tau)^3|\mathcal{D}|}} \cdot \int_0^\infty \exp\left(\frac{-\mathbf{u}^T \cdot \mathcal{D}^{-1} \cdot \mathbf{u}}{4\tau} r^2\right) dr. \quad (3.10)$$

This integral is of a known form that can be analytically computed (Zwillinger, 2003):

$$\int_0^\infty \exp(-\alpha r^\beta) r^n dr = \frac{\Gamma(\frac{n+1}{\beta})}{\beta \alpha^{\frac{n+1}{\beta}}},$$

with $\Gamma(\cdot)$ being the gamma function. Using $n = 0$, $\beta = 2$ and \mathcal{Z} a normalization constant, we obtain:

$$ODF(\mathbf{u})_{DTI} = \frac{\mathcal{Z}}{\sqrt{\mathbf{u}^T \cdot \mathcal{D}^{-1} \cdot \mathbf{u}}}. \quad (3.11)$$

The difference between the Q-ball and DTI ODFs is that the latter is single-peaked due to the Gaussian assumption. This is illustrated in Fig. 3.3, where the ODFs estimated from human diffusion tensor and Q-ball imaging data in the same region are presented.




Figure 3.3: DTI and Q-ball ODFs plotted for (roughly) the same coronal slice. The crossing between the corpus callosum and the corona radiata at the region of the centrum semiovale is shown magnified for each case. Q-ball ODFs have clearly two peaks in this region. All ODFs have been sharpened (min-max normalized (Tuch, 2004)). DTI ODFs are superimposed on FA maps and Q-ball ODFs on GFA maps.

3.4 Spherical harmonics for diffusion ODF estimation

In this thesis, spherical harmonics will be used to reconstruct Q-ball ODFs. The framework is, therefore, presented here parenthetically.

Let θ ($0 \leq \theta \leq \pi$) and ϕ ($0 \leq \phi < 2\pi$) the zenith and azimuth angles in a spherical coordinate system, as shown in Fig. 3.4. Any point (x, y, z) can then be described by these two angles and vice versa:

$$\begin{aligned}
 \theta &= \arccos(z/r) & x &= r \sin\theta \cos\phi \\
 \phi &= \arctan(y/x) & y &= r \sin\theta \sin\phi \\
 r &= \sqrt{x^2 + y^2 + z^2} & z &= r \cos\theta
 \end{aligned} \quad (3.12)$$




Figure 3.4: Zenith θ and azimuth ϕ angles in a spherical coordinate system.

Spherical harmonics are complex functions of (θ, ϕ) defined on the surface of the unit sphere. A harmonic of order $l \in \{0, 1, 2, \dots\}$ and degree $m \in \{-l, -l+1, \dots, 0, \dots, l-1, l\}$ is defined as (Press et al., 1992):

$$Y_l^m(\theta, \phi) = \begin{cases} \sqrt{\frac{2l+1}{4\pi} \cdot \frac{(l-m)!}{(l+m)!}} \cdot P_l^m(\cos\theta) \cdot e^{im\phi} & \text{if } m \geq 0 \\ (-1)^{-m} Y_l^{-m*}(\theta, \phi) & \text{if } m < 0, \end{cases} \quad (3.13)$$

with Y^* the complex conjugate and P_l^m the associated Legendre polynomials:


$$P_l^m(x) = (-1)^m (1-x^2)^{m/2} \left(\frac{d^m}{dx} P_l(x) \right), \quad m \geq 0 \quad (3.14)$$

$$P_l(x) = \frac{1}{2^l l!} \left(\frac{d^l}{dx} (x^2 - 1)^l \right).$$

The harmonics of the first few orders are shown in Fig. 3.5. Spherical harmonics (SH) form an orthonormal basis, such that every complex function f on the unit sphere can be written as a linear combination of SHs (Alexander et al., 2002):

$$f(\theta, \phi) = \sum_{l=0}^{\infty} \sum_{m=-l}^l c_{lm} Y_l^m(\theta, \phi). \quad (3.15)$$

The coefficients of the above SH series can be obtained from the inverse transform,


 Figure 3.5: Real part squared of spherical harmonics up to $l=4$ order.

called the spherical harmonic transform (SHT):

$$c_{lm} = \int_0^{2\pi} \int_0^\pi f(\theta, \phi) Y_l^{m*}(\theta, \phi) \sin\theta \, d\theta d\phi . \quad (3.16)$$

A harmonic function of order l contains polynomial terms up to order l . Thus, the harmonic orders l effectively correspond to frequency modes on the unit sphere (Frank, 2002). Higher orders represent higher frequency components. Relatively smooth functions can be described using SHs of low order. Therefore, in practice, the SH series is truncated to a maximum order l_{max} .

Real spherical harmonics

The measured diffusion signal can be considered as a function on the sphere. Therefore, it can be represented as a SH series. Spherical harmonics of even order l exhibit antipodal symmetry, i.e. $Y_l^m(\pi - \theta, \phi + \pi) = Y_l^m(\theta, \phi)$, while of odd order l are antipodally anti-symmetric, i.e. $Y_l^m(\pi - \theta, \phi + \pi) = -Y_l^m(\theta, \phi)$. Since the diffusion signal is antipodally symmetric, harmonics of even order only are used to represent it. Further-

more, the diffusion signal attenuation is real. A modified SH basis that comprises of real spherical harmonic functions exists (Descoteaux et al., 2007):

$$\begin{aligned}
 YR_l^m(\theta, \phi) &= \begin{cases} \sqrt{\frac{2l+1}{2\pi}} \cdot \frac{(l-m)!}{(l+m)!} \cdot P_l^m(\cos\theta) \cdot \sin(m\phi) & \text{if } m > 0 \\ Y_l^0(\theta, \phi) & \text{if } m = 0 \\ (-1)^{-m} \sqrt{\frac{2l+1}{2\pi}} \cdot \frac{(l+m)!}{(l-m)!} \cdot P_l^{-m}(\cos\theta) \cdot \cos(-m\phi) & \text{if } m < 0 \end{cases} \\
 &= \begin{cases} \sqrt{2} \cdot \text{Im}[Y_l^m(\theta, \phi)] & \text{if } m > 0 \\ Y_l^0(\theta, \phi) & \text{if } m = 0 \\ \sqrt{2} \cdot (-1)^{-m} \cdot \text{Re}[Y_l^{-m}(\theta, \phi)] & \text{if } m < 0. \end{cases} \quad (3.17)
 \end{aligned}$$

Following the convention of (Descoteaux et al., 2007), we can define a single harmonic index $j = (l^2 + l + 2)/2 + m$ with $l \in \{0, 2, 4, 6, \dots, l_{\max}\}$ and $m \in \{-l, -l + 1, \dots, 0, \dots, l - 1, l\}$. Setting $Y_j(\theta, \phi) = YR_l^m(\theta, \phi)$, we can describe the real, symmetric diffusion signal attenuation using even real SHs:

$$E(\theta, \phi) = \sum_{j=1}^R c_j Y_j(\theta, \phi), \quad (3.18)$$

with $R = (l_{\max} + 1)(l_{\max} + 2)/2$ the number of SH coefficients c_j . The DW signal is measured along M diffusion-sensitizing directions (θ_i, ϕ_i) . Using these measurements, $R \leq M$ coefficients can be estimated using linear least squares (Alexander et al., 2002). If \mathbf{E} is a $M \times 1$ vector containing all the measurements $E(\theta_i, \phi_i)$, \mathbf{C} a $R \times 1$ vector with the unknown SH coefficients and \mathbf{Y} a $M \times R$ matrix containing the values of the SH functions at points (θ_i, ϕ_i) , then we can write:

$$\begin{aligned}
 \begin{bmatrix} E(\theta_1, \phi_1) \\ E(\theta_2, \phi_2) \\ \vdots \\ E(\theta_M, \phi_M) \end{bmatrix} &= \begin{bmatrix} Y_1(\theta_1, \phi_1) & Y_2(\theta_1, \phi_1) & \cdots & Y_R(\theta_1, \phi_1) \\ Y_1(\theta_2, \phi_2) & Y_2(\theta_2, \phi_2) & \cdots & Y_R(\theta_2, \phi_2) \\ \vdots & \vdots & \ddots & \vdots \\ Y_1(\theta_M, \phi_M) & Y_2(\theta_M, \phi_M) & \cdots & Y_R(\theta_M, \phi_M) \end{bmatrix} \cdot \begin{bmatrix} c_1 \\ c_2 \\ \vdots \\ c_R \end{bmatrix} \implies \\
 \mathbf{E} &= \mathbf{Y} \cdot \mathbf{C} \implies \mathbf{C} = (\mathbf{Y}^T \mathbf{Y})^{-1} \mathbf{Y}^T \mathbf{E}. \quad (3.19)
 \end{aligned}$$

This solution provides a discrete approximation to the exact solution of Eq. (3.16).

ODF representation using spherical harmonics

A useful relationship exists that can be used to represent ODFs analytically in terms of spherical harmonics. Both Hess et al (Hess et al., 2006) and Descoteaux et al (Descoteaux et al., 2007) have proved that:

$$\int \delta_D(\mathbf{u} \cdot \mathbf{q}) \cdot Y_j(\mathbf{q}) d\mathbf{q} = 2\pi P_{l_j}(0) \cdot Y_j(\mathbf{u}) , \quad (3.20)$$

where P_{l_j} the Legendre polynomial and l_j the order associated with the j^{th} element of the modified SH basis. If we express the signal attenuation in spherical harmonics as in Eq. (3.18) and compute the coefficients c_j through Eq. (3.19), the ODF can then be estimated as:

$$\begin{aligned} ODF(\mathbf{u}) &\approx \int \delta_D(\mathbf{u} \cdot \mathbf{q}) \cdot E(\mathbf{q}) d\mathbf{q} = \int \delta_D(\mathbf{u} \cdot \mathbf{q}) \cdot \sum_{j=1}^R c_j Y_j(\mathbf{q}) d\mathbf{q} \implies \\ ODF(\mathbf{u}) &\approx \sum_{j=1}^R c_j \cdot 2\pi P_{l_j}(0) \cdot Y_j(\mathbf{u}) . \end{aligned} \quad (3.21)$$

3.5 Other HARDI-based approaches

Persistent angular structure and Diffusion orientation transform

Even after sharpening, ODFs are rounded and smooth. However, the fibre orientation profiles are expected to be relatively spiky. Towards this direction, the persistent angular structure (PAS) MRI computes an angular feature of the diffusion propagator that has a sharp profile (Jansons and Alexander, 2003). The PAS function \tilde{p} is defined on the unit sphere as a non-linear combination of cosines that captures the angular structure of the diffusion scatter pattern:

$$\tilde{p}(\mathbf{u}) = \exp \left(\lambda_0 + \sum_{j=1}^M \lambda_j \cos(r \mathbf{g}_j \cdot \mathbf{u}) \right) . \quad (3.22)$$

M is the number of measurements, \mathbf{g}_j are the measurement directions, r is a smoothing parameter and λ 's are the unknown coefficients. To estimate them, the Fourier transform of the PAS is fitted to the spherically sampled signal attenuations. The peaks of the

PAS indicate fibre orientations. PASMRI requires fewer samples than QBI to achieve the same accuracy in resolving fibre crossings (Alexander, 2005a). Its main drawback is the high computational requirements, however recent work has shown optimized computation time (Sakaie, 2008).

The diffusion orientation transform (DOT) (Ozarslan et al., 2006) calculates a variant of the diffusion ODF. The angular structure $P(R\mathbf{u})$ of the diffusion propagator is estimated at a fixed radius R , contrary to Q-ball imaging that integrates over all possible radii. HARDI data, sampled on a single q-space shell, are utilized and DOT is applied to the ADC values $-\ln[S(\mathbf{g})/S(0)]/b$ corresponding to the measurements. By increasing R , DOT can provide finer estimates of the underlying fibre orientations in crossing regions.

Resolving the ADC profile

A group of methods have particularly focused on resolving the ADC profile using HARDI data. The ADC, as a function on the sphere, has been represented either through spherical harmonics (Descoteaux et al., 2006) or higher order tensors (Ozarslan and Mareci, 2003):

$$ADC(\mathbf{g}) \equiv -\frac{1}{b} \ln \left[\frac{S(\mathbf{g})}{S(0)} \right] = \sum_{l=0}^{\infty} \sum_{m=-l}^l c_{lm} Y_l^m(\mathbf{g}) = \sum_{i_1=1}^3 \sum_{i_2=1}^3 \cdots \sum_{i_l=1}^3 \mathbf{D}_{i_1 i_2 \dots i_l} g_{i_1} g_{i_2} \cdots g_{i_l} , \quad (3.23)$$

with \mathbf{D} being a symmetric, rank l tensor that comprises of 3^l elements, of which only $(l+1)(l+2)/2$ are independent (Ozarslan and Mareci, 2003). Estimation of a higher order tensor of rank l is termed Generalized DTI and reduces to DTI for $l=2$. It is equivalent to fitting a spherical harmonic series of maximum order l . As shown by Ozarslan et al (Ozarslan and Mareci, 2003), the spherical harmonics coefficients up to order l are algebraically related to the elements of a rank l tensor. Higher order tensors can provide new tensor-derived scalar parameters, such as a generalized anisotropy index, which performs better in voxels of complex structure than the DTI FA. However, estimating fibre orientations from ADC profiles, especially in crossing regions, is still far from obvious.

3.6 Spherical deconvolution and fibre ODF

Spherical deconvolution methods utilize HARDI data, as well. However, they are presented here separately, due to the large number of existing implementations.

Even if features of the diffusion propagator, such as the diffusion ODF and the PAS function, have been empirically found to provide orientation estimates, the function of interest for tractography is the distribution of fibre orientations. Spherical deconvolution (Alexander, 2005b; Anderson, 2005; Dell'Acqua et al., 2007; Tournier et al., 2004) allows the direct estimation of such a function, known as the fibre orientation distribution function (fODF). The measured signal $E(\mathbf{q})$ is treated as the convolution of the fODF F with the signal $R(\mathbf{q}, \mathbf{u})$ obtained at a measurement wavevector \mathbf{q} from a single fibre population with orientation \mathbf{u} :

$$E = F \otimes R \quad \Rightarrow \quad E(\mathbf{q}) = \int F(\mathbf{u})R(\mathbf{q}, \mathbf{u})d\mathbf{u}. \quad (3.24)$$

If the single fibre response R is known, then the unknown fODF F can be estimated using deconvolution methods. This approach assumes that there is no exchange between the mixing compartments, within a voxel (Tournier et al., 2004).

The deconvolution proposed by Tournier et al (Tournier et al., 2004) is transformed to the unit sphere domain. Both E and F are represented by spherical harmonics, while R is estimated from the most anisotropic voxels in the brain and represented by rotational harmonics (Healy et al., 1998). Noticing that a spherical harmonic transform is effectively a Fourier transform on the sphere allows the representation of convolution with a set of $(l_{max} + 2)/2$ matrix multiplications $\mathbf{E}^l = \mathbf{R}^{ll} \cdot \mathbf{F}^l, l = 0, 2, \dots, l_{max}$ (Healy et al., 1998). Vector \mathbf{E}^l comprises of the l^{th} order spherical harmonic coefficients estimated from the attenuated signal, vector \mathbf{F}^l of the unknown l^{th} order spherical harmonic coefficients of the fODF and matrix \mathbf{R}^{ll} of the l^{th} order rotational harmonics coefficients estimated from the single fibre response. Therefore, deconvolution is reduced to a set of matrix inversions, one per utilized harmonic order l (Tournier et al., 2004).

Tournier et al (Tournier et al., 2004) assume a common single fibre response across all voxels. To improve on this assumption, Anderson derives an individual response function R for each voxel based on a multi-tensor model (Anderson, 2005). R depends

on the mean and transverse diffusivities estimated from the data. Deconvolution is performed using spherical harmonic representations for all three functions E , R and F that allows an analytic calculation of the fODF coefficients.

Utilizing truncated spherical harmonics for deconvolution introduces artefacts (Anderson, 2005; Tournier et al., 2004) that appear as small-amplitude, high frequency lobes near the origin. This is due to the fact that the fODF is a very sharp function, ideally a sum of delta functions. Furthermore, non-negativity of the estimated fODF, caused by noise, harmonics truncation or an inappropriate response function is not guaranteed using the methods described before. Both issues are tackled in (Dell'Acqua et al., 2007), using an iterative algorithm to perform the deconvolution. In (Tournier et al., 2004), noise artifacts are low-pass filtered by attenuating the coefficients of higher harmonic orders. In (Tournier et al., 2007), the authors impose zero fODF amplitude as a constraint along the directions with negative fODF values. Refined estimates of the fibre orientation distribution are obtained in an iterative fashion, until convergence is achieved. In (Alexander, 2005b), a non-linear representation of the fODF is utilized. This ensures a sharp profile and positive definiteness of the estimated function.

A Bayesian framework for performing deconvolution is introduced in (Kaden et al., 2007). This estimates not only the fODF in each voxel, but also its uncertainty given the data, by drawing a random sample of fODFs. To achieve that a model-based approach is followed. The fODF is modelled as a mixture of N Bingham distributions, each of which represents a single fibre population. The signal is then expressed as the convolution integral of this Bingham mixture with a single fibre response kernel, common to all voxels. The drawback of this technique is that a model selection problem arises for finding the N that best fits the data.

FODF and diffusion ODF

In a recent study, Descoteaux et al (Descoteaux et al., 2009) developed a formal relationship between the fODF and the diffusion ODF. The diffusion ODF in a voxel can be described as the convolution of the fODF F with the diffusion ODF R_{ODF} of a single fibre population: $ODF = F \otimes R_{ODF}$. Thus, the ODF deconvolution will provide an estimate of the fODF. Using real spherical harmonics, described in Section 3.4, an analytic

solution exists:

$$F(\mathbf{u}) = \sum_{j=1}^R f_j \cdot Y_j(\mathbf{u}) = \sum_{j=1}^R \frac{c_j}{r_j} \cdot 2\pi P_{l_j}(0) \cdot Y_j(\mathbf{u}) , \quad (3.25)$$

with f_j and c_j being the spherical harmonic coefficients for the fODF and the measured signal attenuation, respectively, and r_j being coefficients that depend on the single fibre response R_{ODF} (they are given analytically in (Descoteaux, 2008)). It should be pointed out that $c_j 2\pi P_{l_j}(0)$ are the coefficients of the diffusion ODF and the fODF coefficients can be obtained with a single scaling factor. A comparison between DTI ellipsoids, diffusion ODFs and fODFs for a single population, a 60° and a 90° crossing are shown in Fig. 3.6.




Figure 3.6: DTI ellipsoids, Q-ball diffusion ODFs and fODFs estimated on simulated data of (from top to bottom): a single fibre bundle, a 60° crossing and a 90° crossing. ODFs sharpened through min-max normalization (Tuch, 2004) are also presented. FODFs were obtained through ODF deconvolution (Descoteaux et al., 2009). Simulation parameters: $b=3000$ s/mm², 61 directions, individual bundle FA=0.8).

3.7 Diffusion spectrum imaging and other q-space approaches

Despite HARDI being the most popular scanning protocol for resolving fibre crossings, Diffusion Spectrum Imaging (DSI) (Wedeen et al., 2005) offers another alternative, by

allowing direct estimation of the diffusion propagator P , as explained in section 2.4. The diffusion ODF can then be calculated by radial integration of the propagator (Eq. 3.6). Exactly as in Q-ball imaging, the DSI ODF peaks provide estimates of fibre orientations. However, the DSI ODF provides a better estimation of the radial projection of P compared to the Q-ball ODF, which approximates it by the Funk-Radon transform of the measurements. On the other hand, sampling a q-space grid is very time consuming, making DSI impractical for certain applications. Even the most recent optimized DSI protocols (Wedeen et al., 2008) require roughly 50 minutes of scanning per subject to achieve a moderate spatial resolution.

Composite hindered and restricted model of diffusion

Compared to DSI, the composite hindered and restricted model of diffusion (CHARMED) (Assaf and Basser, 2005; Assaf et al., 2004) requires a smaller number of q-space samples. It has been found that the measured signal along a certain direction, as a function of $|\mathbf{q}|$ deviates from its monoexponential decay character used in DTI, particularly at high $|\mathbf{q}|$ values and in WM (Clark and Le Bihan, 2000). Instead, the signal follows a multi-exponential decay and diffusion deviates from the assumed Gaussian behaviour. Commonly, two different diffusive compartments are evident, known as the fast and slow compartments (Clark and Le Bihan, 2000). The former has a higher volume fraction and ADC than the latter, which contributes to the signal only at high $|\mathbf{q}|$. In the CHARMED framework, a different model is assigned to each of the two compartments, which are treated as extra-axonal and intra-axonal spaces (Assaf and Basser, 2005; Assaf et al., 2004). A model of hindered diffusion describes diffusion in the extra-axonal space and a sum of models of restricted diffusion describes intra-axonal diffusion. The signal attenuation is:

$$E(\mathbf{q}) = f_h E_h(\mathbf{q}) + \sum_{k=1}^N f_{rk} E_{rk}(\mathbf{q}) . \quad (3.26)$$

The explicit models for E_h and E_r are presented in (Assaf et al., 2004). Briefly, the hindered model is Gaussian (i.e. free diffusion) and the restricted non-Gaussian, describing diffusion within a cylinder. E_r dominates at high $|\mathbf{q}|$ and describes much more anisotropic diffusion than E_h , in agreement with the findings in (Clark and Le Bihan,

2000) that the slow compartment is more anisotropic than the fast. The number N of restricted compartments used in the model defines the maximum number of crossing fibre orientations that can be resolved. The scanning protocol is described in (Assaf and Basser, 2005) and comprises of multiple spherical shells with the sampling angular resolution increasing with $|\mathbf{q}|$.

Diffusion kurtosis imaging

Another q-space approach that requires fewer samples than DSI is diffusion kurtosis imaging (DKI) (Jensen et al., 2005). DKI estimates the covariance (2nd moment) and kurtosis (4th moment) of the diffusion propagator P , rather than P itself. The kurtosis quantifies deviation of a probability density function from Gaussian behaviour. Therefore, it increases with structure complexity and heterogeneity and can therefore be used as a measure of tissue structure (Jensen et al., 2005). It is in general small along the fibres and large across them. The signal equation is modified to:

$$S(b) = S(0) \exp(-bD_{app} + \frac{1}{6}b^2D_{app}^2K_{app}) , \quad (3.27)$$

with D_{app} and K_{app} the apparent diffusion and kurtosis coefficients along the direction of the applied diffusion sensitizing gradient. The above equation is valid for relatively small b values ($<3000 \text{ s/mm}^2$). If enough measurements are obtained, the 2nd order diffusion and 4th order kurtosis tensors can be estimated from the apparent coefficients (Lu et al., 2006). Scalar measures derived from the kurtosis tensor have been shown to provide enhanced contrast between different tissue types compared to their DTI counterparts.

Orientation information can be also obtained from DKI. Taking the Fourier transform of the signal in Eq. (3.27) and then its radial projection gives the diffusion kurtosis ODF (Lazar et al., 2008). This has both Gaussian and non-Gaussian components. The latter can reliably resolve fibre crossings of small crossing angles, where traditional HARDI methods, such as Q-ball imaging, fail.

4

White matter tractography

Through diffusion MRI, structural features of tissue can be resolved. Especially for brain white matter (WM), orientations of fibre bundles can be estimated within each voxel, as discussed in the previous chapter. Over the last years, a group of algorithms have been developed that utilize these orientation estimates to reconstruct WM tracts in the brain. Commonly known as *tractography* methods, they provide a tool to resolve major neuronal fibre bundles non-invasively and in-vivo (Catani et al., 2002; Mori et al., 2005), as shown in Fig. 4.1.

Tractography algorithms effectively reconstruct the paths in the brain, along which water diffusion is least hindered (Behrens and Jbabdi, 2009). In the case of WM, molecules are more hindered in a systematic way by the axonal membranes and the myelin sheaths (Beaulieu, 2002). They will, thus, diffuse along rather than across the neuronal axons. This preferred diffusion directionality can indirectly infer geometrical features of the underlying axonal bundles. Therefore, tractography methods provide information on the route of group of axons, rather than on individual axonal processes (note that individual subcortical axons have a diameter in the order of a few μm , while the DW imaging resolution is in the order of 1-2 mm).




Figure 4.1: Major white matter tracts reconstructed using DW imaging and tractography (FACT reconstruction with the brute-force approach (Huang et al., 2004) was utilized on DTI images). An axial (left) and a sagittal (right) view are shown with the brain outline in grey. Blue: Genu and splenium of corpus callosum, Yellow: Cingulum, Green: Superior longitudinal fasciculus, Red: Inferior longitudinal fasciculus, Orange: Inferior fronto-occipital fasciculus, Turquoise: Optic radiation, Purple: Part of corona radiata, Dark orange: Anterior thalamic radiation.

Most tractography techniques can be grouped in two categories, as pointed out in (Behrens and Jbabdi, 2009). Local approaches propagate a curve from a starting (seed) point using locally greedy criteria, i.e. tracking sequentially through orientation estimates in adjacent voxels. Global approaches identify the best path between two points of interest, according to some optimization criterion, rather than identifying paths arising from a single point. In the next sections, representative methods from each category are discussed. A third category is also presented, comprising of algorithms that simulate the diffusion process or solve the diffusion equation to reconstruct WM tracts.

4.1 Local approaches

Deterministic streamline tractography

Fibre orientations \mathbf{e} can be estimated in each voxel from DW images. For now, we will assume a single orientation per voxel, as is the case for DTI, where the principal eigenvector of the diffusion tensor is used in each voxel. However, the approaches

discussed here can be generalized to include multiple intravoxel orientations, as shown in (Chao et al., 2008; Descoteaux et al., 2009; Wedeen et al., 2008).

The orientation estimates form a vector field. Assuming a white matter tract can be represented by a 3D space curve, the simplest tractography approaches estimate the curve $\mathbf{r}(s)$ that is tangent to the orientation vector field (Basser et al., 2000). Such a curve is called a streamline. Streamline tractography is then governed by the differential equation:

$$\frac{d\mathbf{r}(s)}{ds} = \mathbf{e}(\mathbf{r}(s)), \quad \mathbf{r}(0) = \mathbf{r}_0. \quad (4.1)$$

The vector $\mathbf{r}(s)$ gives a point on the curve, s is the distance along the curve and \mathbf{r}_0 a starting (seed) point. The above differential equation is continuous, however orientation estimates are available only for discrete locations. It can thus be solved numerically, by approximating the rate of change $d\mathbf{r}(s)/ds$ with a finite difference. This gives the Euler approximation, first used in (Conturo et al., 1999):

$$\mathbf{r}(s_{n+1}) = \mathbf{r}(s_n) + h\mathbf{e}(\mathbf{r}(s_n)), \quad (4.2)$$

with h being a small step size (small relative to the curvature of the trajectory), during which the orientation is assumed unchanged. From the above equation, one can observe that orientation estimates are needed for any arbitrary location $\mathbf{r}(s_n)$. However, measurements are available only on specific locations. If we assume that the estimates lie at the centre of each voxel, we can apply interpolation techniques to "guess" what an estimate would be at a different location. A trilinear interpolation scheme can be employed (Tench, 2003). Another popular approach is to employ a nearest neighbour interpolation, which effectively means that the same orientation estimate is followed over the entire voxel (Mori et al., 1999). Known as fibre assignment by continuous tracking (FACT), it has been amongst the first tractography algorithms. Due to the lack of a smooth interpolant, FACT returns more rough streamlines, as shown in Fig. 4.2.

Regardless of the propagation method used, the complete trajectory from a seed is estimated in two steps, one with $\mathbf{e}(\mathbf{r}_0)$ and a second with $-\mathbf{e}(\mathbf{r}_0)$. Propagation continues until some stopping criteria are met. Tractography is usually stopped: a) when a low anisotropy region is encountered (e.g. FA drops below 0.2) to avoid propagation within CSF or GM regions and/or b) a large angle changes occur (e.g. larger than 45°) between




Figure 4.2: Streamline tractography performed using FACT (left) and Euler method with trilinear interpolation (right). Wide arrows represent the fibre orientation estimate per voxel, while the black curved arrow the tract trajectory. Small black arrows on the right panel represent the interpolated vector field. Voxels that are identified as belonging to the tract of interest are coloured with grey.

successive steps to avoid unrealistically sharp turns (Basser et al., 2000; Mori et al., 1999). An example of streamline propagation within a DTI fibre orientation field in the corpus callosum tract is shown in Fig. 4.3.

The approaches discussed so far tend to underestimate fibre branching. From a seed point one curve will be reconstructed, therefore when a branching is encountered one of the branching paths will be followed. A remedy to this problem has been proposed in (Huang et al., 2004) via the brute-force approach. Instead of propagating curves from a seed region, curves are seeded from all image voxels and the ones that traverse the region of interest are kept.

Streamline methods can be also utilized when multiple orientations are estimated in a voxel. These can be obtained, for instance, as the principal eigenvectors of multiple tensors fitted to the data (Tuch et al., 2002), as the local peaks of the diffusion ODF estimated using DSI (Wedeen et al., 2005) or QBI (Tuch, 2004), or as the local peaks of the fibre ODFs (Tournier et al., 2004) (see (Alexander, 2005a) for a review). An approach to deal with these cases, chooses, upon entering a voxel, the orientation that produces the smallest curvature with the incoming path (Parker and Alexander, 2003; Wedeen et al., 2008). Another approach follows all orientations that do not exceed a curvature threshold, by initiating a new streamline per orientation (Chao et al., 2008; Descoteaux et al., 2009).




Figure 4.3: Left: DTI fibre orientation estimates in the region of corpus callosum superimposed on FA values. The orientational coherence is evident, especially in regions of high anisotropy. Right: A streamline generated within the vector field using the Euler method with trilinear interpolation. The seed point is indicated with the blue dot.

Streamline tractography utilizes only the fibre orientation estimates. In general, more information is available upon post-processing of DW images. For example, with DTI a diffusion tensor is computed. Tensor deflection tractography is a FACT variant that uses the whole DTI tensor rather than just its principal eigenvector to determine the direction of curve propagation (Lazar et al., 2003). In each voxel, the outgoing direction vector $\mathbf{e}(\mathbf{r}(s))$ is determined by the application of the respective diffusion tensor \mathcal{D} to the incoming path direction vector, i.e. $\mathbf{e}(\mathbf{r}(s_n)) = \mathcal{D} \cdot \mathbf{e}(\mathbf{r}(s_{n-1}))$. The vector $\mathbf{e}(\mathbf{r}(s_{n-1}))$ is deflected towards the direction of the principal eigenvector of \mathcal{D} , with the deflection angle increasing with tensor anisotropy. For perfectly spherical tensors the deflection angle is zero, while for oblate tensors, the deflection is towards the ellipsoid plane. Tensor deflection tractography is less sensitive to noise than streamline, as shown in simulations for straight tracts (Lazar et al., 2003). It can also propagate through regions of perpendicular fibre crossings, where the principal eigenvector of the underlying oblate tensor is meaningless. However, tensor deflection underestimates curvature for curved tracts. Furthermore, TEND results should be interpreted carefully, as an incoming direction that coincides with any of the tensor eigenvectors will not be deflected by that tensor; even if it is perpendicular to the principal eigen-

vector of a highly prolate tensor (Lazar et al., 2003).

All the methods presented in this section are deterministic and provide binary connectivity information. A voxel **B** can be either connected or not connected to the seed **S**, depending on whether a streamline from **S** passes through **B**. To tackle this issue, a group of probabilistic tractography approaches have been developed and these are presented in the following section.

Probabilistic streamline tractography

Despite its success in delineating many major WM tracts (Catani et al., 2002; Mori et al., 2005), deterministic streamline tractography provides no intrinsic way to assign a confidence measure to a reconstructed path. Such a measure is useful, since streamlines are prone to errors caused by experimental noise (Lazar and Alexander, 2003) and potentially by the method used for orientation estimation. Therefore, when a streamline is reconstructed, it would be useful to know how reproducible it is. Unless there is a systematic error, the higher the reproducibility of a streamline, the higher the chance to reflect true structural features of the underlying tissue rather than noise-induced, fictitious pathways.

The ideal way to assess streamline reproducibility would be to repeat the MRI experiment many times $N_R \rightarrow \infty$ and for each of these experiments reconstruct a streamline from the same seed point. Overlapping between these N_R streamlines will give an intrinsic reproducibility measure. Given that experiment repetition is not feasible in practice, probabilistic tractography approaches aim to tackle this issue in an indirect way.

Introduced at the same time by Behrens et al (Behrens et al., 2003) and Parker et al (Parker et al., 2003), probabilistic tractography estimates a spatial distribution of streamlines arising from a single seed rather than a single streamline. From a starting point **S**, N_R streamlines are generated. In each propagation step of each streamline, a random perturbation of the underlying fibre orientation estimate is followed. Perturbations are generated using functions that characterize the uncertainty in the fibre orientation within voxel. Once N_R streamlines are generated in that fashion, a probabilistic index of connectivity (PICO) between the seed **S** and an arbitrary point **B** is

defined as $p = M_R/N_R$, where M_R is the number of streamlines that traverse \mathbf{B} . The number of repetitions N_R should be large enough to achieve a converged value for p . An example of the index of connectivity calculated using a probabilistic approach is shown in Fig. 4.4.




Figure 4.4: Deterministic streamlines (left) and probabilistic index of connectivity (right) for a seed region of interest within the pyramidal tract at the level of the pons. Sagittal sections are shown. Results, generated using the CAMINO toolkit (Cook et al., 2006), are superimposed on a fractional anisotropy image.

The probabilistic framework is also capable of resolving fibre branching configurations, due to the repeated curve propagation, each time with different fibre orientation perturbations. On the contrary, as discussed before, traditional streamline tractography follows a single branch when a branching is encountered.

Probabilistic approaches mainly differ in the way the orientation uncertainty is assessed. A number of methods use a Bayesian framework to fit a model and calculate the posterior probability of the model parameters (Behrens et al., 2003, 2007; Friman et al., 2006; Zhang et al., 2009). A random sample of N_R fibre orientation sets can then be drawn in each voxel from the distribution of the orientation parameters; each orientation set comprising of $N \geq 1$ orientations, depending on the reconstruction method employed. In (Behrens et al., 2003, 2007) the Monte Carlo-Markov chain (MCMC) approach is used to sample the orientation posterior distribution. In (Friman et al., 2006), the posterior is computed numerically, after simplifying its Bayesian expression using dirac priors on the nuisance model parameters. In (Zhang et al., 2009), particle filtering is employed to draw a random sample from the orientation posterior.

An alternative to characterize uncertainty in fibre orientation estimates is bootstrapping (Pajevic and Basser, 2003). The bootstrap method is a non-parametric sta-

tistical procedure that estimates the distribution of a given parameter by randomly drawing samples from a set of repeated measurements. DW acquisitions of M directions are repeated r times (instead of 1 acquisition of M directions). Then M images are drawn randomly, with replacement, from this superset of measurements and fibre orientations are estimated. That will be a single bootstrap sample. Drawing N_R bootstrap samples, will give a distribution for the fibre orientations. Bootstrap tractography (Jones and Pierpaoli, 2005; Lazar and Alexander, 2005) utilizes this framework and for a given seed point generates a streamline per bootstrap sample. Then a visitation index can be defined in a similar fashion to PICO. The advantage of bootstrap tractography is that no ad-hoc assumptions are made on the noise and is sensitive to all sources of variability that affect the acquired dataset. Therefore, effects that cannot be modelled parametrically, such as the effects of physiologic noise and of the scanning system instabilities, will be considered.

Conventional bootstrap requires repeated acquisitions, i.e. a subject needs to be scanned more than once with the same scanning parameters. A recent simulation study showed that at least $r = 5$ repeats are needed for DTI data, so that the estimated distributions of tensor-derived parameters are close to the true ones (O’Gorman and Jones, 2006). Model-based residual bootstrap offers an alternative, since it requires only a single data acquisition (Berman et al., 2008; Haroon et al., 2009; Jones, 2008). A model needs to be fitted to the data and the model residuals are calculated. A single bootstrap sample can then generated by permuting freely the residuals between all model predicted values. The model is then fitted to the new dataset, new residuals are calculated, which are then added randomly to the model predicted values. The process is repeated N_R times to obtain N_R bootstrap samples used subsequently for probabilistic tractography. A special case of residual bootstrap is wild bootstrap (Jones, 2008), where only the sign of the residuals is randomly permuted and each residual is then added back to the respective datum. Wild bootstrap is suitable when the measured data points have different variances (Whitcher et al., 2008).

Apart from Bayesian and bootstrap approaches, other groups use a calibration procedure to empirically predict the orientation uncertainty (Parker and Alexander, 2003, 2005; Parker et al., 2003). In (Parker et al., 2003) the orientation uncertainty in DTI data is determined as a function of the degree of anisotropy FA. Higher FA values mean

higher confidence on the principal eigenvector of the tensor. In (Parker and Alexander, 2003), Monte-Carlo simulations are utilized to predict the orientation uncertainty for different tensor shapes and different number of tensors within a voxel. A lookup table is created, indexed by the absolute and relative tensor eigenvalues. At a given SNR, for each table entry, many noisy realizations of the signal are simulated using a single or a two-tensor model. The orientations are estimated for each realization and the set of orientations are used to estimate the parameters of an empirical orientation distribution. A Normal distribution is assumed for the orientation of each tensor in (Parker and Alexander, 2003) and a Watson distribution in (Cook et al., 2004). A similar calibration procedure is described in (Parker and Alexander, 2005), where the PAS function (Jansons and Alexander, 2003) is utilized to resolve fibre crossings. Data are simulated for each entry of the lookup table and the peaks of the PAS function are estimated for each simulated dataset. The mean sharpness of each peak, evaluated by the trace of the peak Hessian, is used to estimate the standard deviation of the -assumed as- normally distributed orientation. In (Seunarine et al., 2007), a Bingham distribution is used for the orientation estimates, obtained as the peaks of the ODF or PAS function, and the anisotropy of the peak is mapped to the distribution parameters. The peak anisotropy is calculated from the eigenvalues of the peak Hessian matrix and its consideration accounts for fanning and bending structures. It should be pointed out that a common requirement for all the above calibration methods is that an estimate of the SNR of the acquired data is needed to drive the simulations.

4.2 Global tractography approaches

A drawback of probabilistic tractography approaches is that the probabilistic index of connectivity decreases with distance from the seed point, due to the cumulative effect of the uncertainties in propagation (Parker et al., 2003). Furthermore, since they are streamline-based, they suffer from additive dispersion error caused by local, noise-induced artifacts.

Global tractography approaches try to overcome these limitations by a) providing indices of connectivity that are not inherently distance-dependent and b) increasing immunity against noise. The latter is achieved by finding the optimal path between

two voxels, according to a global property that contains information along the whole tract; rather than following orientation estimates in a locally greedy sense. Contrary to streamlines, the reconstructed paths are not necessarily tangent at all points to the orientation vector field. However, they satisfy optimally the global criterion, compared to any other path linking the two voxels of interest (Behrens and Jbabdi, 2009).

Towards this direction, Jbabdi et al (Jbabdi et al., 2007) have developed a global Bayesian model to derive the posterior probability of connections between two regions of interest (ROIs). All possible fibre paths between the two ROIs are represented by piecewise cubic splines, each defined by a set of control points. Sampling from the posterior of the paths is then equivalent to sampling from the posterior of the control points. The posterior distribution of the paths depends on local diffusion properties, estimated using Behrens's ball and stick model (Behrens et al., 2003), and also on prior knowledge of existing connectivity between the two ROIs. The path trajectories are compatible with the local fibre orientations in regions with low uncertainty estimates. In regions with high uncertainty, the global connectivity information constrains the local parameter estimation and affects the path sampling. Even if the framework assumes a single fibre bundle connecting the ROI pair, it is promising for tasks such as connectivity-based tissue parcellation.

Front evolution tractography

Front evolution techniques were amongst the first global approaches. Fast marching tractography has been applied to DTI data (Parker et al., 2002) and propagates a front from a seed towards neighbouring voxels. The speed of propagation v_{ij} between voxels \mathbf{i} and \mathbf{j} is maximum when the local fibre orientations \mathbf{e}_1 are collinear with the normal to the front \mathbf{n} and minimum when the orientations are perpendicular to the normal front or to each other, as defined in (Parker et al., 2002):

$$v_{ij}(\mathbf{n}) = \frac{1}{1 - \min(|\mathbf{e}_1(\mathbf{i}) \cdot \mathbf{n}|, |\mathbf{e}_1(\mathbf{j}) \cdot \mathbf{n}|, |\mathbf{e}_1(\mathbf{i}) \cdot \mathbf{e}_1(\mathbf{j})|)} . \quad (4.3)$$

The front expands from the seed neighbours to the next neighbouring nodes with speeds determined by the local fibre orientations. As the front propagates, a time of front arrival can be associated with each visited voxel. Once all image voxels have been

traversed by the front, paths of connection can be obtained going backwards in the map of front arrival times. Starting from an arbitrary voxel, a gradient descent algorithm can find the fastest route back to the seed. A connectivity index can be associated with each path, representing either the weakest link along the path or the agreement between the path tangents and the underlying vector orientation field.

The fast marching algorithm (Sethian, 1996) governs front propagation. Three sets of voxels are defined, voxels already visited by the front, voxels lying just outside the front (the narrow band) and the remaining voxels, as shown in Fig. 4.5. Members of the narrow band are candidates for inclusion in the front, which propagates in each step into the voxel with the smallest arrival time. As mentioned before, earliest arrival times are expected where orientational coherence is high, i.e. in white matter tracts. Fig. 4.5 shows an example of the front arrival time map obtained after seeding in the genu of the corpus callosum.




Figure 4.5: Left: The fast marching algorithm divides available voxels in three sets, the ones already visited by the front (dark grey), the narrow band (light grey) and the away voxels (white). Propagation occurs from a front voxel to a narrow band voxel, which then is also included in the front. Right: An axial slice of the arrival time map obtained with fast marching tractography, after seeding in the genu of the corpus callosum (FA threshold=0.15). Early arrival times occur mainly within the genu.

Fast marching tractography has been found promising in reconstructing known

anatomical connections, using DTI to get fibre orientation estimates (Parker et al., 2002). Branching pathways are inherently considered by the algorithm, due to front propagation towards all directions. To deal with crossing configurations, different propagation speed functions have been employed. In (Staempfli et al., 2006), voxels containing crossings are identified in DTI data using Westin's planar index (Westin et al., 2002). When the speed value to a crossing voxel needs to be calculated, a modified version of Eq. (4.3) is used, with the third eigenvector of the diffusion tensor rather than its principal eigenvector. The latter in crossing regions points anywhere within the crossing plane, while the former is normal to that plane. Terms, such as $1 - |\mathbf{e}_1(\mathbf{i}) \cdot \mathbf{e}_3(\mathbf{j})|$ can assess collinearity between the principal eigenvector of voxel \mathbf{i} and the crossing plane in voxel \mathbf{j} . In a similar attempt to resolve crossings in Q-ball images, Campbell et al (Campbell et al., 2005) utilizes the diffusion ODFs as speed functions. The front propagation speed along a direction \mathbf{n} is evaluated as a normalized ODF value along \mathbf{n} .

Contrary to voxel-to-voxel propagation utilized in fast marching, a front evolution method that allows continuous propagation in both orientational and spatial fields is presented in (Tournier et al., 2003). An empirical, DTI-derived orientation distribution function is used to drive the propagation. In an iterative fashion, child fronts are generated by stepping away from parent points along the orientations of highest probability. These child fronts are merged to form a surface from which new child fronts will be generated.

Within a different framework, Jackowski et al (Jackowski et al., 2005) solve the anisotropic front propagation equation:

$$\|\nabla T(\mathbf{x})\| \cdot v(\mathbf{x}, \nabla T) = 1, \quad (4.4)$$

with T being the front arrival time at point \mathbf{x} and v the anisotropic propagation speed. The equation is solved numerically for T , using an iterative algorithm. The ADC value $(\mathbf{n}^T \cdot \mathcal{D} \cdot \mathbf{n})$, obtained from the diffusion tensor, is weighted by the FA value and is used as a propagation speed along direction \mathbf{n} . In (Fletcher et al., 2007), front propagation is seeded individually from two regions of interest to study the connectivity of the ROI pair. The propagation equation (4.4) is solved twice, once per seed ROI, using diffu-

sion tensor-derived ODFs (i.e. functions of the inverse diffusion tensor) as propagation speeds. For each image voxel \mathbf{x} , the sum of the two arrival times obtained from the two propagations is used as the total minimal cost of the path between the two ROIs, passing through \mathbf{x} .

Graph-based tractography

In more recent tractography studies, a weighted graph representation of the image has been utilized (Iturria-Medina et al., 2007; Lifshits et al., 2009; Zalesky, 2008). Each image voxel becomes a node in the graph and there are arcs connecting pairs of neighbouring voxels. Arcs are assigned weights, that can be representative of any type of tissue structural information, available in the two connected voxels. Anatomical paths are then defined as chains with successive elements being neighbouring voxels. The weights of the arcs comprising a path are used to determine the path strength. The strongest path between any image voxel and a seed can then be identified using algorithms that search the image graph exhaustively. The strength of this strongest path (or functions of it) is treated as a relative connectivity index. Using this graph representation, the tractography task is formulated as a "shortest path" problem (Cormen et al., 2002; Dijkstra, 1959), which can be solved very efficiently and fast.

The graph-based methods found in the literature differ mainly in the way the arc weights are determined. The flexibility in adjusting the weights is another advantage of these methods. In (Zalesky, 2008), the Bayesian framework of (Friman et al., 2006) is employed to estimate within voxel the posterior distribution of the fibre orientation. This distribution is then used to define pseudo-probability terms of observing diffusion from a voxel **A** to a neighbouring voxel **B** and vice versa. The arc weights are functions of these terms. In (Iturria-Medina et al., 2007), weights are functions of both diffusivity and structural terms. The pseudo-probabilities of diffusion between two neighbouring voxels are determined using the diffusion orientation distribution functions (ODFs). The structural terms reflect the probability of two neighbouring voxels belonging to the same tissue type and can be obtained from anatomical image modalities, such as T_1 -weighted images. In (Lifshits et al., 2009), random diffusion displacements are drawn from the Gaussian diffusion propagator obtained from DTI. Displacements start from the centre of a voxel \mathbf{i} and last a specific diffusion time. Each displacement is grouped

to the neighbouring voxel j , for which the angle between the displacement vector and the vector connecting the centres of i and j is minimum. The fraction of displacements ending up at j define a weight for the i to j transition.

Global energy minimization tractography

A different global framework is presented in (Kreher et al., 2008a), where all neuronal pathways and the signal from the whole brain volume are considered simultaneously. Each tract is represented as a chain of cylinders, whose position, length and orientation can change. The method tries to find the set of cylinders that best approximate the underlying white matter bundles. This is achieved by minimizing an energy function, mimicking the polymerization process through which chains of monomers are created in an energetically-favourable way. The energy function depicts a) how well the DW data can be explained by a configuration of cylinders and b) how anatomically plausible the cylinder configuration is, considering that fibre bundles should be smooth and continuous. The process starts with a random configuration and the signal obtained from the cylinders, within a voxel, is simulated via a multi-tensor model. The cylinders are then iteratively adjusted to minimize the difference between the simulated and the measured signal and also fulfil the geometric requirements. Despite the very promising results, the extremely high computation time, the large number of parameters to set and the sensitivity of the results to these parameters are drawbacks of this framework.

A similar approach is presented in (Fillard et al., 2009). White matter tracts are represented as chains of short fibre segments, whose position, orientation and quantity can change. An optimal configuration of the segments is obtained by iteratively minimizing an energy potential, such that chains assimilate to fibres. The potential function encourages segments to align with the principal fibre orientations, as depicted by DW-MRI, and to form long chains of low curvature. It further prevents a chain to end within WM, by increasing the number of segments where necessary. Obtaining a configuration that is optimized across the whole brain allows all tracts to be reconstructed simultaneously.

4.3 Diffusion simulation tractography

Since WM tracking effectively reconstructs the diffusing paths of water molecules, methods that simulate the diffusion of water molecules within the brain tissue or directly solve Fick's second law have been developed.

In (Hagmann et al., 2003; Koch et al., 2002), random walk models are used with DTI data to simulate Brownian motion. Trajectories are generated as sequences of jumps. The direction of each jump is suggested by the local diffusive properties. In (Koch et al., 2002), transitions are allowed between the centres of neighbouring voxels. The probability of each jump is representative of the apparent diffusion coefficient along the direction of the vector connecting the voxel centres. Many random walks are generated from a seed and the number of times a voxel is traversed defines its index of connectivity to the seed. In (Hagmann et al., 2003), the direction of each jump is determined by the deflection $\mathcal{D} \cdot \mathbf{r}$ of a random orientation \mathbf{r} , chosen from a uniform distribution over the unit hemisphere (such that forward jumps occur), by the diffusion tensor \mathcal{D} . The technique resembles tensor deflection tractography, but the incorporation of the random variable \mathbf{r} adds to it a probabilistic context. From a single seed, N_R trajectories can be generated and an index of connectivity can be computed.

In (Batchelor et al., 2001), the diffusion equation is solved within WM through the whole brain volume using a finite element approach. The anisotropic version of the diffusion equation is:

$$\frac{dP}{dt} = \nabla \cdot (\mathcal{D} \nabla P), \quad (4.5)$$

with the tensor \mathcal{D} at different spatial locations being estimated from DTI data and initial conditions $P(0)=1$ at the seed and 0 elsewhere. This equation is solved for P . An isotropic version of Eq. (4.5) is also solved, assuming constant isotropic diffusion through the whole brain ($\mathcal{D} = I$). The ratio of the two solutions P_{aniso}/P_{iso} incorporates the effect of tissue structure and is used as an index of connectivity.

Successive diffusion simulations over the entire brain, starting from a seed, are performed in (Kang et al., 2005). The diffusion equation is solved with \mathcal{D} being estimated from DTI data, in order to determine propagation fronts. A predefined threshold on the solution P is used to divide voxels in "zero" and "non-zero" concentration. A front is defined using the latter voxels. A subset of the front voxels become new roots from

where the diffusion equation is again solved to generate new fronts. The subset is chosen based on the distance travelled (the higher the diffusion, the larger the distance traversed from a root and the more probable the connection will be), the fibre orientation and the anisotropy of the front voxels. Once all voxels are visited by a front, white matter tracts can be reconstructed in a backward fashion.

Tractography by simulating fluid flow through a pressure tensor field is performed in (Hageman et al., 2009). The Navier-Stokes equation is solved using a finite element approach, with \mathcal{D} being used as a pressure tensor. A viscosity term is present in the equation. Viscosity increases and slows down propagation when low anisotropy or low orientational coherence across a neighbourhood is encountered. The solution is a fluid velocity vector field. The most likely connection path between two regions is the path that maximizes a function of the magnitude of the velocity and its gradient.

Many of the methods above, tackle limitations of streamline tractography, such as resolving fibre crossing (Hageman et al., 2009) and branching (Batchelor et al., 2001; Hageman et al., 2009; Kang et al., 2005) using DTI data. However, solving a partial differential equation using a finite element approach increases execution time considerably. Furthermore, it is not always straightforward with these approaches to obtain a connectivity map across the whole brain volume (Hageman et al., 2009; Kang et al., 2005), and there is usually a large number of parameters and thresholds to set.

Part III

Original Research

5

A regularized two-tensor model for resolving fibre crossings

Overview

Fibre crossing configurations cannot be resolved by Diffusion Tensor Imaging (DTI), as has been shown in the previous chapters. Given that routine use of high angular resolution diffusion imaging (HARDI) is still tentative, a regularized two-tensor model to resolve fibre crossings from conventional DTI datasets is presented. To overcome the problems of fitting multiple tensors, a geometrically constrained model that exploits the planar DTI diffusion profile in regions with fibre crossings is utilized. A regularization scheme is applied to the estimates of this model to reduce noise artefacts, which can be significant due to the relatively low number of acquired images. A set of basis directions is used to convert the two tensor model to many models of lower dimensionality. Relaxation labelling is utilized to select from amongst these models those that preserve continuity of orientations across neighbours. Spatial regularization improves the orientation estimates of the two-tensor model in simulations. Crossing

orientation estimates in human data, acquired in less than 6 minutes, agree well with a-priori anatomical knowledge. Apart from orientations, revised fractional anisotropy and mean diffusivity indices are computed.

Contributions of this chapter

- The geometrically constrained two-tensor model of (Peled et al., 2006) is derived and described in detail.
- Regularization of orientation estimates using relaxation labelling is presented. A hybrid minimization approach is also introduced, suitable for reducing the non-linear dimensionality of the two-tensor model fitting.
- Two different methods are presented for regularizing the model selection mask. This identifies voxels where the two tensor model should be fitted.
- Results are shown on simulated and human data. Resolved crossing orientations in regions, where a-priori anatomical knowledge exists, agree well with expected orientations.
- The global effect of regularization is shown with tractography. The regularized approach was successful in tracking lateral callosal tracts using conventional DTI data, while both DTI and Q-ball imaging failed.

Publications

Contributions from this chapter have appeared in the following:

- S. N. Sotiropoulos, C. R. Tench, L. Bai, P. S. Morgan, D. P. Auer and C. S. Constantinescu. A regularised two-tensor model fit to low angular resolution diffusion images, *Proceedings of the ISMRM Annual Meeting*, p. 1861, Toronto, Canada, May 2008.
- S. N. Sotiropoulos, L. Bai, P. S. Morgan, D. P. Auer, C. S. Constantinescu and C. R. Tench. A regularized two-tensor model fit to low angular resolution diffusion images using basis directions, *Journal of Magnetic Resonance Imaging*, 28:199-209, 2008.

5.1 Introduction

Despite its potential, the single tensor DTI model (Basser et al., 1994a) cannot describe diffusion profiles that arise from non-coherent fibres coexisting in the same voxel (Alexander et al., 2001). In cases of fibres crossing, DTI wrongly predicts the orientation of fibres and also underestimates anisotropy (Frank, 2001).

To overcome the limitations of the single tensor model, many different methods have been proposed, as presented in detail in Chapter 3. However, most of them have illustrated results on data acquired with protocols that are non-applicable in clinical practice. Diffusion spectrum imaging requires the measurements on a large 3D Cartesian lattice (Wedeen et al., 2005). HARDI-based methods, including Q-ball imaging (Tuch, 2004), multi-tensor models (Behrens et al., 2007; Kreher et al., 2005; Leow et al., 2009; Tuch et al., 2002) and spherical deconvolution (Anderson, 2005; Dell’Acqua et al., 2007; Tournier et al., 2004), require longer acquisition times than conventional DTI for a full brain coverage. Many sampling directions (usually at least 60) and multiple repeats to achieve high enough SNR are usually employed by these methods. The higher b value utilized in some cases increases sensitivity to subject motion, which is another undesirable feature for a clinical setting. It should be pointed out that the constrained spherical deconvolution (Tournier et al., 2007) has good results when applied to relatively low angular resolution DW data. However, only orientational information can be obtained with this method.

To remedy some of the above problems, a geometrically constrained two-tensor model for resolving fibre crossings was introduced in (Peled et al., 2006). The observation that regions with two crossing fibres have a planar DTI diffusion profile and that fibres lie in this plane (Wiegell et al., 2000) was utilized. Using information provided by the single tensor DTI estimates, the model has only four unknown parameters and, thus, requires a small number of measurements. In (Peled et al., 2006) crossings were resolved using imaging protocols of 31 diffusion gradient directions and low b values (750 s/mm²). A regularization scheme for this model will be introduced here, to decrease sensitivity to noise artefacts that can be significant due to the low angular sampling.

Regularization schemes have been used before to reduce random effects of ex-

perimental noise, as well as numerical errors that might arise in the fitting procedure (Coulon et al., 2004; Poupon et al., 2000; Tench et al., 2002). Furthermore, sampling requirements can be reduced and accuracy improved when the a priori knowledge of diffusion continuity across voxels is imposed. The principal eigenvector of the DTI model (Poupon et al., 2000; Tench et al., 2002), all three orthogonal DTI eigenvectors (Tschumperle and Deriche, 2002) and the principal eigenvector and all diffusivities of the single tensor model (Coulon et al., 2004) have been regularized previously.

In this chapter, spatial regularization of the orientation estimates obtained from the simplified two-tensor model of (Peled et al., 2006) is performed. Using a basis set of directions, the problem of fitting a multi-parameter two-tensor model is transformed to solving many models of lower dimensionality. Relaxation labelling (Rosenfeld et al., 1976; Savadjiev et al., 2006; Tench et al., 2002) is then utilized to choose from amongst these models the ones that preserve continuity of fibre orientations in crossing regions. Revised anisotropy and mean diffusivity measures are computed using the parameters of the two-tensor model. The method is tested using numerical simulations, as well as in-vivo DW images acquired in less than 6 minutes.

5.2 A geometrically constrained two-tensor model

Assuming that up to two fibre populations coexist within a voxel and that there is no diffusion exchange between them, the diffusion signal can be modelled as a mixture of two diffusion tensors (Tuch et al., 2002):

$$S(\mathbf{g}_k) \equiv S_k = S_0 \left[f \exp(-b \mathbf{g}_k^T \mathcal{D}_a \mathbf{g}_k) + (1 - f) \exp(-b \mathbf{g}_k^T \mathcal{D}_b \mathbf{g}_k) \right], \quad (5.1)$$

where S_k , $k = 1 \dots M$ is the signal obtained after the application of a diffusion-sensitizing gradient of b value and direction \mathbf{g}_k , $f \in [0, 1]$ is the volume fraction of the diffusion tensor \mathcal{D}_a and S_0 the intensity of the non-diffusion-weighted image. Tensor \mathcal{D}_i ($i = a, b$) is assumed to have eigenvalues λ_{i1} , λ_{i2} and λ_{i3} and respective eigenvectors \mathbf{e}_{i1} , \mathbf{e}_{i2} and \mathbf{e}_{i3} .

The above model has thirteen free parameters. In (Peled et al., 2006), the following phenomenological assumptions are made to reduce this number:

- a) In crossing regions, the tensors \mathcal{D}_a and \mathcal{D}_b are assumed to be cylindrically symmetric ($\lambda_{a2} = \lambda_{a3}$ and $\lambda_{b2} = \lambda_{b3}$).
- b) Both tensors are assumed to share the same eigenvalues and their secondary eigenvalue is set equal to λ_3 , the third eigenvalue of the DTI tensor. So: $\lambda_{a1} = \lambda_{b1} = \lambda_{ab1}$ and $\lambda_{a2} = \lambda_{a3} = \lambda_{b2} = \lambda_{b3} = \lambda_3$.
- c) The principal orientations of \mathcal{D}_a and \mathcal{D}_b are assumed to lie on the plane spanned by the principal \mathbf{e}_1 and secondary \mathbf{e}_2 eigenvectors of the DTI tensor (see Fig. 3.1). This is motivated by the observation that two-fibre crossings are characterized by oblate (planar) DTI profiles (Tuch et al., 2002; Wiegell et al., 2000), i.e. the eigenvalues of the DTI tensor are $\lambda_1 \approx \lambda_2 \gg \lambda_3$, and the crossing occurs on this plane.

In the coordinate system defined by the DTI eigenvectors $[\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3]$, the two tensors can be described as:

$$\tilde{\mathcal{D}}_a = \begin{bmatrix} d_{a1} & d_{a3} & 0 \\ d_{a3} & d_{a2} & 0 \\ 0 & 0 & \lambda_3 \end{bmatrix} \quad \text{and} \quad \tilde{\mathcal{D}}_b = \begin{bmatrix} d_{b1} & d_{b3} & 0 \\ d_{b3} & d_{b2} & 0 \\ 0 & 0 & \lambda_3 \end{bmatrix}, \quad (5.2)$$

since the third eigenvectors $\tilde{\mathbf{e}}_{a3}$ and $\tilde{\mathbf{e}}_{b3}$ are aligned with \mathbf{e}_3 . Furthermore, the principal eigenvector of each tensor in this coordinate system can be described using one angle relative to \mathbf{e}_1 (Fig. 5.1), i.e. $\tilde{\mathbf{e}}_{a1} = [\cos \phi_a \sin \phi_a 0]^T$ and $\tilde{\mathbf{e}}_{b1} = [\cos \phi_b \sin \phi_b 0]^T$. The secondary eigenvectors will be $\tilde{\mathbf{e}}_{a2} = [\cos(\phi_a + \pi/2) \sin(\phi_a + \pi/2) 0]^T = [-\sin \phi_a \cos \phi_a 0]^T$ and $\tilde{\mathbf{e}}_{b2} = [-\sin \phi_b \cos \phi_b 0]^T$. Then, for each tensor $i = a, b$, we can apply its eigen-decomposition relationship to get:

$$\tilde{\mathcal{D}}_i = \tilde{\mathbf{E}}_i \mathbf{L}_i \tilde{\mathbf{E}}_i^T \Rightarrow \tilde{\mathcal{D}}_i = \begin{bmatrix} \cos \phi_i & -\sin \phi_i & 0 \\ \sin \phi_i & \cos \phi_i & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} \lambda_{ab1} & 0 & 0 \\ 0 & \lambda_3 & 0 \\ 0 & 0 & \lambda_3 \end{bmatrix} \begin{bmatrix} \cos \phi_i & \sin \phi_i & 0 \\ -\sin \phi_i & \cos \phi_i & 0 \\ 0 & 0 & 1 \end{bmatrix}. \quad (5.3)$$

From Eqs. (5.2), (5.3), we obtain for $i = a, b$:

$$\begin{aligned} d_{i1} &= \lambda_{ab1} - d_{i3} \tan \phi_i \\ d_{i2} &= \lambda_{ab1} - d_{i3} / \tan \phi_i \\ d_{i3} &= (\lambda_{ab1} - \lambda_3) \sin \phi_i \cos \phi_i. \end{aligned} \quad (5.4)$$




Figure 5.1: Schematic showing the principal orientations \mathbf{e}_{a1} and \mathbf{e}_{b1} of the two tensors \mathcal{D}_a and \mathcal{D}_b respectively, lying on the plane defined by the principal \mathbf{e}_1 and secondary \mathbf{e}_2 eigenvector of the DTI tensor (or the plane that is normal to DTI \mathbf{e}_3).

Since $[\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3]$ has been used as the coordinate system of reference, the gradient directions \mathbf{g}_k are also transformed into this coordinate system. The new gradient vectors $\tilde{\mathbf{g}}_k$ are:

$$\tilde{\mathbf{g}}_k = \mathbf{R} \mathbf{g}_k, \text{ with } \mathbf{R} = \begin{bmatrix} \mathbf{e}_1^T \\ \mathbf{e}_2^T \\ \mathbf{e}_3^T \end{bmatrix}. \quad (5.5)$$

Then, the geometrically constrained two-tensor model of (Peled et al., 2006) is given by:

$$S_k = S_0 \left[f \exp(-b \tilde{\mathbf{g}}_k^T \tilde{\mathcal{D}}_a \tilde{\mathbf{g}}_k) + (1 - f) \exp(-b \tilde{\mathbf{g}}_k^T \tilde{\mathcal{D}}_b \tilde{\mathbf{g}}_k) \right]. \quad (5.6)$$

This model has only four parameters: the volume fraction f of the first tract, the common principal diffusivity in both tracts λ_{ab1} , and the angles ϕ_a and ϕ_b , relative to the principal eigenvector \mathbf{e}_1 of the DTI tensor, in the plane spanned by \mathbf{e}_1 and \mathbf{e}_2 of the two tensors. The parameters can be estimated by minimizing the sum of squared residuals $E = \sum_k (S_k - S'_k)^2$, where S'_k is the model predicted and S_k the measured signal for the k^{th} diffusion-sensitizing gradient. Since the problem is non-linear, Nelder and Mead's downhill simplex algorithm (Press et al., 1992) was used here to minimize E . It should be pointed out that the above model provides tensor estimates in the transformed coordinate system. These tensors should be reoriented to the scanner coordinate system, by multiplying their eigenvectors with \mathbf{R}^{-1} , i.e. $\mathbf{e}_{a1} = \mathbf{R}^{-1} \tilde{\mathbf{e}}_{a1}$.

Similar to all multi-tensor models, a model selection problem arises. Westin's planar index c_p (Westin et al., 2002), calculated from the eigenvalues of the DTI tensor was

used in (Peled et al., 2006) to identify voxels with an oblate diffusion profile. The more perpendicular a crossing and the more equal the volume fractions of the two crossing populations, the more oblate the diffusion profile and the more suitable the model will be. Thus, high values of c_p , in general encourage fitting of the two-tensor model. In (Peled et al., 2006) a minimum threshold of $c_p > 0.2$ was identified using simulations, as fitting the model to configurations with smaller c_p values, produced estimates with larger than 10° orientation errors.

5.3 Spatial regularization using relaxation labelling

Noise artefacts and potential inaccuracies of any model in reflecting the physical reality can introduce errors in the estimates obtained individually in each voxel. The geometrically-constrained two-tensor model is no exception; furthermore, its non-linear nature makes the identification of an estimate that is a global minimum of the objective function E more difficult. However, continuity of diffusive processes across voxels requires that the dominant diffusion orientations are preserved in highly structured media, such as white matter tracts. This low curvature assumption can be imposed using a spatial regularization scheme on the orientation estimates across neighbourhoods of voxels. Such a scheme ensures that local artefacts are corrected, by requiring that local estimates are commensurate with their adjacent ones.

Relaxation labelling (RL) has been employed here for this purpose. Introduced by Rosenfeld et al (Rosenfeld et al., 1976), RL is an iterative method that assigns to a voxel (object in general) \mathbf{i} one of L predefined labels. The assignment of a label $l = 1 \dots L$ is characterized by the weight $P_i(l)$, which reflects how likely is for label l to be the most appropriate amongst all L labels for voxel \mathbf{i} , given the data. It is based on the support $dP_i(l)$ that a label l gets from all voxels belonging to a neighbourhood of \mathbf{i} , F_i^{neigh} . The support function reflects the regularity constraint to be preserved. Using the support values, the weights P can be iteratively modified until convergence. The weights at each iteration $n + 1$, are:

$$P_i^{n+1}(l) = \frac{P_i^n(l)[1 + dP_i^n(l)]}{\sum_{q=1}^L P_i^n(q)[1 + dP_i^n(q)]}, \quad l = 1 \dots L. \quad (5.7)$$

The initial weights $P_i^0(l)$ define a starting point based on the knowledge prior to regularization. In each voxel \mathbf{i} , the label with the highest weight is assumed to be the dominant during iteration n .

Regularizing the orientation vector field

Since it is expected that the fibre orientations are slowly changing across neighbouring voxels, any abrupt change is assumed to be a noise artefact. Therefore, relaxation labelling can be applied to the estimated orientations to preserve directional continuity. When using the geometrically-constrained two-tensor model, two orientations will be estimated in crossing regions. For the remaining voxels, the diffusion tensor model (Basser et al., 1994a) will estimate a single orientation. Given that fitting the non-linear two-tensor model is more difficult than the linear DTI model, RL will be applied in voxels containing crossings to improve the estimates of the two-tensor model.

Labels can be chosen as pairs of orientations $l_{\alpha\beta} = \{\mathbf{u}_\alpha, \mathbf{u}_\beta\}$, $\alpha, \beta = 1 \dots N_B$, with \mathbf{u} being unit vectors chosen from a set Φ of N_B basis directions. Using relaxation labelling, the orientation pair that is most compatible with the orientations in neighbouring voxels will be chosen, as shown in Fig. 5.2.




Figure 5.2: Schematic showing the spatial regularization of orientations in crossing regions. Orientation pairs that are most compatible with the neighbouring orientations are selected.

The geometrically-constrained model defines orientations only in the 2D plane, within which crossing occurs. The basis directions can therefore be restricted to a plane and reduced to a set of N_B basis angles $\Phi = [-\pi/2, -\pi/2 + \delta\omega, -\pi/2 + 2\delta\omega, \dots, \pi/2)$ with $0 < \delta\omega < \pi/2$ defining the fineness of the set. The labels will then be $l_{\alpha\beta} =$

$\{\phi_\alpha, \phi_\beta\}$ with $\alpha = 1 \dots N_B$ and $\beta = \alpha + 1 \dots N_B$, giving in total $L = (N_B^2 - N_B)/2$ labels. The above notation avoids pairs of equal angles and also considers only once the identical $l_{\alpha\beta}$ and $l_{\beta\alpha}$ labels. The orientations corresponding to a label $l_{\alpha\beta}$ will be $\tilde{\mathbf{u}}_\alpha = [\cos \phi_\alpha \ \sin \phi_\alpha \ 0]^T$ and $\tilde{\mathbf{u}}_\beta = [\cos \phi_\beta \ \sin \phi_\beta \ 0]^T$ in the local coordinate system $[\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3]$ of a voxel or $\mathbf{R}^{-1}\tilde{\mathbf{u}}_\alpha$ and $\mathbf{R}^{-1}\tilde{\mathbf{u}}_\beta$ in the scanner's coordinate system (Eq. (5.5)).

A low fibre curvature requirement can then be imposed within a neighbourhood using a support function that promotes slow changes in these two planar orientations. Such a function for a voxel \mathbf{i} and a label l is defined below:

$$dP_i^n(l) = \sum_{j \in F_i^{neigh}} \frac{1}{d_{ij}} T_{ij}^n(l) \quad (5.8)$$

$$T_{ij}^n(l) = \begin{cases} \exp \left(- \left(\frac{\min[\Theta(\mathbf{e}_j, \mathbf{e}_{ial}), \Theta(\mathbf{e}_j, \mathbf{e}_{ibl})]}{\sqrt{2}\sigma_3} \right)^2 \right), & \text{if } \mathbf{j} \text{ is not a planar neighbour} \\ \exp \left(- \left(\frac{\Delta\phi_l - \Delta\phi_j^n}{\sqrt{2}\sigma_1} \right)^2 - \left(\frac{\Theta(\mathbf{v}_{il}, \mathbf{v}_j^n)}{\sqrt{2}\sigma_2} \right)^2 \right), & \text{if } \mathbf{j} \text{ is a planar neighbour,} \end{cases}$$

where d_{ij} is the Euclidean distance of the centres of voxels \mathbf{i} and \mathbf{j} , and \mathbf{e}_j is the principal DTI eigenvector of neighbour $\mathbf{j} \in F_i^{neigh}$. $\Delta\phi_l$ is the angle separation between the orientations of label l , \mathbf{v}_{il} , \mathbf{e}_{ial} and \mathbf{e}_{ibl} are respectively the mean vector of orientations and the two orientations of label l rotated back to scanner's coordinate system from the coordinate system $[\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3]$ of voxel \mathbf{i} . $\Delta\phi_j^n$ and \mathbf{v}_j^n are calculated at each iteration and represent the angle separation and the mean orientation vector of the label with the highest weight in neighbouring voxel \mathbf{j} at the n^{th} iteration; \mathbf{v}_j^n is rotated to the scanner's coordinate system. Finally, function $\Theta(\mathbf{u}, \mathbf{v})$ computes the minimum angle subtended by any combination of vectors \mathbf{u} and $-\mathbf{u}$ with \mathbf{v} and $-\mathbf{v}$. It is defined as $\Theta(\mathbf{u}, \mathbf{v}) = \min[\text{acos}(\mathbf{u} \cdot \mathbf{v}), \pi - \text{acos}(\mathbf{u} \cdot \mathbf{v})]$.

The above support function is a Gaussian to which every neighbouring voxel $\mathbf{j} \in F_i^{neigh}$ contributes. When \mathbf{j} does not have a planar DTI profile, then a single tensor is fitted to it, otherwise the two-tensor model is applied. In the former case, the contribution is peaked when the principal DTI eigenvector of \mathbf{j} is equal to any of the two orientations of label l in \mathbf{i} . In the latter case, a maximum contribution is achieved when the two dominant orientations of \mathbf{j} are equal to the two orientations of label l . To test that we need four comparisons, since the orientation pairs are not sorted. A more

efficient way is shown in Eq. (5.8), where the similarity in the angle separation and the mean vector of the two orientations is instead required to be preserved between neighbours.

Parameters σ_1 , σ_2 and σ_3 determine how much variation is allowed in the neighbourhood. We set $\sigma_3 = 10^\circ = 0.1745$ rads to define variation between neighbouring orientations. Therefore, the allowed variation in the mean orientation will be $\sigma_2 = \sigma_3 = 10^\circ = 0.1745$ rads and in the angle separation $\sigma_1 = 2\sigma_3 = 20^\circ = 0.349$ rads.

The initial weights $P_i^0(l)$ for a voxel \mathbf{i} and label l are calculated from the residuals of the two-tensor model fit. The model described by Eq. (5.6) is fitted but with fixed ϕ_a and ϕ_b , given by the corresponding basis angles of the label l (thus only two parameters are unknown). Fitting of all L labels gives L residuals E_{il} for voxel \mathbf{i} . The smaller the E_{il} , the higher the respective weight of the label. So $P_i^0(l) = 1/(\mathcal{Z} \cdot E_{il})$, with \mathcal{Z} being a normalization constant so that the sum of $P_i^0(l)$ over all L labels is unit.

Since fitting L non-linear models in all voxels with oblate DTI diffusion profile can be computationally expensive, a simpler approach is employed. This is based on the observation that out of the two free parameters f and λ_{ab1} , the fraction f is linearly dependent to the measurements, when keeping the other parameter fixed. Therefore, for a given estimate of λ_{ab1} , an estimate of f can be obtained using a linear least squares scheme. Eventually, L models with non-linear dependence only on the parameter λ_{ab1} are fitted using the hybrid approach described below.

Fitting basis two-tensor models using 1D non-linear optimization

The geometrically-constrained two-tensor model (Eq. (5.6)) can be rewritten as:

$$S_k = S_0 \left[f \left(\exp(-b\tilde{\mathbf{g}}_k^T \tilde{\mathcal{D}}_a \tilde{\mathbf{g}}_k) - \exp(-b\tilde{\mathbf{g}}_k^T \tilde{\mathcal{D}}_b \tilde{\mathbf{g}}_k) \right) + \exp(-b\tilde{\mathbf{g}}_k^T \tilde{\mathcal{D}}_b \tilde{\mathbf{g}}_k) \right] \implies$$

$$S_k = S_0 [f(c_{ka} - c_{kb}) + c_{kb}] , \quad (5.9)$$

with $c_{ka} = \exp(-b\tilde{\mathbf{g}}_k^T \tilde{\mathcal{D}}_a \tilde{\mathbf{g}}_k)$ and $c_{kb} = \exp(-b\tilde{\mathbf{g}}_k^T \tilde{\mathcal{D}}_b \tilde{\mathbf{g}}_k)$ depending on the parameters ϕ_a , ϕ_b and λ_{ab1} . The angles are fixed to ϕ_α and ϕ_β of the label $l_{\alpha\beta}$. If we keep constant the λ_{ab1} , then the volume fraction f is linearly dependent on the measurements. Thus, its value that minimizes the sum of squared residuals E can be obtained using linear

least squares on Eq. (5.9). Having a design matrix with 1 column and elements $A_{k1} = c_{ka} - c_{kb}$ and observations vector with elements $S_k/S_0 - c_{kb}$, the least squares estimate of f will be:

$$f = \frac{\sum_{k=1}^M S_k(c_{ka} - c_{kb}) - S_0 \sum_{k=1}^M c_{kb}(c_{ka} - c_{kb})}{S_0 \sum_{k=1}^M (c_{ka} - c_{kb})^2}. \quad (5.10)$$

For a given least squares estimate of f , the parameter λ_{ab1} can be estimated using a 1D non-linear minimization method. Brent's method (Press et al., 1992) has been employed to search for the value of λ_{ab1} that minimizes the sum of squared residuals E . The new eigenvalue estimate will give a new estimate for the volume fraction, which will give a new estimate for λ_{ab1} and so forth. This iterative process is continued until convergence is achieved (relative error of less than 10^{-3} in both f and λ_{ab1}). As an initial condition used in the first iteration, we set $\lambda_{ab1} = \lambda_1$, the DTI principal eigenvalue.




Figure 5.3: Initial weights $P^0(l_{\alpha\beta})$ of different basis models fitted using the 1D hybrid approach, for a 90° (left) and a 60° (right) crossing. True orientations are shown with blue lines and the angles relative to the DTI principal eigenvector \mathbf{e}_1 are indicated. These angles are predicted well by the label of highest weight. Simulations were noise-free.

This hybrid minimization approach reduces L 2D non-linear regressions to L 1D,

which are faster to solve. Examples of fitting all the L basis models and their residuals E_l are presented in Fig. 5.3. A 90° and a 60° crossing have been simulated using the multi-tensor model of Eq. (5.1) (with each tensor having an FA=0.8, a trace of 2.1×10^{-3} mm²/s and a volume fraction $f=0.5$). The crossing orientations are indicated by the blue lines. All L models, chosen from a basis set Φ with $\delta\omega = 10^\circ$, were fitted using the hybrid approach. For each model, the weight $P^0(l)$ was proportional to the inverse of value of E achieved after fitting (so a high $P^0(l)$ value implied a good fit). We can observe that in both cases, the highest weight labels, $\{80^\circ, -10^\circ\}$ and $\{30^\circ, -30^\circ\}$ respectively, predict well the crossing orientations. Furthermore, the other labels are suppressed, having effectively a zero weight, apart from a few that suggest orientations close to the true ones. Note that, as suggested by the model, the basis angles are defined relative to the DTI principal eigenvector \mathbf{e}_1 .

5.4 Regularizing the model selection mask

Voxels are classified as planar or non-planar via a model selection procedure. Two and one tensors are fitted in the planar and non-planar voxels, respectively. In (Peled et al., 2006), a thresholded planar index ($c_p > 0.2$), calculated from the DTI eigenvalues, was used to identify planar voxels. However, such a procedure is very sensitive to noise. Thus, regularization was applied on the c_p values to increase the robustness of the model selection mask. Two approaches were tried and described below: a) c_p regularization using relaxation labelling and b) c_p non-linear filtering using the edge-preserving Perona and Malik filter (Perona and Malik, 1990).

Relaxation labelling on c_p

The relaxation labelling framework was utilized to regularize the c_p model selection mask. The underlying idea was that voxels surrounded by neighbours with a planar DTI profile are likely to have a planar DTI profile themselves. Two labels ($L = 2$) were defined (*planar* vs *non – planar*) for each voxel \mathbf{i} with initial weights determined by

the $c_{p,i}$ value obtained from the DTI model fit:

$$\begin{aligned} P_i^0(planar) &= 0.5 + \frac{c_{p,i} - thresh}{c_p^{max} - thresh} \times 0.5 \\ P_i^0(non - planar) &= 1 - P_i^0(planar) . \end{aligned} \quad (5.11)$$

According to these equations, voxels with $c_p = thresh$ are considered 50% planar, with this weight increasing as the difference $c_p - thresh$ increases. The parameter $c_p^{max} > thresh$ controls how fast the weight becomes 1 as $c_p - thresh$ goes up, and represents the maximum c_p value. In this study, $thresh = 0.2$ and $c_p^{max} = 0.4$, the latter representing the 95th percentile of c_p values in white matter. For voxels with $c_p > 0.4$, $P_i^0(planar)$ becomes greater than 1 and is thus clipped to 1.

The support function of a label was defined as the sum of the weights of that label in a 3x3x3 neighbourhood. After iterating, the label with the highest weight was assigned to a voxel, resulting into a binary model-selection mask.

The above procedure is not specifically designed to preserve edges between crossing regions and regions containing single-fibre configurations. Under these conditions, contributions from planar and non-planar neighbours compete to define the most appropriate label of a voxel. Thus, an approach that avoids smoothing over edges was also tried.

Non-linear filtering of c_p

The Perona and Malik filter (Perona and Malik, 1990) performs non-linear smoothing in an iterative way that preserves edges. The evolution of the image intensity I at location \mathbf{x} , is described by:

$$\frac{\partial I}{\partial t}(\mathbf{x}) = \nabla \cdot [g(\|\nabla I(\mathbf{x}, t)\|) \nabla I(\mathbf{x}, t)] , \quad (5.12)$$

with $\nabla I(\mathbf{x}, t)$ being the spatial gradient of the image intensity at location \mathbf{x} and time t . The conductance function $g(\cdot)$ was set to $g(s) = \exp[-(s/K)^2]$, as suggested in (Gerig et al., 1992; Perona and Malik, 1990). Finite spatial differences of the image intensity were utilized to discretize the above equation, as shown in (Gerig et al., 1992; Perona

and Malik, 1990):

$$I^{t+\Delta t}(\mathbf{x}) \approx I^t(\mathbf{x}) + \Delta t \cdot \sum_{j \in F_i^{neigh}} \left[\frac{1}{d_{xj}^2} g(|I^t(\mathbf{j}) - I^t(\mathbf{x})|) (I^t(\mathbf{j}) - I^t(\mathbf{x})) \right], \quad (5.13)$$

with \mathbf{j} denoting the neighbouring voxels of \mathbf{x} and d_{xj} the relative distance between the centres of voxels \mathbf{x} and \mathbf{j} . A 3x3x3 neighbourhood was used in 3D (brain data) and a 3x3x1 neighbourhood in 2D (simulated phantoms). Parameter K was obtained as the 85th percentile of the gradient amplitudes over the whole image in each iteration (Gerig et al., 1992; Perona and Malik, 1990). The filter was applied on the c_p image, i.e. $I = c_p$. The initial c_p image was convolved with a 3D Gaussian kernel of $\sigma = 1$, to eliminate large gradient amplitudes caused by noise and discriminate between them and intensity gradients caused by actual edges (Perona and Malik, 1990). The time steps utilized had the maximum values suggested in (Gerig et al., 1992) ($\Delta t = 0.14$ for 2D and $\Delta t = 0.0682$ for 3D).

The above filter iteratively smooths the image, where the gradient magnitude $\|\nabla I\|$ is in the order of K . For much larger gradient values (greater than 2-3 times K), it assumes that an edge exists and no smoothing occurs. Thus, it can potentially preserve edges. After applying the filter to the c_p image, a threshold of 0.2 was used to get a binary model selection mask.

5.5 Data and processing

Simulations

Voxels containing two fibre orientations were numerically generated using Eq. (5.1). Tensors were cylindrically symmetric and had a trace of $2.1 \times 10^{-3} \text{ mm}^2/\text{s}$ and an $\text{FA}=0.8$, representative of highly anisotropic regions (Pierpaoli et al., 1996). Thirty-two directions, uniformly distributed on the surface of a unit sphere (see Appendix), with $b=1000 \text{ s/mm}^2$ were used to generate the noise-free DW signals.

The regularization scheme was validated using a numerical phantom that contained two crossing circular arcs. The crossing tracts had the afore-mentioned eigenvalues and the volume fraction was 0.5. Zero-mean Gaussian noise was added in

quadrature, so that the signal to noise ratio (SNR) was equal to the average non-DW image intensity divided by the standard deviation of noise (Kingsley, 2006a). The non-regularized geometrically- constrained model and the regularized approach were compared at different SNR levels and at different number of diffusion-sensitizing directions. For a given SNR and direction scheme, fifty phantom realizations were obtained, with the two-tensor model being fitted to voxels identified as planar and a single tensor model being used for the remaining ones.

In-vivo images

A whole-brain data set was acquired, with local research ethics committee approval, from a healthy male subject who gave informed consent. Scans were performed using a single-shot, echo-planar, diffusion-weighted sequence (acquisition matrix 112x112 with in-plane resolution $2 \times 2 \text{ mm}^2$, interpolated during reconstruction to 224×224 , TE=60 ms, TR=9500 ms) in a Philips 3T Achieva clinical imaging system (Philips Medical Systems, Best, The Netherlands). A parallel imaging factor of 2 was used. Three non-DW images were acquired and averaged, while diffusion weighting was applied along 32 uniformly distributed directions with $b=1000 \text{ s/mm}^2$. 52 slices were acquired with a thickness of 2 mm. Total imaging time was less than 6 minutes.

Images were corrected for eddy current distortion using FSL (Smith et al., 2004). Brain was extracted using BET (Smith et al., 2004). A binary mask indicating brain voxels was then eroded to remove any boundary voxels that survived skull-stripping. A single tensor model was fitted using a weighted least squares approach (Basser et al., 1994a). A k-means clustering was then applied to the Mean Diffusivity (MD) image to classify voxels into three clusters. Regions belonging to the cluster with the highest MD were excluded from further analysis, since they were considered cerebrospinal fluid (CSF). From the other two clusters, which included brain tissue and brain/CSF interfaces, highly isotropic ($FA < 0.1$) voxels were further excluded. For the remaining voxels, which were primarily within white matter, the planar index was regularized. The L 1-parameter two-tensor models were fitted to all planar voxels. Relaxation labelling was then used to regularize the two orientations in regions of fibre crossing. A step size $\delta\omega = 10^\circ$ was used for the basis set resulting in $N_B=18$ basis directions and $L=153$ labels. Relaxation labelling was stopped after thirty iterations, as simulations

showed that the improvements were mainly obtained within the first ten iterations. Processing time for the whole brain volume was in the order of 2 hours for a 3.2 GHz PC.

Once the orientation estimates were computed and regularized, streamline tractography within the multi-tensor field was performed using CAMINO (Cook et al., 2006). An FA of less than 0.15 and a curvature threshold of 45° were used as stopping criteria. For comparison, tractography using DTI estimates, non-regularized two-tensor estimates, as well as Q-ball estimates was performed. Q-ball ODFs were reconstructed using spherical harmonics as basis (Descoteaux et al., 2007; Hess et al., 2006) (harmonics up to 4^{th} order were fitted).

5.6 Results

Simulations

To test how well the hybrid minimization approach works, it was tried on various crossing configurations. Estimated fibre orientations are shown in Fig. 5.4, for different simulated scenarios. No noise was added in these simulations to identify problems caused by the non-linear model fitting. A basis set with $\delta\omega = 2.5^\circ$ was used resulting to $N_B=72$ basis directions and $L=2556$ models (labels). All these models were fitted in each case and the one with the smallest residual error E was chosen. The orientations of this best model are shown as gray bars and are superimposed with the true orientations. A maximum orientation error of 2.63° and a maximum volume fraction error of 9% were observed in the low volume fraction ($f = 0.2$) configurations, caused by the discreteness of the basis set and the non-linear fitting procedure. The maximum error in the principal eigenvalue was 4.4%. Fitting the 4-parameter geometrically constrained model in each case (using downhill simplex (Press et al., 1992) with multiple restarts), gave comparable maximum errors in orientation, volume fraction and principal eigenvalue of 1.18° , 7.5% and 3.9%, respectively.

The approaches for regularizing the model selection mask were tried in a phantom with two curved, crossing tracts. Fig. 5.5 shows the results obtained using non-linear smoothing and relaxation labelling on the c_p values of a phantom realization at




Figure 5.4: Estimated orientations (gray bars) using the 1-parameter hybrid minimization approach superimposed on true orientations (black lines). Results for different angle separations (60° - 90°) and volume fractions f (0.2-0.5) are shown. In all cases presented, c_p was greater than 0.2.

SNR=10. We can observe that simple thresholding, as performed in (Peled et al., 2006), induces many false positives compared to the noise-free mask. Both smoothing techniques eliminated most of the noise artefacts and resulted in a mask located within the crossing region. The non-linear filtering exhibited slightly better behaviour across the edges and will be therefore used from now on.

The regularization of fibre orientations was then tested against noise, as shown in Fig. 5.6. For different SNRs, the non-regularized two-tensor model (Peled et al., 2006) and the regularized approach were compared. Dark gray indicates regions of fibre crossings, as identified by the model selection procedure, where a two-tensor model was applied. A thresholded c_p mask and a regularized c_p mask were utilized for model selection by the non-regularized and the regularized approaches, respectively. The smooth orientation change imposed by relaxation labelling is evident and restores the spatial continuity of orientation estimates in the crossing region, even for low SNR conditions.

Fig. 5.7 presents the evolution of label weights during regularization. The weights in a voxel of the phantom's crossing region (Fig. 5.6) are plotted for different number of iterations. The SNR was 10 and $\delta\omega = 10^\circ$. As shown in the figure, the dominant label for that voxel at iteration $n=0$ was $l_i=(-70^\circ, 10^\circ)$. However due to noise, the weight of l_i was only slightly higher than the weights of other labels. Furthermore, l_i was




Figure 5.5: Regularizing the c_p -based model selection mask using non-linear filtering and relaxation labelling. A crossing was simulated at SNR=10. Both techniques returned a binary mask, indicating voxels where two tensors should be fitted. The regularized masks resembled the noise-free one, as opposed to the mask obtained from simple c_p thresholding. The evolution of the masks after n filtering iterations is shown. For the non-linear filter, the raw c_p values are also shown at different iterations. At $n=30$ both filters had converged.

not the correct, noise-free label. Imposing the neighbourhood information, allowed the weight of the correct label $l_c = (-30^\circ, 40^\circ)$ to increase after the first iteration of relaxation labelling. After three iterations l_c became the dominant label and l_i was totally suppressed. Subsequently, the weight of l_c increased towards unity, while the weights of all other labels decreased to zero.




Figure 5.6: Effect of regularization on planar index and orientation estimates. A numerical phantom of two crossing circular arcs was used. The geometrically constrained two-tensor model and the regularized two-tensor approach are compared at various noise levels. Dark gray indicates voxels identified as planar in each case. For the regularized cases, c_p was non-linearly filtered. Thirty iterations were used for relaxation labelling, applied to the orientations ($\delta\omega = 2.5^\circ$).




Figure 5.7: Evolution of label weights during relaxation labelling. A voxel within the phantom's crossing region (Fig. 5.6) was chosen. The initial label $l_i = (-70^\circ, 10^\circ)$ is very sensitive to perturbations, as its weight is slightly higher than the weights of other labels. Imposing the neighbourhood information, makes another label $l_c = (-30^\circ, 40^\circ)$ dominant and suppresses the others. This is because l_c promotes orientational coherence across the neighbourhood, more than any other label.

Quantification of improvement due to regularization

A quantitative evaluation of the improvement achieved with regularization is presented in Fig. 5.8. Monte-Carlo simulations were performed and random noise at given SNR levels was added to the crossing phantom. For each phantom realization, three schemes were used, the non-regularized two-tensor model and the regularized two-tensor approach with step sizes $\delta\omega = 10^\circ$ and $\delta\omega = 2.5^\circ$. Errors in orientation, volume fraction and principal eigenvalue of all voxels identified as planar were averaged. Whisker plots that show the distribution of these mean errors across fifty phantom realizations are presented. In all cases, the regularized schemes improved the orientation estimates. Improvements were larger for smaller step sizes. Using the regularized orientation estimates to refit the model to the data, revised estimates for the volume fraction and the principal eigenvalue were obtained. The revised volume fractions were also improved, while no significant difference was observed for the eigenvalues.




Figure 5.8: Whisker plots of the mean errors across the crossing region of the phantom in (a) orientation, (b) volume fraction and (c) principal eigenvalue. For each SNR level, 50 experiments were performed. Results in each SNR case refer from left to right to the non-regularized two-tensor model, the regularized two-tensor approach with $\delta\omega = 10^\circ$ and the regularized approach with $\delta\omega = 2.5^\circ$. The lines of each box correspond to the lower quartile, the median and the upper quartile of a sample, while the whiskers represent the range of the sample.

The effect of regularization on orientation estimates was also assessed against the number of diffusion-sensitizing directions, at different SNR levels (Fig. 5.9). As before, the non-regularized two-tensor model was compared with the regularized approach with step sizes 10° and 2.5° . Uniformly distributed directions were generated on the surface of a sphere, as described in the Appendix. For each direction scheme and SNR,

the orientation errors from fifty phantom realizations were averaged. For each realization, the orientation error was the mean across the crossing region. We can observe that for 15-32 directions, regularization halved orientation errors, while the 24 and 32 direction schemes provided a good compromise between acquisition time (4-6 minutes) and performance (mean error 5°) for a reasonable SNR of 20. For large SNRs and/or number of directions the benefit from orientation regularization becomes small.




Figure 5.9: Mean orientation error as a function of the number of diffusion-sensitizing directions (6-64) and SNR (10-40). For each direction scheme and SNR level, 50 simulations were performed using the crossing phantom. Mean orientation errors across these simulations were obtained using the non-regularized two-tensor model (No-RL, black lines), the regularized approach with $\delta\omega = 10^\circ$ (RL-10°, light gray lines) and the regularized approach with $\delta\omega = 2.5^\circ$ (RL-2.5°, dark gray lines). Uniformly distributed directions were produced as described in the Appendix.

In-vivo human data

Fig. 5.10 shows the regularization effect on the c_p calculated from in-vivo DW images. Fig. 5.10b presents a thresholded c_p mask ($c_p > 0.2$), as suggested in (Peled et al., 2006) and used by the non-regularized approach. The effect of relaxation labelling and non-linear smoothing is shown in Fig. 5.10c and 5.10d, respectively. The regularized binary masks resemble the thresholded c_p image, but has fewer discontinuities. Furthermore, the non-linear smoothing is slightly more conservative at the boundaries.

The orientation estimates obtained from the regularized two-tensor approach were compared against a-priori anatomical knowledge, in regions known to consist of crossing tracts (Fig. 5.11). The orientations are shown colour-coded, with the colour scheme shown in Fig. 5.11d. The length of each orientation vector is scaled by the single tensor FA, in regions with one fibre orientation and the FA computed from the revised eigenvalues of the two-tensor model in crossing regions. In Fig. 5.11a, the superior




Figure 5.10: Effect of regularization on planar index c_p . (a) Raw c_p axial image, (b) Thresholded c_p image ($c_p > 0.2$), (c) Binary mask from regularized c_p using relaxation labelling (after 50 iterations), (d) Binary mask from regularized c_p using non-linear smoothing (after 30 iterations)

part of the corpus callosum, with a high medio-lateral and superior-inferior orientation is resolved along with the cingulum that runs posterior-anteriorly. This is an apparent crossing, caused by the limited MRI resolution. In Fig. 5.11b, a crossing at the centrum semiovale is resolved. The corona radiata fibres that run primarily inferior-superiorly cross with longitudinal fasciculi. In the same region, corona radiata crosses with the corpus callosum, as indicated in Fig. 5.11c. Finally, the crossing between the corticospinal tract and the transverse pontine fibres that have a strong medio-lateral orientation is presented in 5.11d. In all cases, the regularized approach has returned orientations that change smoothly across the crossing region.

To test the global effect of regularization, streamline tractography was performed. The tracking algorithm was seeded from two different regions of interest (ROIs), one within the anterior body of the corpus callosum (Fig. 5.12) and the other lying in the internal capsule of the right hemisphere (Fig. 5.13). Orientation estimates from four different techniques were utilized, with tractography being seeded from the same ROIs in all cases. For the corpus callosum, tractography with DTI estimates resolved only medial tracts, similar to tractography using Q-ball imaging estimates, which was not able to resolve any crossings due to the low b -value and low angular resolution of the data. On the contrary, both two-tensor approaches reconstructed tracts through the centrum semiovale and exhibited greater fanning of the callosal fibres to the cortex. However, the regularized approach gave much smoother results and larger lateral coverage, compared to the non-regularized model. It did not propagate in the corona radiata and did not produce any curves within the inferior fronto-occipital fasciculus (see sagittal per-




Figure 5.11: Colour-coded orientations resolved from the regularized two-tensor model. In each case, a T2-weighted image is shown on the left, along with a colour-coded image of the principal DTI eigenvector. The colour-coding scheme is shown on the surface of a sphere in (d). The region delineated with the box is shown magnified on the right with the regularized two-tensor orientation estimates. (a) Crossing between the superior corpus callosum and the cingulum. (b) Crossing between the corona radiata and longitudinal fasciculi at the centrum semiovale. (c) Crossing between the corona radiata and the corpus callosum at the centrum semiovale. (d) Crossing between the corticospinal tract and the the transverse pontine fibres at the pons. Orientations are superimposed on the single tensor FA, in regions with one fibre orientation and the FA computed from the revised eigenvalues of the two-tensor model in crossing regions.

spective); these artefacts being evident for the non-regularized approach.

For the second ROI, DTI could not resolve the crossing between the corticospinal tract with the transverse pontine fibres. Thus, erroneous connections were identified to the contralateral hemisphere. Tractography with Q-ball and non-regularized two-tensor orientations had also similar problems, they exhibited however greater fanning to the cortex. The regularized two-tensor model resolved the crossing at the pons in a better way and also allowed greater fanning to the cortex, compared to DTI.

Apart from orientations, the two-tensor model provides revised eigenvalues in crossing regions. Fig. 5.14 shows revised fractional anisotropy (FA) and mean diffusivity (MD) images calculated from the new eigenvalue estimates. More specifically, both FA and MD were equal to their DTI values in non-planar regions, while the new prin-




Figure 5.12: Streamline tractography using orientation estimates of four different techniques (from left to right): DTI, Q-ball, non-regularized two-tensor model, regularized two-tensor approach. A seed ROI in the anterior body of the corpus callosum is used. Streamlines are superimposed on DTI FA maps and shown in both coronal (top) and sagittal (bottom) perspective.




Figure 5.13: Streamline tractography using orientation estimates of four different techniques (from left to right): DTI, Q-ball, non-regularized two-tensor model, regularized two-tensor approach. A seed ROI in the posterior limb of the internal capsule is used. Streamlines are superimposed on DTI FA maps and shown in both coronal (top) and sagittal (bottom) perspective.

cipal and secondary eigenvalues of the regularized two-tensor approach were used in planar voxels. Revised FA values were higher in regions identified as containing crossing fibres, compared to the FA value from the single tensor fit in the same region (white arrow), which is underestimated (Frank, 2001). Fig. 5.14b shows histograms of the FA values obtained using the single and two tensor models, in regions where fibre crossing was identified. The FA mean and standard deviation were 0.4426 and 0.103 for

the single tensor model, 0.6961 and 0.11 for the two tensor model. On the other hand, the MD images were identical, with an MD of roughly $0.7 \times 10^{-3} \text{ mm}^2/\text{s}$ in the brain parenchyma, as expected for a normal individual (Basser and Jones, 2002).




Figure 5.14: Revised FA and MD images using the regularized two-tensor estimates in planar regions. (a) Fractional anisotropy (FA) maps, (b) FA histograms in planar regions and (c) mean diffusivity (MD) images obtained from the single tensor DTI model (left) and the regularized two-tensor model (right). The revised FA and MD images (a), (c) on the right had the same value as the images on the left in non-planar regions, in planar regions though, they were calculated from the eigenvalues of either of the two resolved populations. Gaussian smoothing ($\sigma=0.8$) was applied to both FA images.

5.7 Discussion

A regularized two-tensor approach has been presented that resolves fibre crossings using conventional DTI data, acquired with low b value and relatively low number of DW directions. Compared to the 4-parameter geometrically constrained two-tensor model presented in (Peled et al., 2006), a) the problem of fitting a multi-parameter non-linear model was transformed to fitting many models of lower dimensionality using a set of basis directions, b) the model selection mask, based on the planar index c_p , was regularized and c) the orientation estimates in regions of fibre crossings were obtained using the lower dimensionality models and were regularized using relaxation labelling.

The geometrically constrained two-tensor model of (Peled et al., 2006) can be prone to noise (Fig. 5.6). Therefore, the orientation estimates were regularized to increase noise immunity, while keeping the advantage of low scanning requirements. Under the low curvature assumption, the two orientations resolved from the model were required to slowly change across neighbouring voxels, thus preserving the continuity of diffusion pathways. The planar index was also regularized to improve the model selection procedure (Fig. 5.5). This is important given that the model selection mask defines boundary conditions for the regularization of the orientations. Simulations showed that the regularized approach was more successful in resolving the crossing regions and the respective orientations, particularly in cases of low SNR (Fig. 5.6). Tractography results (Fig. 5.12 and 5.13) further supported these findings in-vivo and showed the benefits from regularization in a more global scale.

The regularized two-tensor approach provides information about orientations, diffusivity and anisotropy using a typical DTI scanning protocol (<6 minutes of scanning for a 128x128x52 volume). The recent super-resolved spherical deconvolution (Tournier et al., 2007) has shown satisfactory results in resolving fibre crossings using DTI data acquired in 4.5 minutes, however, only orientational estimates are given by this technique. PAS-MRI (Jansons and Alexander, 2003) has also shown good results using relatively low angular resolution DTI data, however it has very high computational requirements. For comparison, a Q-ball reconstruction was attempted on the data, using spherical harmonics as a basis (Hess et al., 2006). The resolving power of such a reconstruction was smaller compared to the proposed approach, as shown by the tractography results of Fig. 5.12, due to the low b value and the low angular sampling.

The idea of using a basis set to resolve fibre crossings has been also utilized in (Ramirez-Manzanares and Rivera, 2006). The DTI diffusion tensor was expressed as a linear combination of predefined basis tensors with the volume fractions of these tensors being spatially regularized. In planar regions, an oblate DTI tensor was decomposed in multiple compartments with orientations compatible with the neighbourhood. The eigenvalues and anisotropy of these basis tensors was the same across all voxels. The regularized two-tensor approach presented here allows different eigenvalues for different voxels. Furthermore, it is based on a two-tensor model, i.e. the

diffusion propagator is modelled as a sum of two zero-mean trivariate Gaussians; contrary to (Ramirez-Manzanares and Rivera, 2006) that keep the DTI assumption and decompose the single tensor to different compartments.

In (Ding et al., 2005; Parker et al., 2000) the raw image data were smoothed using non-linear filtering techniques. This can, in principle, remove noise artefacts and recover orientational coherence in white matter. However, there are artefacts caused by fitting non-linear models, which the above methods do not take into account. Fitting a multi-tensor model suffers from numerical instabilities, due to the multiple local minima of the objective function E (Tuch et al., 2002) and there are also identifiability issues (Zhou et al., 2008). Smoothing directly the estimates of these models, as done here, considers both types of artefacts.

A model selection procedure was necessary to identify regions where DTI fails. Residuals of the single tensor fit (Tuch et al., 2002), statistical tests that compare the residuals of single and multiple tensor models (Kreher et al., 2005) and the power of spherical harmonics coefficients fitted to the ADC profile (Alexander et al., 2002) have been utilized before. In this study, regions with a planar diffusion profile were considered inappropriate for DTI, as suggested by previous studies (Peled et al., 2006; Staempfli et al., 2006; Tuch et al., 2002; Wiegell et al., 2000) and a two-tensor-model was used instead. Planarity was quantified using the planar index c_p . To reduce noise artefacts, c_p was spatially regularized using two different approaches that gave comparable results. The regularized c_p appeared more resistant to noise compared to a mask calculated using spherical harmonics (Alexander et al., 2002) with the sampling scheme employed here (results not shown).

Assumptions and potential improvements

The approach presented in this chapter is governed by certain assumptions. A maximum number of two crossing tracts can be resolved. In cases of three crossing populations the DTI diffusion profile will be mostly spherical and will not be picked up by the model selection procedure. Such cases can be identified using Westin's spherical index c_s (Westin et al., 2002) within white matter and a three-tensor model can be fitted.

Utilizing a basis set of directions discretized the solution space and, therefore, re-

duced model accuracy. Reducing the step size $\delta\omega$ could improve accuracy. In a noisy environment, though, a finer set would provide better fits to noisy data, thus having a questionable advantage. Indeed, for cases of low SNR and/or number of DW directions, regularized orientation estimates were not significantly better when a step size of 10° or 2.5° was used (Fig. 5.8a and 5.9). Since decreasing the step size increased computation time geometrically, 10° was considered a good compromise between accuracy and speed.

An assumption of the geometrically constrained model was that the two crossing tensors are cylindrically symmetric and have the same eigenvalues. Scenarios of crossing tensors with different principal eigenvalues were simulated. This dissimilarity reduced c_p . For c_p values used mainly here ($c_p > 0.2$), errors in the model estimates were low, even in cases of 25% difference between the principal eigenvalues. For a 50% difference, estimates were poor for volume fraction and eigenvalues, but the maximum error for orientations was in the order of 10° . Therefore, differences between eigenvalues influenced mainly the volume fraction and eigenvalue estimates and much less the orientation estimates.

Assumptions were also made for regularization. The basis orientations that are compared between neighbouring voxels lie on different planes, as these are defined by the respective third DTI eigenvectors \mathbf{e}_3 . It is assumed that planes (i.e. \mathbf{e}_3) are slowly changing between neighbours, thus the compared orientations can be assumed roughly coplanar. Indeed, the median difference in \mathbf{e}_3 between neighbouring voxels, identified as planar by the regularized approach, was 9.82° across the whole brain volume. It should be pointed out that even if compared orientations are not perfectly coplanar, higher support will be given to orientations that promote smoother changes across neighbours relative to orientations that cause abrupt turns.

The orientations of the voxels identified as containing a single tract are not regularized; in these regions estimates have low uncertainties (Jones, 2003) and are therefore more resistant to noise than voxels with a planar diffusion profile (Fig. 5.6). However, the orientation estimates from single tract voxels act as boundary conditions for the regularization scheme. Thus, a potential improvement would be to regularize the single tensor eigenvectors, e.g. as in (Coulon et al., 2004; Tschumperle and Deriche, 2002), prior to application of the two-tensor model.

Summary

In summary, a method to regularize a two-tensor model has been presented, for better resolving fibre crossings using DTI data. Crossing orientations were estimated utilizing geometrically constrained two-tensor models of low dimensionality, applied to voxels with a planar diffusion profile. Using a regularization scheme, both the model selection procedure and orientation estimates were improved in images acquired with conventional DTI protocols. Tractography results showed the achieved improvement, as well as the smaller resolving power of Q-ball reconstruction when applied to this type of data. Apart from orientations, revised anisotropy and diffusivity information were also obtained. The ability to reconstruct complex intra-voxel structures robustly with conventional DW imaging techniques can be of importance in research, as well as in clinical environments.

Appendix A: Generating uniformly distributed directions on a unit sphere

Evenly spaced directions on the unit sphere were generated by maximizing the minimum angle between any pair of nearest neighbouring vectors. N directions, restricted on the hemisphere, were obtained as following:

- Start with N arbitrary unit vectors $\mathbf{v}_i = \{\theta_i, \phi_i\} = [\sin \theta_i \cos \phi_i \ \sin \theta_i \sin \phi_i \ \cos \theta_i]^T$, $i = 1 \dots N$, where $0 \leq \theta \leq \pi$ and $0 \leq \phi \leq 2\pi$ are the spherical angles.
- For each vector \mathbf{v}_i :
 - * Generate vectors $\mathbf{u}_1 = \{\theta_i, \phi_i\}$, $\mathbf{u}_2 = \{\theta_i + \delta, \phi_i\}$, $\mathbf{u}_3 = \{\theta_i - \delta, \phi_i\}$, $\mathbf{u}_4 = \{\theta_i, \phi_i + \delta\}$ and $\mathbf{u}_5 = \{\theta_i, \phi_i - \delta\}$, where δ is a small angle (1° used here).
 - * Compute the maximum dot products $md_k = \max_{j \neq i}(|\mathbf{u}_k \cdot \mathbf{v}_j|)$, with $k = 1 \dots 5$, representing the minimum angle between each of the vectors \mathbf{u} with the other $N - 1$ vectors \mathbf{v}_j .
 - * Set $\mathbf{v}_i = \mathbf{u}_k$, where k is determined by selecting the minimum md_k .
 - * Iterate until convergence.

Appendix B: Simulating the diffusion-weighted signal

The procedure described in (Kingsley, 2006b) was followed to simulated the diffusion-weighted signal. In general, the multi-tensor model (Tuch et al., 2002) was used to simulate the noise-free signal obtained from N Gaussian compartments using a gradient of b value and \mathbf{g}_k direction:

$$S_k = S_0 \sum_{n=1}^N f_n \exp(-b \mathbf{g}_k^T \mathcal{D}_n \mathbf{g}_k), \text{ with } \sum_{n=1}^N f_n = 1. \quad (5.14)$$

For $N=1$ this model reduces to the DTI model (Basser et al., 1994a). Noise was then added to the signal S_k in quadrature (Kingsley, 2006b). To simulate a given signal to noise ratio (SNR), the noise standard deviation was $\sigma = S_0/\text{SNR}$ and a noisy realization S'_k of the signal was:

$$S'_k = \sqrt{(S_k + \nu_R)^2 + \nu_I^2}, \quad (5.15)$$

with ν_R and ν_I being randomly drawn from a zero-mean Normal distribution with variance σ^2 , as shown in (Press et al., 1992).

To construct a diffusion tensor \mathcal{D} with eigenvalues $\lambda_1 \geq \lambda_2 \geq \lambda_3$ and principal orientation (θ, ϕ) , θ being the (zenith) angle from the z axis and ϕ being the (azimuth) angle on the xy plane (see Fig. 3.4), the following rotation matrices were used (Skare et al., 2000):

$$\mathcal{D} = \mathbf{R}^T \cdot \mathbf{L} \cdot \mathbf{R} \quad \text{with} \quad \mathbf{L} = \begin{bmatrix} \lambda_3 & 0 & 0 \\ 0 & \lambda_2 & 0 \\ 0 & 0 & \lambda_1 \end{bmatrix} \quad \text{and} \quad \mathbf{R} = \begin{bmatrix} \cos\theta\cos\phi & \cos\theta\sin\phi & -\sin\theta \\ -\sin\phi & \cos\phi & 0 \\ \sin\theta\cos\phi & \sin\theta\sin\phi & \cos\theta \end{bmatrix}.$$

For a given FA value and mean diffusivity MD(=trace/3), the eigenvalues of a cylindrically-symmetric tensor were obtained using (Jones, 2004):

$$\begin{aligned} \lambda_1 &= \text{MD} \left(1 + \frac{2\text{FA}}{\sqrt{3 - 2\text{FA}^2}} \right) \\ \lambda_2 &= \lambda_3 = \text{MD} \left(1 - \frac{\text{FA}}{\sqrt{3 - 2\text{FA}^2}} \right). \end{aligned} \quad (5.16)$$

Simulating tract geometries

In order to simulate a certain curve geometry (e.g. circular), tensors should be parallel to the tangent direction of the curve. Assuming that a curve \mathbf{r} is described parametrically by $\mathbf{r}(s) = [x(s) \ y(s) \ z(s)]^T$, the tangent to it at position $\mathbf{r}(s)$ will be $\mathbf{e}(\mathbf{r}(s)) = [x'(s) \ y'(s) \ z'(s)]^T$. Making a tensor's principal eigenvector \mathbf{e}_1 equal to \mathbf{e} at locations defined by the image grid, will allow the simulation of the DW signal that would be obtained from the certain geometry. Once the vector \mathbf{e}_1 is defined, a unit vector \mathbf{e}_2 can be chosen so that $\mathbf{e}_1 \cdot \mathbf{e}_2 = 0$. Then, the third eigenvector will be $\mathbf{e}_3 = \mathbf{e}_1 \times \mathbf{e}_2$. Having all the eigenvectors defined, the tensor will be $\mathcal{D} = \lambda_1 \mathbf{e}_1 \mathbf{e}_1^T + \lambda_2 \mathbf{e}_2 \mathbf{e}_2^T + \lambda_3 \mathbf{e}_3 \mathbf{e}_3^T$.

To simulate a circle of radius α centred at (x_o, y_o) and lying on the xy plane, we have:

$$\begin{pmatrix} x(\theta) = x_o + \alpha \cos \theta \\ y(\theta) = y_o + \alpha \sin \theta \end{pmatrix} \implies \begin{pmatrix} x'(\theta) = -\alpha \sin \theta \\ y'(\theta) = \alpha \cos \theta \end{pmatrix}, \quad (5.17)$$

with $\theta \in [0, 2\pi)$ being the angle from the x axis. If (x_o, y_o) is a fixed location on the

image grid, then a voxel with centre coordinates (x_1, y_1) will be on the circle with radius $\alpha = \sqrt{(x_1 - x_o)^2 + (y_1 - y_o)^2}$ and at angle $\theta = \text{atan}[(y_1 - y_o)/(x_1 - x_o)]$ ($\theta = \pi/2$ if $x_1 = x_o$). The tangent vector at (x_1, y_1) will then be $\mathbf{e} = [-\alpha \sin \theta \ \alpha \cos \theta]^T$.

6

Fuzzy Anatomical Connectedness of the Brain

Overview

Medical images are by nature fuzzy and diffusion-weighted MR images are no exception. In this chapter, a global, fuzzy algorithm for assessing white matter connectivity in the brain is presented. The proposed method considers anatomical paths as chains of linked neighbouring voxels. Links between neighbours are assigned weights using the respective fibre orientation estimates. By checking all possible paths between any two voxels, a connectedness value is assigned, representative of the weakest link of the strongest path connecting the voxel pair. Multiple orientations within a voxel can be incorporated, thus allowing the utilization of fibre crossing information, while fibre branching is inherently considered. Under the assumption that paths connected strongly to a seed will exhibit adequate orientational coherence, fuzzy connectedness values offer a relative measure of path feasibility.

Contributions of this chapter

- Fuzzy connectedness, an image segmentation algorithm, is modified and adapted to the WM tractography problem.
- Fuzzy connectedness tractography is presented as a global algorithm and is shown to work with diffusion MRI methods that provide any number of fibre orientation estimates per voxel. Thus, it can propagate through fibre crossings.
- Validation of the new algorithm is presented through simulations and in-vivo data. Diffusion tensor and Q-ball images are used as working examples. Results with orientations obtained from the two-tensor approach of the previous chapter are also presented.
- A qualitative comparison is performed with probabilistic tractography and the recent graph-based tractography.

Publications

Contributions from this chapter have appeared in the following:

- S. N. Sotiropoulos, L. Bai and C. R. Tench. Fuzzy anatomical connectedness of the brain using single and multiple fibre orientations obtained from diffusion MRI, *Computerized Medical Imaging and Graphics*, in press, 2009.
- S. N. Sotiropoulos, C. R. Tench and L. Bai. In-vivo brain anatomical connectivity using diffusion magnetic resonance imaging and fuzzy connectedness, *Proceedings of the IEEE International Conference on Bioinformatics and Bioengineering (BIBE)*, pp. 1-8, Athens, Greece, October 2008.
- S. N. Sotiropoulos, C. R. Tench, P. S. Morgan and L. Bai. Combining Q-ball imaging and fuzzy connectedness: An approach to distributed brain tractography, *Medical Image Understanding and Analysis (MIUA) Proceedings*, pp. 24-28, Dundee, Scotland, July 2008.
- S. N. Sotiropoulos, C. R. Tench and L. Bai, "Fuzzy anatomical connectedness using diffusion MRI: An approach to tractography of the brain", *Proceedings of the ISMRM Annual Meeting*, p.1846, Toronto, Canada, May 2008.

6.1 Introduction

Commonly used streamline tractography (Basser et al., 2000; Conturo et al., 1999; Mori et al., 1999) utilizes hard binary relationships to assess white matter connections, with two voxels being considered either connected or not connected. However, medical images are by nature fuzzy (Udupa and Saha, 2003) and DW images are no exception. Experimental noise, hardware limitations, limited spatial resolution, and partial volume artefacts are some of the factors that contribute to the fuzziness of the images. Therefore, a fuzzy framework seems more appropriate. In this chapter, a fuzzy algorithm for assessing anatomical connectivity is presented.

The probabilistic framework has been proposed as an alternative to the deterministic streamline tractography. Probabilistic techniques (Behrens et al., 2003; Friman et al., 2006; Hagmann et al., 2003; Jones and Pierpaoli, 2005; Parker et al., 2003) generate streamlines in a Monte-Carlo fashion and define an index of connectivity using visitation maps. A drawback of these techniques is the inevitable reduction of the connectivity index with distance from the seed (Behrens et al., 2003; Parker et al., 2003). Moreover, the repetitive streamline generation increases execution time, while the connectivity values can depend on the total number of streamlines launched.

All streamline-based methods, deterministic and probabilistic, determine white matter bundles using locally greedy criteria, i.e. tracking sequentially through orientation estimates in adjacent voxels. Another way of tract reconstruction is to identify the best path between two voxels of interest, according to some global criterion, rather than identifying paths arising from a single voxel. These global approaches are less sensitive to local artefacts introduced by various noise sources (Behrens and Jbabdi, 2009). Front propagation techniques (Campbell et al., 2005; Jackowski et al., 2005; Parker et al., 2002; Staempfli et al., 2006; Tournier et al., 2003) are an example. They identify a best path from a seed to all other brain voxels by evolving a surface front from this seed. A distributed index of connectivity can then be defined using either functions of the front arrival times in each voxel or the agreement between the surface trajectory and the underlying fibre orientation estimates. Distributed indices of connectivity are also calculated by graph-based tractography techniques (Iturria-Medina et al., 2007; Zalesky, 2008). These methods utilize graph theory to exhaustively search the image for the strongest path connecting any two voxels of interest.

In this chapter, a new global tractography method using a fuzzy framework for assessing brain anatomical connectivity is introduced. The new approach inherently considers fibre crossing information. It utilizes fuzzy connectedness (Udupa and Saha, 2003; Udupa and Samarasekera, 1996), an algorithm initially proposed for image segmentation, to tackle the tractography problem. Anatomical paths are considered as chains of linked neighbouring voxels. A connectedness value is assigned between any voxel and a seed of interest by checking all the possible paths that connect them on the discrete image grid and finding the strongest. The connectedness value reflects the weakest link along this strongest path. It quantifies the orientational coherence along a path and provides a relative measure for path feasibility, under the condition that fibre orientations change smoothly along an anatomical path. Therefore, it is expected that paths exhibiting high orientational coherence will acquire high connectedness values relative to the background.

Compared to other tractography methods, the algorithm provides converged connectivity values that do not drop systematically with the distance from the seed, in a single iteration and for all image voxels. Furthermore, path propagation can be performed with relatively high angular resolution and combined in a single step with connectivity assignment. Algorithm implementation with dynamic programming results in fast execution. The new method inherently accounts for fibre branching. To reduce the chance of false positives when propagating through crossing regions, multiple links are considered -when appropriate- between neighbouring voxels to account for the existence of multiple tracts. The proposed method can be combined with any DW-MRI reconstruction approach that provides $N \geq 1$ fibre orientation estimates. Results in simulated data, as well as in human data obtained from Diffusion Tensor Imaging (DTI) (Pierpaoli et al., 1996) and Q-ball Imaging (Tuch, 2004), are presented.

6.2 Fuzzy connectedness framework

Given a 3-dimensional digital image, a local fuzzy relation can be defined between any two neighbouring voxels \mathbf{i} and \mathbf{j} of the image. This relation is called affinity (Udupa and Samarasekera, 1996). The strength of the affinity is given by the function $\mu(\mathbf{i}, \mathbf{j}) (\mu : Z^3 \times Z^3 \rightarrow [0, 1])$, determined by the product of a) the similarity of

the image intensities or intensity-based features at i and j with b) the adjacency of the two voxels. Normally, the adjacency is a hard binary relationship with a value of 1 for neighbouring and 0 for non-neighbouring voxels. Therefore, non-neighbouring voxels have a zero affinity. A standard implementation uses a $3 \times 3 \times 3$ neighbourhood, which gives 26 non-zero affinities for every voxel i .

The fuzzy connectedness algorithm (FCA) (Udupa and Samarasekera, 1996) utilizes the local affinity values to assign strengths $FC(a, b)$ of a global fuzzy relation called fuzzy connectedness, between any pair of voxels a and b . Fuzzy connectedness takes values in $[0,1]$ and represents the weakest link of the strongest path connecting the two voxels of interest. There are many possible paths that connect a and b on the discrete image grid. Each path can be considered a chain of voxels, with successive elements being neighbouring. In this framework, the smallest affinity μ along a path, i.e. the weakest link of the chain, determines the path strength. Considering all possible paths connecting a and b , the strength of the strongest will be the connectedness $FC(a, b)$. An example is shown in Fig. 6.1, with three putative paths connecting the voxels of interest. Affinities are shown colour-coded and the strength of each path is determined by the smallest affinity along it.




Figure 6.1: Example of applying fuzzy connectedness on a simple spatial grid. To find the connectedness FC between voxels a and b , all the possible paths that connect them are explored, assume Paths 1, 2 and 3 in this example. Each path is a chain with links between neighbouring voxels. The strength of each link is given by the affinity of the linked voxels, shown in gray scale. The smallest affinity along a path determines the strength of the path. The strength of the strongest path will be the connectedness $FC(a, b)$.

Using FCA with a seed voxel or region of interest (ROI) and an affinity function, a connectedness value $FC(\mathbf{a}, \mathbf{s})$ between the seed \mathbf{s} and every other image voxel \mathbf{a} can be computed. In the following sections, this connectedness value is referred as $FC(\mathbf{a})$ for simplicity, since there is always a seed \mathbf{s} .

The fuzzy connectedness algorithm can be implemented efficiently using graph-searching methods, which result in execution at almost interactive speeds (Nyul et al., 2003). Instead of searching explicitly for each possible path between \mathbf{s} and all the other voxels, a recursive scheme is used, along with a queue as an auxiliary data structure. An FC value of zero is assigned initially to all voxels, apart from the seed that has a connectedness of one and is inserted into the queue. In each iteration, the voxel \mathbf{i} with the maximum FC value is dequeued and its neighbours \mathbf{j} are considered. For each \mathbf{j} , a trial value is computed as the minimum between the FC value of the focal voxel \mathbf{i} and the affinity $\mu(\mathbf{i}, \mathbf{j})$ between the two. This computes the weakest link (strength) of the path from \mathbf{s} to \mathbf{j} through \mathbf{i} . If the trial value is larger than the current FC value at \mathbf{j} , then this becomes the new fuzzy connectedness value of \mathbf{j} ; meaning that the path from \mathbf{s} through \mathbf{i} is the strongest so far to \mathbf{j} , such that its weakest link becomes the connectedness of \mathbf{j} . When such an update occurs, voxel \mathbf{j} enters the queue, unless it is already queued. Iterations stop when the queue becomes empty, i.e. optimal connectedness values have been calculated for all voxels.

6.3 Fuzzy connectedness tractography

To apply FCA to WM tractography, two main aspects of the algorithm should be considered and modified: the affinity function and the path generation procedure. The affinity is determined by the similarity of the fibre orientation estimates between neighbouring voxels, as derived from the intensities of the DW images. More details on the affinity function are presented below.

Regarding path generation, FCA takes into account all possible links that are legal voxel connections on the discrete image grid. However, this does not guarantee that the paths are anatomically realistic. Lateral and backward propagation (180° turns) needs to be avoided and propagation should be restricted along the WM fibre orientations. A modification is required to force forward path propagation. A *memory* property for each

voxel is introduced. Once a voxel \mathbf{a} is identified as part of a path, its previous voxel \mathbf{c} in the path chain, through which \mathbf{a} is connected to the seed in the strongest way, is stored. Knowing the vector $\mathbf{r}_{\mathbf{ca}}$ that connects voxel \mathbf{c} to \mathbf{a} , propagation can be constrained only to directions \mathbf{n} , such that $\mathbf{r}_{\mathbf{ca}} \cdot \mathbf{n} > 0$, prohibiting lateral or backward path propagation. Seed voxels, i.e. voxels without memory, are treated differently. All neighbours, both forward and backward, are considered for the seeds to initiate propagation in both directions. Different aspects of the fuzzy connectedness tractography (FCT) algorithm are presented in the following sections.

Affinity and connectedness based on DTI orientation estimates

Assuming that anatomical paths are smooth curves, fibre orientation estimates along a path are expected to vary smoothly. Therefore, the affinity between neighbouring voxels \mathbf{i} and \mathbf{j} is assessed using a coherence condition on the orientation estimates obtained from diffusion MRI. For now, a single fibre orientation estimate is assumed to be available for each voxel, as is the case for diffusion tensor imaging (DTI). Then, the affinity function is a modified version of the symmetric voxel linking function suggested in (Parker et al., 2002; Poupon et al., 2000):

$$\mu(\mathbf{i}, \mathbf{j}) = \mu(\mathbf{j}, \mathbf{i}) = \frac{1}{\mathcal{Z} \cdot (1 - \min(|\mathbf{e}(\mathbf{i}) \cdot \mathbf{n}(\mathbf{i}, \mathbf{j})|, |\mathbf{e}(\mathbf{j}) \cdot \mathbf{n}(\mathbf{i}, \mathbf{j})|, |\mathbf{e}(\mathbf{i}) \cdot \mathbf{e}(\mathbf{j})|))} , \quad (6.1)$$

where $\mathbf{e}(\mathbf{i})$ is the fibre orientation estimate at voxel \mathbf{i} (i.e. the principal DTI eigenvector), $\mathbf{n}(\mathbf{i}, \mathbf{j})$ is the unit vector connecting the centres of voxels \mathbf{i} and \mathbf{j} and \mathcal{Z} a normalization constant to keep affinities in the $[0,1]$ interval. \mathcal{Z} is chosen as the maximum of the non-normalized affinity values across the whole 3D image. The above function is peaked when all three vectors are collinear, i.e. when the orientation vector of voxel \mathbf{i} points to the centre of voxel \mathbf{j} and the orientation vector of voxel \mathbf{j} points to the centre of voxel \mathbf{i} . In this case, the affinity as presented in Eq. (6.1) becomes infinite, which is handled in practice by subtracting the $\min(\cdot)$ function from a number slightly larger than 1 (1.001 used here).

An example of the affinity values obtained with Eq. (6.1) is shown in Fig. 6.2. The affinity for two neighbours \mathbf{i} and \mathbf{j} is plotted as a function of the respective fibre orientations. For simplicity, orientations are varied only within plane. As expected, the




Figure 6.2: Affinity values for a pair of neighbouring voxels i and j . The normalized affinity μ is plotted as a function of the fibre orientations (shown with blue lines) at the two voxels. Only 2D orientations are considered for simplicity. The angles ϕ are relative to the horizontal axis. Colour code is light for high and dark for low.

maximum affinity is reached when both orientations become colinear with the connecting vector \mathbf{n} (i.e. $\phi_i = \phi_j = 90^\circ$).

The FCT algorithm utilizes the above affinity function to calculate fuzzy connectedness values $FC(\mathbf{a})$ between every image voxel \mathbf{a} and the seed voxel(s) \mathbf{s} . FCT also outputs the Memory array that keeps for every voxel \mathbf{a} its previous through which \mathbf{a} is connected to the seed in a strongest way. A priority queue is used as an auxiliary structure and the elements are sorted so that the one with the largest FC value is always the first queue element. Pseudo-code is provided below, for the case of a 3x3x3 neighbourhood.

Algorithm 6.1 Fuzzy Connectedness Tractography with DTI fibre orientations

```
1: for all image voxels a do
2:    $FC(a)=0$ ;  $Memory(a)=undefined$ ;
3: for all seed voxels s do
4:    $FC(s)=1$ ; Enqueue s;
5: while Queue  $\neq Empty$  do
6:   Dequeue element i with maximum  $FC$  value;
7:    $prev=Memory(i)$ ;
8:   for all neighbours j of i do
9:      $f_{min}=\min[FC(i),\mu(i,j)]$ ;
10:    if  $prev \neq undefined$  then
11:       $n_1 = i - prev$  and  $n_2 = j - i$ ;
12:       $dot = n_1 \cdot n_2$ ;
13:      if  $dot \leq 0$  then
14:         $f_{min} = 0$ ;
15:      if  $f_{min} > FC(j)$  then
16:         $FC(j)=f_{min}$ ;
17:         $Memory(j)=i$ ;
18:        if  $j \notin Queue$  then
19:          Enqueue j;
20:        else
21:          Update the Queue key of j;
```

Increasing angular resolution

Affinities are calculated between each image voxel **i** and its immediate neighbours. The neighbourhood size effectively determines the number of candidate propagation directions **n**. The implementation above uses the 3x3x3 neighbourhood, i.e. 26 neighbours in 3D. The angular resolution of path propagation can be enhanced by increasing the neighbourhood size to 5x5x5, including next nearest neighbours, as shown in Fig. 6.3. This gives 124 neighbours in 3D. In this case, though, neighbours that are not directly connected to the focal voxel **i** exist, such as voxel **j** in Fig. 6.3b. To avoid discontinuities in paths, at least one nearest neighbour of the focal voxel is required to be in the path. More specifically, for each next nearest neighbour **j** (light gray in Fig. 6.3b), the two nearest neighbours **c** and **d** (darker gray in Fig. 6.3b) that are traversed by the vector connecting the focal voxel **i** to **j** are considered. Then, the one with the highest connectedness FC_h value to the seed is chosen between **c** and **d**. If this connectedness is smaller than the connectedness value of **j** and propagation is decided from **i** to **j**, the assignment $FC_h=FC(j)$ occurs.




Figure 6.3: Different neighbourhood systems of focal voxel i , 3x3x3 in (a) and 5x5x5 in (b). Only neighbours within the same slice as i are presented. Black arrows represent candidate directions for path propagation. Dark gray: nearest neighbours of voxel i , light gray: next nearest neighbours of voxel i .

Incorporating fibre crossing information

What happens when a crossing occurs and more than one fibre orientation estimate are available for a voxel? To accommodate such scenarios, FCT can be modified accordingly. Assuming that up to $N \geq 1$ fibre orientations exist per voxel, the affinity function between voxels i and j can be generalized to an affinity between the l^{th} orientation of i and the k^{th} orientation of j :

$$\mu_{lk}(i, j) = \mu_{kl}(j, i) = \frac{1}{\mathcal{Z} \cdot (1 - \min(|\mathbf{e}_l(i) \cdot \mathbf{n}(i, j)|, |\mathbf{e}_k(j) \cdot \mathbf{n}(i, j)|, |\mathbf{e}_l(i) \cdot \mathbf{e}_k(j)|))} \cdot \quad (6.2)$$

Similarly as before, this affinity function is peaked when the l^{th} orientation vector of voxel i points to the centre of voxel j and the k^{th} orientation vector of voxel j points to the centre of voxel i . If L ($1 \leq L \leq N$) orientations exist in i and K ($1 \leq K \leq N$) orientations in j , the indices l and k range from 1 to L and 1 to K respectively, giving $L \times K$ affinity values between the two voxels.

To account for the multiple affinity values between two voxels, each of the orientations per voxel is treated as a separate element in the fuzzy connectedness framework. Then, the structural element of the algorithm becomes (i, l) , the l^{th} orientation of voxel i , rather than voxel i itself. Therefore, FCT does not compute a scalar connectedness value $FC(i)$ for a voxel i , but a vector $FC(i)$ with length L ; l^{th} vector entry $FC_l(i)$ corre-

sponds to the FC value of the l^{th} orientation. The final FC value for the specific voxel \mathbf{i} will be given by the maximum of these L values. The generalized FCT algorithm is given in the Appendix.

6.4 Data and processing

Simulations

A phantom containing a circular tract (Fig. 6.4a) and another containing two curved crossing tracts (Fig. 6.4b) were numerically generated. The tracts were embedded within a diffusively isotropic surround, to simulate gray matter (GM), where fibre orientations were random. The noise-free DW signals in each voxel were simulated as described in (Kingsley, 2006b), using a multi-tensor model:

$$S(\mathbf{g}_k) = S(0) \cdot \sum_{i=1}^N f_i \exp(-b_k \mathbf{g}_k^T \mathcal{D}_i \mathbf{g}_k), \quad \text{with} \quad \sum_{i=1}^N f_i = 1, \quad (6.3)$$

where N is the number of tracts simulated, \mathcal{D}_i is the diffusion tensor simulating the diffusion profile along the i^{th} tract, having a volume fraction f_i , and \mathbf{g}_k is the direction vector of the $k^{th} = 1 \dots M$ diffusion-sensitizing magnetic field gradient. Sixty-one ($M = 61$) evenly spaced gradient directions (Cook et al., 2007), $b=1000 \text{ s/mm}^2$ for the circular and $b=3000 \text{ s/mm}^2$ for the crossing phantom were used. The number of tracts was set to $N = 1$ for all voxels, apart from the voxels contained in the crossing region of the second phantom, where $N = 2$. The tensors in simulated WM tracts had a fractional anisotropy $FA=0.8$ (Pierpaoli and Basser, 1996), representative of regions within the highly anisotropic corpus callosum, and an orientation tangent to the desired simulated shape at each point. The tensors in GM voxels had an $FA=0.25$ and random orientation. All tensors had a trace of $2.1 \times 10^{-3} \text{ mm}^2/\text{s}$ (Pierpaoli et al., 1996).

Zero-mean Gaussian noise was added in quadrature, as described in (Kingsley, 2006b) to simulate the Rician nature of MRI noise (Gudbjartsson and Patz, 1995). The signal to noise ratio (SNR) was defined as the average intensity of the non-diffusion weighted image $S(0)$ divided by the standard deviation of the noise. Given the simulated DW data, fibre orientation estimates were obtained. For the first phantom the Dif-

fusion Tensor Imaging (DTI) model (Basser et al., 1994a; Pierpaoli et al., 1996) was used to get a single fibre orientation per voxel, as estimated from the principal eigenvector of the diffusion tensor. For the second phantom Q-ball imaging (QBI) (Tuch, 2004) was utilized. The peaks of the Q-ball ODF gave the orientation estimates. The ODFs were computed using even spherical harmonics up to 4th order as basis (Descoteaux et al., 2007; Hess et al., 2006), as explained in section 3.4. The CAMINO diffusion toolkit was then used (Cook et al., 2006) to find the ODF peaks. A maximum of $N=2$ orientations were calculated in each voxel. FCT was applied to both phantoms to test its performance in single ($N=1$) and multiple orientation ($N=2$) fields.

In-vivo images

A DTI (Pierpaoli et al., 1996) and a Q-ball (Tuch, 2004) whole-brain scan of a healthy male subject were performed; local research ethics committee approval and informed consent were obtained. A single-shot, spin-echo, echo-planar, diffusion-weighted sequence was used (acquisition matrix 112×112 with in-plane resolution $2 \times 2 \text{ mm}^2$, interpolated during reconstruction to 224×224) in a Philips 3T Achieva clinical imaging system (Philips Medical Systems, Best, The Netherlands). A parallel imaging factor of 2 was used. Six $b=0 \text{ s/mm}^2$ images were acquired and averaged. Diffusion weighting was applied in 61 evenly spaced directions (Cook et al., 2007) with $b=1000 \text{ s/mm}^2$ for the DTI scan ($TE=57.6 \text{ ms}$, $TR=10990 \text{ ms}$) and $b=3000 \text{ s/mm}^2$ for the Q-ball scan ($TE=72 \text{ ms}$, $TR=15292 \text{ ms}$). 52 slices were acquired with a thickness of 2 mm. Total imaging time was roughly 35 minutes.

Images were corrected for eddy current distortion using FSL's diffusion toolbox (Smith et al., 2004). Brain was extracted using BET (Smith et al., 2004). Corrected images were then tri-linearly interpolated along the out-of-slice axis to get isotropic voxels with $1 \times 1 \times 1 \text{ mm}^3$ dimensions. Fibre orientations were estimated in each case as described in the simulations section. For the DTI dataset, orientations were also estimated using the regularized two-tensor model presented in the previous chapter.

FCT was performed on both the DTI and Q-ball datasets. Voxels with a small degree of fractional anisotropy ($FA < 0.2$) were excluded from the analysis, assuming that these represent GM or CSF. For comparison purposes, streamline tractography using a modified FACT algorithm (Mori et al., 1999) was performed ($FA \text{ threshold}=0.2$, cur-

vature threshold= 45°), as it is implemented in the CAMINO toolkit (Cook et al., 2006). Q-ball orientation estimates were used, so that streamlines could propagate through crossings. The distributed graph-based tractography was also implemented, as described in (Iturria-Medina et al., 2007), using the whole QBI ODFs. Finally, multi-tensor ($N=2$) probabilistic tractography (Parker and Alexander, 2003) was performed using CAMINO (Cook et al., 2006).

6.5 Results

Simulations

Fig. 6.4a presents fuzzy connectedness maps obtained using FCT on the circular phantom. The gray scale corresponds to FC values between each voxel and the seed, which is indicated by the white arrow. For comparison, the binary connectivity index obtained from streamline tractography for the same seed is shown. The distributed nature of FCT compared to streamline tractography is evident. FCT identified strong paths arising from the seed, along with other weaker connections. WM voxels with coherent fibre orientations were distinguished from the non-coherent GM background. The effect of increasing angular resolution using more neighbours is also shown in Fig. 6.4a. As expected, a larger neighbourhood assisted the algorithm in capturing more curvature and increased the connectedness values.

The ability of the method to go through fibre crossings and utilize multiple fibre orientations per voxel is shown in Fig. 6.4b. When seeding within the individual tracts, paths in the appropriate tract were identified. When seeding in the crossing region, paths belonging to both tracts were found.

To test the behavior of FCT against noise, 100 Monte-Carlo simulations were performed using the circular phantom at a given SNR level. A linear ROI, indicated in Fig. 6.4a with a black solid line (left subfigure), defined a cross-section of the phantom. The central part of this cross-section (7 voxels) was within the WM circular ring, while its edges (5 voxels each) were within the GM region. The fuzzy connectedness values FC of the voxels along this ROI were studied across the 100 simulations. The seed voxel was the same across simulations, at the location indicated by the white ar-




Figure 6.4: Fuzzy connectedness (FC) maps obtained from simulated phantoms ($SNR=20$). Connectedness values are shown in gray scale. Seed voxels are indicated with white arrows. (a) From left to right: Fibre orientation estimates for the circular phantom, FC map obtained using a $3 \times 3 \times 3$ neighbourhood, FC map obtained using a $5 \times 5 \times 5$ neighbourhood and the same seed, Streamline generated from the same seed and corresponding binary connectivity index (white=connected, black=not connected). The black solid line in the orientations plot defines a cross-section used for quantitative simulations (see Fig. 6.5). (b) From left to right: fibre orientation estimates for the crossing phantom (cropped and magnified to show crossing region), FC maps obtained using a $5 \times 5 \times 5$ neighbourhood and three different seed points. In both (a) and (b), orientation estimates are superimposed on diffusion anisotropy values, which are high for WM voxels, low for GM voxels and intermediate for voxels in the crossing region.

row in Fig. 6.4a. Fig. 6.5 shows the mean and the standard deviation of the FC values at each of these discrete voxel locations, using the $3 \times 3 \times 3$ and the $5 \times 5 \times 5$ neighborhoods. The noise-free connectedness values along the ROI are shown by a dashed line, which is continuous, just for visualization purposes. The SNR in these simulations was 15, representative of the noise conditions in a DW-MRI acquisition. The plots show the robustness of the calculated values against noise and their good agreement with the noise-free ones. Given that the simulated tract is very curved and that the utilized affinity (Eq. 6.2) penalizes curvature, it is expected that the connectedness of the stronger paths will be smaller than one. Deviation from this ideal value is reduced in the case of the $5 \times 5 \times 5$ neighborhood (Fig. 6.5b). Compared to the values obtained with a $3 \times 3 \times 3$ neighborhood (Fig. 6.5a), the connectedness of the voxels belonging to the circular ring is almost doubled. Furthermore, the contrast between the more central and stronger paths and the surrounding weaker paths is enhanced with the $5 \times 5 \times 5$ neighborhood.

Nevertheless, using both neighborhood sizes, FCT differentiated regions that were coherently connected to the seed from the non-coherent ones.




Figure 6.5: Mean (circle) and standard deviation (horizontal bars) of fuzzy connectedness values across 100 simulations (SNR=15), at voxel locations defined by the linear ROI of Fig. 6.4a (black solid line). Voxel counting starts from the top left corner of the phantom and increases towards the phantom centre. *FC* values were obtained using (a) a 3x3x3 neighbourhood, (b) a 5x5x5 neighbourhood. Noise-free *FC* values are indicated by the dashed line.

In-vivo DW images

The performance of FCT in different brain regions is shown in Fig. 6.7 - Fig. 6.10. A 5x5x5 neighborhood was used in all cases. Typical execution times for a seed ROI were in the order of 1 minute for DTI FCT and 2 minutes for Q-ball FCT on a 3.2 GHz PC.

An axial ROI within the body of the corpus callosum seeded the algorithm in Fig. 6.7. The orientation estimates utilized are shown in Fig. 6.6. The corpus callosum crosses with the corona radiata and the superior longitudinal fasciculus at a region known as the centrum semiovale (Woolsey et al., 2008). As shown by the white arrows, Q-ball imaging resolves multiple crossing directions in this area.

In Fig. 6.7a, coronal maximum intensity projections (MIP) of the raw connectedness values *FC* are shown when the DTI and the Q-ball orientation estimates are utilized, respectively. (The intensity of a voxel on a MIP is the maximum intensity at this voxel location across all the slices parallel to the projection plane). For comparison, streamlines generated within the multi-fibre field of the Q-ball reconstruction are presented. The colour-coded connectedness maps had higher intensities in regions that

were relatively more likely to belong to the corpus callosum. The distributed nature of the algorithm is evident, since all image voxels are assigned a connectedness value.




Figure 6.6: DTI and Q-ball orientation estimates for roughly the same coronal slice, zoomed into the region of the corpus callosum. The crossings of the callosal with other tracts in the centrum semiovale (white arrows) are resolved only by Q-ball imaging. In both cases, orientations are superimposed on an FA map. Voxels with $FA < 0.2$ are masked out for better visualization.




Figure 6.7: Fuzzy connectedness (FC) maps obtained from an ROI within the corpus callosum. (a) From left to right: Coronal maximum intensity projection (MIP) of the raw FC values using the DTI orientation estimates, Coronal maximum intensity projection of the raw FC values using the Q-ball orientation estimates, Streamlines generated using the Q-ball orientation estimates, (b) MIP of thresholded FC values using the DTI estimates, along with coronal slices, (c) MIP of thresholded FC values using the Q-ball estimates, along with coronal slices. In all cases, the FC maps are superimposed on diffusion fractional anisotropy (FA) images. Seed locations are indicated by white arrows.

In Fig. 6.7b and Fig. 6.7c, thresholds were set to keep the top connected voxels and remove the relatively homogeneous FC values of the background. The thresholds were determined from the histograms of the FC values, where sharp discontinuities indicated the transition from the background to the top connected voxels (Fig. 6.8). These thresholds corresponded roughly to the 97.5th percentile of the distribution of the FC values for DTI FCT (Fig. 6.7b) and to the 95th percentile for Q-ball FCT (Fig. 6.7c), reflecting the higher coverage of the callosal tracts by the latter method. A MIP of thresholded connectedness values, along with the FC values on different coronal slices is presented in each case. DTI FCT (Fig. 6.7b) captured the medial portion of the corpus callosum, as expected when using the DTI orientations (Mori et al., 2005). However, more branching was captured compared to the medial portion resolved from streamline tractography. Q-ball FCT (Fig. 6.7c) went through the crossing at the cen-




Figure 6.8: Histogram-based detection of background connectedness threshold. The histogram of the *FC* values obtained with Q-ball FCT when seeding in the corpus callosum (Fig. 6.7) is shown. A very sharp transition in the histogram, pointed by the black arrow, was used to identify the background connectedness value.

trum semiovale and the results captured lateral projections to the cortex, apart from the medial callosal paths. It should be pointed out that the resolved crossings in Fig. 6.7c were not perfectly symmetric in the two hemispheres and greater fanning was evident in the left hemisphere. This was caused by asymmetric Q-ball reconstruction in the contralateral regions of the centrum semiovale, possibly due to noise.

A similar organization of tractography maps is presented in Fig. 6.9, with the seed being a coronal ROI in the cingulum. In this case, streamlines (Fig. 6.9a) were stopped at the level of the splenium of the corpus callosum and did not continue towards the parahippocampal gyrus, probably due to noise artefacts. The distributed nature of FCT, however, allowed the algorithm to continue, with Q-ball FCT (Fig. 6.9c) having a slightly higher connectedness through that point than DTI FCT (Fig. 6.9b). Given that there are not many fibre crossings along this tract, only small differences can be observed when single and multiple orientations per voxel are used (Fig. 6.9b and Fig. 6.9c). Q-ball FCT (Fig. 6.9c) assigned higher values in the subgenual area, where partial volume artefacts distort the DTI orientations. Furthermore, the body of the cingulum was slightly thicker (Fig. 6.9c) compared to the DTI FCT result (Fig. 6.9b), since the apparent (due to limited resolution) crossing between the cingulum and the corpus




Figure 6.9: Fuzzy connectedness maps obtained from a coronal ROI within the cingulum. (a) From left to right: Sagittal maximum intensity projection (MIP) of the raw *FC* values using the DTI orientation estimates, Sagittal maximum intensity projection of the raw *FC* values using the Q-ball orientation estimates, Streamlines generated using the Q-ball estimates, (b) MIP of thresholded *FC* values using the DTI estimates, along with axial slices, (c) MIP of thresholded *FC* values using the Q-ball estimates, along with axial slices. In all cases, the *FC* maps are superimposed on diffusion fractional anisotropy (FA) images. Seed locations are indicated by white arrows.

callosum cannot be resolved by DTI. For the same reason, paths with relatively low connectedness arising from the main body of the cingulum towards the cortex that were evident with DTI FCT were not observed with Q-ball FCT.

MIPs of thresholded FC values in other brain regions are shown in Fig. 6.10. In all cases DTI orientations were utilized, apart from Fig. 6.10a and Fig. 6.10c, where Q-ball estimates were used. In Fig. 6.10a, a sagittal ROI within the medial portion of the superior longitudinal fasciculus was the seed and only Q-ball orientations supported the existence of paths arising from the parietal cortex and ending up in motor cortical regions, in agreement with (Behrens et al., 2007). In Fig. 6.10b a single seed voxel was selected in the inferior part of the fornix, close to the amygdaloid complex. Fornix paths were reconstructed along with some paths in the anterior thalamic radiation, which had though smaller connectedness values. In Fig. 6.10c an axial seed ROI at the upper




Figure 6.10: Fuzzy connectedness maps obtained using various seeds in the brain. Maximum intensity projections of thresholded *FC* values are shown in each case. The seeds are placed, (a) at the medial portion of the superior longitudinal fasciculus, (b) at the inferior part of the fornix, (c) at the upper pons within the pyramidal tract, (d) at the splenium of the corpus callosum, and (e) at the apex of Meyer's loop. DTI orientations are used in all cases, apart from (a) and (c) where Q-ball estimates are utilized. Seed locations are indicated by white arrows.

pons was used and corona radiata paths were depicted. In Fig. 6.10d a single voxel in the splenium of the corpus callosum seeded the algorithm. Finally, a seed placed at the apex of Meyer's loop was found to be strongly connected to the optic radiation fanning to the occipital cortex and the lateral geniculate nucleus of the thalamus (Fig. 6.10e).

Apart from connectedness maps, paths generated by FCT can be reconstructed using the memory information. The strongest path connecting any voxel to a seed can be found in a backward fashion. Starting from an end point, one can recursively find the previous voxels and fibre orientations within those voxels that lead to the seed. By setting a threshold to the *FC* values, the paths that correspond to the top connected voxels can be plotted. An example is shown in Fig. 6.11a, where the most strongly connected paths generated using FCT from a seed ROI in the splenium of the corpus callosum are presented. Paths were generated by starting from the voxels with the top 0.5% *FC* values. For each voxel belonging to a path, the respective fibre orientation estimate was plotted. For comparison, streamlines generated using CAMINO (Cook et al., 2006) are shown in Fig. 6.11b. FCT paths exhibit greater branching.

Another example is shown in Fig. 6.12 that presents paths belonging to different brainstem tracts. During the backward propagation, a moving average filter was utilized to smooth the trajectories. The tracts of interest were the middle, superior and inferior cerebellar peduncle, the medial lemniscus and the corticospinal tract. FCT was




Figure 6.11: (a) Top 0.5% paths generated using FCT from a seed ROI in the splenium of the corpus callosum. DTI orientation estimates ($N=1$) were utilized. For each voxel belonging to a path, the respective fibre orientation is plotted. Paths were generated backwards starting from the top connected voxels back to the seed. (b) Streamlines generated using the same ROI. Paths are superimposed on FA maps.




Figure 6.12: Reconstruction of paths in the brainstem using FCT. Different colours are used for paths belonging to different WM tracts. Five different seed voxels were used at the middle cerebellar peduncle (red), at the right inferior cerebellar peduncle (orange), at the right superior cerebellar peduncle (green), at the right medial lemniscus (yellow) and at the right corticospinal tract (blue). The seed voxels are indicated with a sphere. Only the top paths connected to each of the seeds are shown. DTI orientation estimates were used. The right thalamus rendered and in white colour and an axial slice of the FA map are shown for anatomical reference.

applied individually for each of the five tracts, using one seed voxel per tract, and only the paths arising from the top 1% connected voxels were plotted (for the medial lemniscus and superior cerebellar peduncle a smaller threshold was used due to the smaller size of these tracts). These paths agree well with a-priori anatomical knowledge of the brainstem (Mori et al., 2005) and verify the anatomically realistic propagation of the algorithm.

FCT results were compared against the results of two other quantitative tractography algorithms, the latest graph-based algorithm (Iturria-Medina et al., 2007) and the multi-tensor ($N=2$) probabilistic tractography (Parker and Alexander, 2003). The graph-based method is a global tractography approach similar to FCT, so a comparison between them is more direct. Multi-tensor probabilistic tractography is a local tracking approach, but it was employed due to the popularity of probabilistic streamline techniques (Behrens and Jbabdi, 2009). All methods utilized the Q-ball data and an axial ROI within the body of the corpus callosum as a seed.

Maximum intensity projections of the *FC* values, of the path strengths obtained from graph-based tractography (GT) and of the index of connectivity from probabilistic tractography are presented in Fig. 6.13. We can observe that GT (Fig. 6.13b) resolves lateral portions of the corpus callosum (indicated by black arrows), but the strengths assigned to them are relatively low. Furthermore, paths belonging to the crossing corona




Figure 6.13: (a) Coronal maximum intensity projection of the raw *FC* values using the Q-ball FCT and an axial seed ROI in the body of the corpus callosum. (b) Coronal maximum intensity projection of the path strengths calculated using graph-based tractography and the Q-ball ODFs to define graph weights. (c) Coronal maximum intensity projection of the index of connectivity calculated using two-tensor probabilistic tractography. Colour coding is qualitatively the same in all cases with dark red being low and white being high.

radiata are assigned a high strength, meaning that propagation through the crossing at the centrum semiovale has given rise to false positive connections. For multi-tensor probabilistic tractography (Fig. 6.13c) lateral callosal tracts on the left hemisphere can be differentiated from the background. However, a higher index of connectivity is assigned to the crossing corona radiata tracts on the right hemisphere compared to the right callosal tracts. Furthermore, the inherent reduction of connectivity with distance from the seed due to the repetitive Monte-Carlo sampling is evident, especially in the right hemisphere. On the other hand, FCT propagated through the crossing bilaterally and assigned relatively high values only to corpus callosum paths. This is indicative of the algorithm's ability to propagate correctly through crossing configurations. It should be pointed out that apart from using the same data for all three algorithms, FCT and GT also utilized exactly the same ODFs. However, FCT has benefitted from the utilization of multiple affinities (i.e. links) between neighbours in crossing regions that reflect the existence of multiple fibre populations.

FCT can be combined with any technique that provides orientation estimates. So far, DTI and Q-ball orientations were utilized. The geometrically-constrained two-tensor model and the regularized two-tensor approach were also employed to estimate fibre orientations on the low b value DTI dataset. Fig. 6.14 shows fuzzy connectedness maps using these estimates. Coronal maximum intensity projections are shown of thresholded FC values with: Q-ball orientations, orientations from the geometrically-constrained two-tensor model of (Peled et al., 2006) and orientations from the regularized two-tensor model. Regularization for this HARDI dataset improved more the c_p model selection mask, rather than the orientations themselves. Lateral callosal paths were resolved bilaterally in all cases.




Figure 6.14: Coronal maximum intensity projection of thresholded FC values using (from left to right): the Q-ball orientations from the high b value dataset, the geometrically-constrained two-tensor model orientations from the low b value dataset and the regularized two-tensor model orientations from the low b value dataset ($\delta\omega = 2.5^\circ$, 30 iterations). An axial seed ROI in the body of the corpus callosum was used in all cases. Results are superimposed on fractional anisotropy maps.

6.6 Discussion

A new fuzzy framework to assess white matter anatomical connectivity in the brain was presented. Fuzzy connectedness tractography is a global tracking approach that utilizes any number of fibre orientations per voxel to produce distributed connectedness maps, so that every image voxel is eventually connected to the seed through some path. A connectedness value for a voxel characterizes how feasible the path connecting it to the seed is, given an affinity function. This value is computed by trying all possible valid connections between a voxel and the seed and finding the strongest. The smallest affinity, i.e. the weakest link, along this strongest path will be the connectedness of the voxel to the seed. The affinity employed here reflects the anatomical assumption of slowly varying fibre orientations along WM paths. Thus, the computed connectedness values provide a relative measure for path feasibility that depends on the orientational coherence along a path. The algorithm differentiates between voxels that exhibit orientational coherence to the seed and the background.

Comparison with other tractography algorithms

Apart from its fuzzy nature, the algorithm expands other distributed tractography methods (Campbell et al., 2005; Iturria-Medina et al., 2007; Jackowski et al., 2005; Parker et al., 2002; Staempfli et al., 2006) in two ways: a) FCT propagates paths using a relatively high angular resolution and b) it can utilize any number of fibre orientations per voxel and propagate through fibre crossings. It can thus be combined with any DW reconstruction technique that provides multiple orientation estimates, such as multi-tensor models (Tuch et al., 2002), spherical deconvolution (Tournier et al., 2004), persistent angular structure (Jansons and Alexander, 2003) and diffusion spectrum imaging (Wedeen et al., 2005).

In this chapter, Q-ball imaging and multi-tensor models were used as working examples to estimate multiple orientations per voxel. Q-ball imaging has been combined with other types of distributed tractography, such as front evolution (Campbell et al., 2005) and graph-based tractography (Iturria-Medina et al., 2007). However, both previous studies define a single weight between any two voxels, utilizing the whole Q-ball ODF, regardless of the number of fibre populations being present in each voxel. In the case of a fibre crossing, this may allow propagation towards all crossing orientations and artificially elevate the connectivity index of crossing tracts over paths belonging to the tract of interest, as shown in Fig. 6.13b. The approach followed here, defines -if needed- many affinities between two voxels and propagates more strongly along the appropriate orientation in the case of crossings. The analog to this method for front evolution techniques (Campbell et al., 2005; Jackowski et al., 2005; Parker et al., 2002; Staempfli et al., 2006; Tournier et al., 2003) would be the initialization and propagation of multiple fronts in crossing regions; this is not tackled by the algorithms found in the literature.

FCT produced reasonable connectedness maps in many different WM regions (Fig. 6.7 - Fig. 6.10). Given that validation of a tractography algorithm is in general difficult and impossible on a single subject basis, advanced streamline tractography using Q-ball orientation estimates and the same seed ROIs was performed; this at least shows that the FCT results are commensurate with the, now familiar from the literature, streamline results (Catani et al., 2002; Mori et al., 2005). Connectedness maps expanded the connectivity picture from streamline tractography (Fig. 6.7 - Fig. 6.9), ex-

hibiting greater branching. Even when only the DTI orientations were used, the results seem more distributed than the Q-ball streamlines, as shown in the cingulum (Fig. 6.9). It should be pointed out, that FCT operates on the discrete image grid and connects voxels centre-to-centre, in order to identify the strongest path between two voxels according to a global criterion. Streamline tractography can move anywhere within a voxel using locally greedy criteria. This can also contribute to the differences observed between the two techniques. Furthermore, when multiple orientations are available in a single voxel, the streamline tractography method employed here selects one of the available orientations. FCT will consider all available orientations and assign a connectedness value to each of them, depending on how coherently connected they are to the seed.

Compared to probabilistic streamline tractography approaches (Behrens et al., 2003; Friman et al., 2006; Jones and Pierpaoli, 2005; Parker et al., 2003), FCT assigns a fuzzy rather than a probabilistic index of connectivity. The latter gives the number of times a voxel has been traversed by a tract over repetitive streamline propagations in a Monte-Carlo fashion and is inherently sensitive to the distance from the seed (Fig. 6.13c). The former quantifies the fuzziness of paths and provides a relative measure of connectivity, given an affinity function. The affinity reflects the anatomical assumption of smoothness in WM paths and quantifies the coherence of adjacent fibre orientations. It can be influenced by factors that affect the estimation of orientations, such as limited spatial resolution, noise, diffusion profile reconstruction limitations and partial volume artefacts, and these factors can contribute to the observed spread of the FC values. Clearly, fuzzy and probabilistic indices are different and their interpretation in terms of absolute anatomical connectivity remains an open question for further investigation. Nevertheless, the fuzzy framework presented here extends the binary framework of streamline tractography, where only two outcomes are allowed ("not connected" vs "connected"), and relates to it more naturally rather than to probabilistic approaches.

Similar to all tractography algorithms, FCT can be influenced by unresolved partial volume artefacts. For example, false positives are evident in Fig. 6.10d, where paths of the inferior longitudinal fasciculus are shown connected to Meyer's loop. However, in most cases these artefacts were characterized by relatively lower connectedness values.

Fuzzy connectedness and Dijkstra's algorithm

An interesting relationship can be established between DTI FCT (Algorithm 6.1) that uses one affinity per voxel pair and graph-based tractography (Iturria-Medina et al., 2007). The former utilizes the fuzzy connectedness algorithm (Udupa and Samarasekera, 1996), the latter uses Dijkstra's algorithm (Dijkstra, 1959), a commonly used graph-searching method that solves the shortest path problem between graph nodes (Cormen et al., 2002).

As noted in (Carvalho et al., 1999; Nyul et al., 2003), fuzzy connectedness can be seen as a modified version of Dijkstra's algorithm. If we refer to Algorithm 6.1, once an element \mathbf{i} is dequeued (line 6 of Algorithm 6.1), $FC(\mathbf{i})$ has reached its optimal value and will no longer be updated (Nyul et al., 2003). This is because, once \mathbf{i} is dequeued, it has the maximum FC value amongst all queue elements and any subsequent voxel \mathbf{j} that is updated will have $FC(\mathbf{j}) \leq FC(\mathbf{i})$ (Note that $FC(\mathbf{j}) = f_{min} = \min[FC(\mathbf{i}), \mu(\mathbf{i}, \mathbf{j})] \leq FC(\mathbf{i})$). The same stands for Dijkstra's algorithm; by the time an element is dequeued, it will have reached its optimal strength.

Therefore, when one seed voxel \mathbf{s} is used, the only implementation difference between the two algorithms is the relaxation condition $f_{min} = \min[FC(\mathbf{i}), \mu(\mathbf{i}, \mathbf{j})]$ (line 9 of Algorithm 6.1). Replacing this condition with $f_{min} = FC(\mathbf{i}) \cdot \mu(\mathbf{i}, \mathbf{j})$ and interpreting FC and μ as path strength and graph weight respectively (rather than fuzzy connectedness and affinity), converts FCT to the graph-based tractography of (Iturria-Medina et al., 2007).

What is, then, the performance difference between the two methods? Imagine a scenario where the weakest affinity μ_x of a path is at a point \mathbf{x} , near the seed and this is generally smaller than most of the affinities in the image grid (e.g. due to a very curved transition at \mathbf{x}). FCT will then identify this weakest link and will assign to \mathbf{x} and to all voxels that follow with higher affinities, the same connectedness value equal to μ_x . That will make differentiation between paths after \mathbf{x} difficult. On the other hand, graph-based tractography, due to its multiplicative relaxation condition, can assign a different strength to each voxel after \mathbf{x} (unless a weight of 1 exists) and, thus, will not suffer from such an issue.

However, it has been shown (Fig. 6.13) and discussed that when multiple orienta-

tions are estimated within voxel, FCT treats them more appropriately due to the multiple affinities defined. In the next chapter, the benefits from combining the advantages of the two methods will be illustrated.

Assumptions and future work

FCT is governed by certain assumptions, which should be pointed out. The local ODF maxima have been used as orientation estimates from Q-ball images, an approach followed in many recent studies (e.g. (Haroon et al., 2009)). This approach has been shown to introduce bias in the estimated orientations, due to the finite width of each peak (Zhan and Yang, 2006). However, obtaining orientation estimates from the ODF peaks has been utilized here as an example. As pointed out before, FCT will work with any method that provides fibre orientations, such as the tensor decomposition approach described in (Schultz and Seidel, 2008) that does not bias the estimates.

Furthermore, orientational uncertainty, which is the basis for probabilistic tractography (Behrens et al., 2003; Friman et al., 2006; Jones and Pierpaoli, 2005; Parker et al., 2003) is not currently considered. Voxel similarity is assessed using strictly the fibre orientation estimates. However, the algorithm is independent of the affinity function, therefore more advanced affinities that incorporate uncertainty or whole ODF information can be used. An example towards this direction is shown in the next chapter.

The discreteness of the algorithm both in the spatial and the orientational domain may potentially cause problems, especially when very sharp turns are present within the range of a very few voxels. To overcome these situations, an interpolated image grid was used and the angular resolution of path propagation was increased (Fig. 6.3). As it has been shown, the results were satisfactory in many WM regions. A future extension of the current framework might consider more than one location within a voxel, so that not only centre-centre connections exist. A further remedy will be the incorporation of curvature in the affinity function, which currently penalizes deviations from straight line connections.

Identifying the background connectedness intensity and keeping only the relevant paths is a challenging issue in distributed tractography techniques. In this study, a histogram-driven threshold was used to choose the most strongly connected voxels to

the seed. More advanced methods that involve hypothesis testing (Morris et al., 2008) may be used.

The fact that the connectedness values are based on the weakest link along a path can make them sensitive to noise. Variations of the fuzzy connectedness algorithm exist with increased robustness against noise. Scale-based fuzzy connectedness (Saha et al., 2000) is an example and scale-based FCT can be a direct extension of the algorithm.

Summary

A fuzzy distributed algorithm to assess white matter connectivity from diffusion MRI orientation estimates was presented. Fuzzy connectedness tractography can use any number of fibre orientations per voxel and can therefore be combined with any model-based or model-free reconstruction technique that provides such estimates. In this chapter, its performance was illustrated using imaging data from two popular protocols, diffusion tensor and Q-ball imaging. The algorithm checks all valid paths on the discrete image grid that connect two voxels of interest and assigns a connectedness value between them, representative of the weakest link of the strongest path. Under the assumption that WM tracts exhibit orientational coherence, it can differentiate between regions that are highly connected to a seed and paths that are not. Its flexibility in the way neighbouring voxel similarity is assessed, allows the incorporation of different affinity functions and image modalities that may even be application-specific.

Appendix A: Fuzzy connectedness tractography algorithm

The FCT algorithm is based on the label-setting fuzzy connectedness algorithm presented in (Nyul et al., 2003). A priority queue, implemented as a d -ary ($d=4$) heap (Kruse and Ryba, 2000), is used as an auxiliary data structure. Each queue element (\mathbf{i}, l) comprises of the coordinates of voxel \mathbf{i} and an index l identifying which of the L orientations of voxel \mathbf{i} the element refers to. For each of the queue elements, a fuzzy connectedness value $FC_l(\mathbf{i})$ is stored and elements are sorted so that the one with the maximum connectedness value is always the first element of the queue. The presented implementation utilizes $N \geq 1$ fibre orientations per voxel and a $5 \times 5 \times 5$ neighbourhood. Neighbouring elements are sorted, so that nearest neighbours are considered first and next nearest neighbours follow during each propagation step. The algorithm utilizes a function **Update()** defined below.

Function 6.2 Update($\mathbf{j}, k, \mathbf{i}, l, \text{val}$)

- 1: $FC_k(\mathbf{j}) = \text{val};$
 - 2: $\text{Memory}_k(\mathbf{j}) = (\mathbf{i}, l);$
 - 3: **if** $(\mathbf{j}, k) \notin \text{Queue}$ **then**
 - 4: Enqueue $(\mathbf{j}, k);$
 - 5: **else**
 - 6: Change the Queue key of $(\mathbf{j}, k);$
-

Algorithm 6.3 Fuzzy Connectedness Tractography

```
1: for all image voxels a do
2:   for each orientation  $n$  of a do
3:      $FC_n(\mathbf{a})=0$ ;  $\text{Memory}_n(\mathbf{a}) = \text{undefined}$ ;
4:   for all seed voxels s do
5:     for each orientation  $n$  of s do
6:        $FC_n(\mathbf{s})=1$ ; Enqueue (s, $n$ );
7:   while Queue  $\neq$  Empty do
8:     Dequeue element (i, $l$ ) with maximum  $FC$  value;
9:     Use  $\text{Memory}_l(\mathbf{i})$  to determine valid neighbours;
10:    for all valid neighbours j of i do
11:      for each orientation  $k$  of j do
12:         $f_{\min}=\min[FC_l(\mathbf{i}),\mu_{lk}(\mathbf{i},\mathbf{j})]$ ;
13:        if  $f_{\min}>FC_k(\mathbf{j})$  then
14:          if j is a nearest neighbour then
15:            Update(j,  $k$ , i,  $l$ ,  $f_{\min}$ );
16:          if j is a next nearest neighbour then
17:            c,d = nearest neighbours connecting j to i;
18:            Choose amongst orientations in c, the  $p$  that is closest to  $k$ ;
19:            Choose amongst orientations in d, the  $q$  that is closest to  $k$ ;
20:            (z, $t$ ) = Pick from (c, $p$ ), (d, $q$ ) the one with higher  $FC$ ;
21:            Update(j,  $k$ , z,  $t$ ,  $f_{\min}$ );
22:            if  $f_{\min}>FC_d(\mathbf{z})$  then
23:              Update(z,  $t$ , i,  $l$ ,  $f_{\min}$ );
24:  for all image voxels a do
25:     $FC(\mathbf{a})=\max_n[FC_n(\mathbf{a})]$ ;
```

7

Graphs and multigraphs for white matter tractography

Overview

In the previous chapter, connections between fuzzy connectedness and graph-based tractography were discussed. The benefit from the consideration of multiple voxel affinities in crossing regions was illustrated. In this chapter, a similar idea is applied to graph-based tractography algorithms, to achieve robust propagation through crossing fibre configurations, as depicted by Q-ball imaging ODFs. First, graph theory for tract reconstruction is presented. Then, building up on this theory, a new algorithm that utilizes a *multigraph* representation of the brain volume is derived. Each image voxel is treated as a graph node and graph arcs connect neighbouring voxels. Weights representative of both structural and diffusivity features are assigned to each arc. To account for the existence of crossing fibre populations within a voxel, multiple weights between neighbouring voxels are defined -where appropriate-, each representative of a different combination of single-fibre compartments. The new structure, termed a

multigraph, is searched exhaustively, but efficiently, to find the strongest paths and assign connectivity strengths between a seed and all the other image voxels. Compared to existing graph-based tracking methods, the new approach improves the connection strengths through fibre crossing regions, reducing the strengths of paths that are less anatomically plausible.

Contributions of this chapter

- White matter reconstruction using elements of graph theory is presented.
- The multigraph-based representation of a brain volume is introduced. Multigraph-based tractography is described as a generalization of existing methods that propagates more robustly through fibre crossings.
- A quantitative comparison of the new method is performed with graph-based and fast marching tractography.
- In-vivo tracking results are obtained on Q-ball imaging data. Using the path strength as an index of connectivity, thalamic parcellation is performed.

Publications

Contributions from this chapter have appeared in the following:

- S. N. Sotiropoulos, L. Bai, P. S. Morgan, C. S. Constantinescu and C. R. Tench. Brain tractography using Q-ball imaging and graph theory: Improved connectivities through fibre crossings via a model-based approach, *NeuroImage*, 49:2444-2456, 2010.
- S. N. Sotiropoulos, C. R. Tench, P. S. Morgan and L. Bai. Robust graph-based tracking through crossing fibre configurations, *Proceedings of the IEEE International Symposium on Biomedical Imaging (ISBI)*, pp. 1394-97, Boston, USA, June 2009.
- S. N. Sotiropoulos, L. Bai, P. S. Morgan and C. R. Tench. Graph-based tractography for robust propagation through complex fibre configurations, *Proceedings of the ISMRM Annual Meeting*, p. 852, Honolulu, USA, April 2009.

7.1 Introduction

Global tractography techniques offer an alternative to local, streamline-based approaches (Behrens and Jbabdi, 2009). Front propagation methods (Campbell et al., 2005; Parker et al., 2002) were the first that defined optimal paths according to a global criterion. In more recent studies, a weighted graph representation of the image has been utilized (Iturria-Medina et al., 2007; Lifshits et al., 2009; Zalesky, 2008) to search for optimal paths. These are identified using modified versions of shortest-path algorithms (Cormen et al., 2002). The graph-based methods, as well as most of the front propagation methods, are inherently discrete in the orientation and spatial fields. However, compared to the streamline-based approaches and their probabilistic counterparts, they combine a) converged indices of connectivity to a seed for all image voxels (Iturria-Medina et al., 2007; Parker et al., 2002), b) connectivities that do not drop systematically with the distance from the seed (Iturria-Medina et al., 2007; Parker et al., 2002), c) inherent ability to incorporate information from other imaging modalities (Iturria-Medina et al., 2007), d) inherent ability to deal with fibre branching in a single-pass execution (Iturria-Medina et al., 2007; Parker et al., 2002) and e) relatively short execution times (Parker et al., 2002).

Despite the potential of such methods, only a few studies that utilize Q-ball imaging and non-streamline tractography exist (Campbell et al., 2005; Iturria-Medina et al., 2007). For these studies, propagation in each step is determined using the whole Q-ball ODF of the current voxel, and in crossing regions connections are distributed equally towards all crossing directions, regardless of the path propagated so far. Therefore, anatomically unlikely connections may be identified. In this chapter, this limitation is discussed and a new graph-based tractography algorithm that deals more appropriately with crossing regions is presented. The algorithm takes into account multiple fibre populations within a voxel, when partial volume exists, by treating the image as a multigraph. It is then possible to distribute the connectivities in a weighted manner, with the most appropriate fibre population (or populations) obtaining the highest weight. The aim is thus to reduce the connection strengths of paths from a seed that are less anatomically plausible, without using exclusion/inclusion ROIs that require some prior knowledge. Results are presented on simulated data as well as on a group of healthy human subjects and quantitative comparisons are performed with both Q-ball

based front evolution tractography (Campbell et al., 2005) and the recent graph-based tractography (Iturria-Medina et al., 2007).

7.2 Brain tractography using graph theory

An image can be considered as a non-directed weighted graph $G = [V, A]$, with V being the set of graph nodes and A the set of graph arcs. Each voxel \mathbf{i} in the image is a graph node and a neighbourhood system is chosen to define the set of neighbours \mathbf{F}_i^{neigh} of \mathbf{i} . Arcs, a , exist between neighbouring voxels and weights $w \in \mathfrak{R}$, are assigned to each arc reflecting both diffusivity and structural information of the connecting voxels. According to (Iturria-Medina et al., 2007) the weight of the arc between two neighbours \mathbf{i} and \mathbf{j} is symmetric and defined as:

$$w[a(\mathbf{i}, \mathbf{j})] \equiv w_{ij} = w_{ji} = P_{mat}(\mathbf{i}) \cdot P_{mat}(\mathbf{j}) [P_{Diff}(\mathbf{i}, \mathbf{r}_{ij}) + P_{Diff}(\mathbf{j}, \mathbf{r}_{ji})] . \quad (7.1)$$

The term $P_{mat}(\mathbf{i})$ represents the probability of voxel \mathbf{i} belonging to a specific tissue type. It can be computed by performing probabilistic tissue segmentation on structural images. $P_{Diff}(\mathbf{i}, \mathbf{r}_{ij})$ is a pseudo-probability term of observing diffusion from \mathbf{i} along the direction \mathbf{r}_{ij} that connects the centres of \mathbf{i} and \mathbf{j} . It can be computed by integrating the diffusion ODF over a solid angle ω around the vector \mathbf{r}_{ij} . The angle ω is determined by the neighbourhood system employed, i.e. the number of arcs arising from a node (for a 3x3x3 neighbourhood, $\omega = 4\pi/26$). Then:

$$P_{Diff}(\mathbf{i}, \mathbf{r}_{ij}) = \frac{1}{\mathcal{Z}} \int_{\omega} ODF(\mathbf{i}, \mathbf{r}) dS \approx \frac{1}{\mathcal{Z}} \sum_{q=1}^{N_Q} ODF(\mathbf{i}, \mathbf{r}_q) \Delta S_q , \quad (7.2)$$

where the integral is approximated by a sum of ODF terms evaluated at N_Q points \mathbf{r}_q contained in the solid angle cone and obtained from an icosahedral tessellation. The normalization constant \mathcal{Z} ensures that the maximum value of the set $\{P_{Diff}(\mathbf{i}, \mathbf{r}_{ij}), \mathbf{j} \in \mathbf{F}_i^{neigh}\}$ is 0.5. Using the above definitions the arc weights are $w \in [0, 1]$. Note that in (Zalesky, 2008), the weights are defined by integrating the Bayesian posterior probability of fibre orientation rather than the ODF.

Given the weights, paths can be defined in the image graph as chains of neigh-

bouring voxels. The strength M of a path C that comprises of X nodes $\mathbf{i}_1, \mathbf{i}_2, \dots, \mathbf{i}_X$ or equivalently of $X-1$ arcs $a(\mathbf{i}_1, \mathbf{i}_2), a(\mathbf{i}_2, \mathbf{i}_3), \dots, a(\mathbf{i}_{X-1}, \mathbf{i}_X)$ is given by:

$$M_C \equiv M_{i_1 \dots i_X} = M_{i_X \dots i_1} = w_{i_1 i_2} \cdot w_{i_2 i_3} \cdot \dots w_{i_{X-1} i_X} . \quad (7.3)$$

Many paths exist between any two voxels. The strongest path between them can be found by searching the graph exhaustively to maximize M . In (Iturria-Medina et al., 2007; Zalesky, 2008) a modified version of Dijkstra's algorithm (Dijkstra, 1959) is used to perform this search. To force anatomically realistic path propagation, a 90° maximum curvature threshold is imposed between two successive path steps. The strongest path between a seed and all image voxels can be found and the strengths of these paths form a new image. In (Zalesky, 2008) this strength map is treated as a connectivity image, while in (Iturria-Medina et al., 2007) strengths are utilized to keep the most relevant paths propagating from a seed and exclude the background. The smallest arc weight along each of these non-background paths is then used as a connectivity index.

Incorporating Q-ball ODFs - Problem outline

Regardless of the chosen connectivity index, strengths reflect the way paths are propagated and are used to infer anatomically relevant paths from a seed (Iturria-Medina et al., 2007; Zalesky, 2008). The question addressed here is whether the calculated strength maps allow anatomical inference, when Q-ball derived ODFs are used with the previous framework to propagate through crossing regions.

One would anticipate that replacing DTI ODFs, as used in previous studies (Iturria-Medina et al., 2007), with multi-peaked Q-ball ODFs should allow more appropriate propagation through crossing regions. A simulated example of a perpendicular crossing shown in (Iturria-Medina et al., 2007) verifies the expectation for this particular situation.

However, the algorithm is not always successful in describing the expected picture. Fig. 7.1 shows the strength maps obtained using Q-ball ODFs in the previous framework, for two simulated straight crossings of 70° (Fig. 7.1b) and 60° (Fig. 7.1c), in a noise-free environment. For each voxel \mathbf{i} the strength M of the strongest path connecting \mathbf{i} to the seed is plotted. In both cases, a seed ROI is placed within the horizontal

tract. For the 70° crossing, propagation strength is correctly high for the horizontal tract, but it is also undesirably elevated for the diagonal tract. Even worse, for the 60° crossing, propagation to the diagonal tract is favoured and connectivity is restricted within the horizontal tract. This problem is absent for a perpendicular crossing due to the 90° curvature threshold. It should be pointed out that in the case of DTI ODFs, which lack directionality in crossing regions, propagation to all crossing tracts might be acceptable.




Figure 7.1: (a) Diffusion ODFs calculated for a phantom with a 70° crossing. ODFs are colour-coded with the colour scheme plotted on the surface of a sphere. Strength maps for (b) 70° and (c) 60° crossing phantom obtained using conventional graph-based tractography and the seed ROIs indicated with the black arrow. Strengths are colour-coded as shown in the legend.

The problems highlighted above arise in crossing regions from the consideration of the whole ODF in each propagation step. Q-ball ODFs may comprise of many fibre compartments. Once a path reaches a voxel where fibre crossings occur, it can propagate towards any or all crossing tracts equally regardless of its history; at least within the 90° turning threshold. In the example of Fig. 7.2a, all weights w_{ji} , w_{if} and w_{id} will be high due to the relative ODF configuration. As a result, the path **j-i-d** will have an undesirably high strength. The problem can sometimes be masked by the discrete nature of path propagation, however it is in general present even for high crossing angles (70°), as shown in Fig. 7.1. Note that the ODF-driven fast marching tractography (Campbell et al., 2005) will be influenced by the same factor, since the whole ODF is considered for the speed calculation. This will be illustrated in the results section.




Figure 7.2: Graph structure in the case of a crossing in (a) conventional graph-based tractography (Iturria-Medina et al., 2007), (b) multigraph-based tractography. For each node the diffusion ODF is plotted. All Q-ball ODFs are sharpened (raised to the third power). The $N=2$ ODF components plotted for voxel i in (b) correspond to the two Gaussian ODFs estimated using a two-tensor model in this voxel.

7.3 Multigraph-based tractography

To alleviate the problem, the weight calculation, the graph structure, and the graph searching routine are modified. It is assumed that ODFs with N peaks correspond to N crossing populations and can be decomposed to N single-peaked $ODF_n (n = 1..N)$ components, each representative of a different population; this is illustrated in the following section. Weights $w_{i_n j_l}$ are then defined between the $n^{th} (n = 1..N)$ fibre population of voxel i and the $l^{th} (l = 1..L)$ population of voxel j , N and L being the number of ODF peaks of the two voxels respectively. This gives $N \times L$ weights between i and j . The diffusivity terms used to compute graph weights are obtained by integrating the n^{th} and l^{th} single-peaked ODF components of voxels i and j , for example:

$$P_{Diff-n}(\mathbf{i}, \mathbf{r}_{ij}) = \frac{1}{\mathcal{B}} \int_{\omega} ODF_n(\mathbf{i}, \mathbf{r}) dS, \quad (7.4)$$

where the normalization constant \mathcal{B} ensures that the set $\{P_{Diff-n}(\mathbf{i}, \mathbf{r}_{ij}), \mathbf{j} \in \mathbf{F}_i^{neigh}, n \in [1, N]\}$ has a maximum value of 0.5. The new graph weights are:

$$w_{i_n j_l} = w_{j_l i_n} = P_{mat}(\mathbf{i}) \cdot P_{mat}(\mathbf{j}) [P_{Diff-n}(\mathbf{i}, \mathbf{r}_{ij}) + P_{Diff-l}(\mathbf{j}, \mathbf{r}_{ji})]. \quad (7.5)$$

To account for the multiple weights between neighbours, a voxel i is represented in the image graph by N instances, one per population (Fig. 7.2b). Each voxel instance

\mathbf{i}_n is treated as a separate element in the graph that keeps the same graph connections as \mathbf{i} . Due to the existence of multiple instances, rather than calculating a single strength value between each image voxel \mathbf{i} and the seed \mathbf{s} , a strength vector $\mathbf{M}_{i\dots s}$ with length N is calculated per voxel; n^{th} vector entry $M_{i\dots s}^n$ corresponds to the strength of the strongest path connecting the n^{th} fibre population \mathbf{i}_n to \mathbf{s} .

The path propagation follows the idea of Dijkstra's algorithm (Dijkstra, 1959). The instances of all image voxels are initially inserted into a queue and are assigned a zero strength. The seed is also inserted with a strength of one. In each iteration of the algorithm, the instance \mathbf{i}_n with the highest strength is removed from the queue. For \mathbf{i}_n , the strength of the paths from the seed \mathbf{s} , through \mathbf{i}_n , to all instances \mathbf{j}_l of neighbours $\mathbf{j} \in \mathbf{F}_i^{neigh}$ are evaluated (Eq. (7.3)). If a new path strength is found to be larger than the previous (i.e. of a path that did not go through \mathbf{i}_n), the best path connecting \mathbf{j}_l to the seed \mathbf{s} is updated and the instance \mathbf{i}_n is included in this path. Only a subset \mathbf{V}_i^{neigh} of the neighbours of \mathbf{i} that require a curvature change of less than 90° are considered. Effectively, all neighbouring instances $l = 1 \dots L_j, \mathbf{j} \in \mathbf{V}_i^{neigh}$ (L_j being the number of instances in \mathbf{j}) are tried; this gives $\sum_j L_j, \mathbf{j} \in \mathbf{V}_i^{neigh}$ combinations for instance \mathbf{i}_n . Iterations continue until all instances are removed from the queue and strength vectors $\mathbf{M}_{i\dots s}$ are calculated for all voxels \mathbf{i} in the image. Then, a post processing step finds for each voxel \mathbf{i} the maximum of the $M_{i\dots s}^n$ values to obtain a scalar map of connection strengths.

The calculated path strengths depend on the alignment of ODFs along the path, with higher strengths meaning better alignment. The consideration of multiple instances where appropriate, effectively allows a low curvature prior to be imposed through fibre crossings.

The graph searching algorithm presented in (Iturria-Medina et al., 2007) is modified to account for the changes described above. The presence of multiple weights between two nodes is equivalent to the presence of multiple arcs connecting the two nodes. This generalized graph structure is known as a multigraph (Zwillinger, 2003) and the new method is called multigraph-based tractography (MGT). MGT tests in crossing regions all the possible combinations between fibre populations and, depending on the seed location, allows the appropriate ones dominate. In the example of Fig. 7.2b, the path $\mathbf{j}_1 - \mathbf{i}_1 - \mathbf{f}_1$ will now have a much larger strength than the path $\mathbf{j}_1 - \mathbf{i}_2 - \mathbf{d}_1$, which is less anatomically plausible. A detailed description of the algorithm is given

in the Appendix.

7.4 Estimation of multiple single-peaked ODFs within a voxel

As shown before, given a set of Q-ball images (Tuch, 2004) the diffusion orientation distribution function (ODF) at voxel \mathbf{i} can be estimated (Descoteaux et al., 2007). ODFs with N peaks are assumed to correspond to N crossing fibre populations. Under this assumption, an appropriate single-fibre component ODF_n can be assigned to the n^{th} population. A model-based approach to estimate these different components is proposed here.

The peaks of the ODF are detected using a finite difference scheme, similar to (Descoteaux et al., 2009). Uniformly distributed points on the sphere are obtained using an icosahedral tessellation. The ODF is sharpened (raised to the 3^{rd} power) and the ODF values for each of the tessellation points are evaluated. Candidate peaks are identified at points, where the ODF is larger than the neighbouring ODF values and also exceeds a threshold (50% of the maximum ODF value used here). Since noise artefacts can increase the number of detectable local maxima, candidate peaks are subsequently filtered. The filtering represents the belief that an ODF at the b values used here, cannot resolve crossings of small angle (Cho et al., 2008). In an iterative fashion, peaks that are separated by an angle less than a minimum threshold α from the others are discarded. In each iteration the angle between all pairs of candidate peaks \mathbf{p} is calculated. For each \mathbf{p} , a counter keeps the number of other peaks that are within an angle α . The peak with the maximum counter value is removed at the end of each iteration. If more than one peaks have the same counter value, then \mathbf{p} with the smallest ODF value is removed. Iterations are continued until the counter value of all remaining peaks is zero. An $\alpha = 45^\circ$ was used.

In voxels with $1 < N \leq 3$ peaks, a mixture model of prolate diffusion tensors (Alexander and Barker, 2005) is fitted to the signal using a Levenberg-Marquardt approach (Press et al., 1992). The number of fitted tensors N is equal to the number of

ODF peaks detected. The model at a given voxel is:

$$S_k = S_0 \sum_{n=1}^N f_n \exp(-b \mathbf{g}_k^T \mathcal{D}_n \mathbf{g}_k), \quad \text{with} \quad \sum_{n=1}^N f_n = 1 \quad \text{and} \quad \mathcal{D}_n = (\lambda_{n1} - \lambda_{n2}) \mathbf{e}_{n1} \mathbf{e}_{n1}^T + \lambda_{n2} \mathbf{I}. \quad (7.6)$$

In the above formulas, S_k is the DW signal obtained with a diffusion-sensitizing magnetic field gradient of b factor and direction \mathbf{g}_k , f_n is the volume fraction of the prolate tensor \mathcal{D}_n with principal eigenvalue, secondary eigenvalue and principal eigenvector λ_{n1} , λ_{n2} and \mathbf{e}_{n1} , respectively. The n^{th} tensor is aligned with the n^{th} ODF peak and only its eigenvalues are estimated (see Appendix). Then, the n^{th} Gaussian ODF component in voxel \mathbf{i} will be (see Eq. (3.11)):

$$ODF_n(\mathbf{i}, \mathbf{u}) = \frac{\mathcal{L}}{\sqrt{\mathbf{u}^T \cdot \mathcal{D}_n^{-1} \cdot \mathbf{u}}}, \quad (7.7)$$

with \mathcal{L} being a normalization constant. In (Iturria-Medina et al., 2007), the ODFs are raised to the third power to achieve sharpening. The same approach has been followed here.

Voxels with $N > 3$ were considered isotropic and their Q-ball ODFs were used without being decomposed. The ODF was used without decomposition for voxels with $N = 1$, as well.

7.5 Increasing angular resolution

Similar to fuzzy connectedness tractography described in the previous chapter, MGT propagates paths from a voxel to its neighbours. Thus, the neighbourhood system determines the number of candidate propagation directions \mathbf{r} . A straightforward implementation uses the 3x3x3 neighbourhood, i.e. 26 neighbours. The angular resolution of path propagation can be increased by increasing the neighbourhood size to 5x5x5, including next nearest neighbours (Fig. 7.3). This gives 124 neighbours, which reduce to 98 non-parallel propagation directions in 3D (e.g. directions to \mathbf{g} and \mathbf{d} are the same). The solid angle is then $\omega = 4\pi/98$. However, neighbours that are not directly connected to the focal voxel \mathbf{i} exist, such as voxel \mathbf{j} in Fig. 7.3. If the weight w_{ij} is the highest amongst the other neighbours of \mathbf{j} , a strongest path will include an \mathbf{i} to \mathbf{j} transi-

tion, leaving a gap in the strength map. To avoid discontinuities in paths, one nearest neighbour of the focal voxel is required to be included in the path. More specifically, for each next nearest neighbour j (light gray in Fig. 7.3), the two nearest neighbours f and g (dark gray in Fig. 7.3) that are traversed by the vector \mathbf{r}_{ij} are considered. From amongst the fibre components of f and g , the one with the highest connection strength M_h to the seed s is then chosen. If the strength M_h is smaller than the strength of j , the value $M_{j...s}$ is assigned to M_h . This assignment, for all identified gaps in the strength map, is performed at the end of the algorithm as a postprocessing step.




Figure 7.3: Different neighbourhood systems of focal voxel i . Left: $3 \times 3 \times 3$ neighbourhood, Right: $5 \times 5 \times 5$ neighbourhood. A subset of within-slice neighbours is presented. Black arrows represent candidate directions \mathbf{r} for path propagation. Nearest neighbours are coloured with dark gray, next nearest neighbours with light gray.

7.6 Q-ball fast marching tractography

For comparison purposes, the Q-ball driven fast marching tractography (Campbell et al., 2005) has been implemented. A $3 \times 3 \times 3$ neighbourhood was employed and ODFs were normalized and thresholded at their mean value (Campbell et al., 2005). Front propagation speed values from a voxel i were computed using the ODF value along directions $\mathbf{r}_{ij}, j \in \mathbf{F}_i^{neigh}$. The fast marching tractography algorithm (Parker et al., 2002) was used to propagate a front from a seed. A path $C_{j...s}$ was then defined for each image voxel j to the seed s by going backwards through the map of front arrival times. Paths that had a curvature change of more than 90° along their length were discarded. The connectivity of voxel j was the smallest speed value along the path $C_{j...s}$ and was used to identify the most relevant paths (Campbell et al., 2005); contrary to graph-based methods that use the path strengths (Iturria-Medina et al., 2007; Zalesky, 2008).

7.7 Data and processing

Simulations

Phantoms containing two straight tracts, crossing at various crossing angles (60° to 90°), were numerically generated (Fig. 7.4). The noise-free DW signals in each voxel were simulated using the DTI model (Basser et al., 1994a) in single-fibre voxels and a mixture model (Eq. (7.6)) in crossing voxels. Sixty-one evenly spaced DW directions (Cook et al., 2007) with $b=3000 \text{ s/mm}^2$ were used. The individual tensors had fractional anisotropy $FA=0.8$ and $\text{trace}=2.1 \times 10^{-3} \text{ mm}^2/\text{s}$.

Zero-mean Gaussian noise was added in quadrature (Kingsley, 2006b) to simulate the Rician nature of the MRI signal (Gudbjartsson and Patz, 1995). Q-ball ODFs were estimated using a real spherical harmonics (SH) basis with Laplace-Beltrami regularization (regularization parameter=0.006) (Descoteaux et al., 2007) and a maximum SH order of 6. Up to two peaks were identified per ODF ($1 \leq N \leq 2$), so that a two-tensor model was fitted in crossing regions. No structural maps were used in the graph weights (Eq. (7.5)) for the simulations. MGT was quantitatively compared against graph-based tractography (GT) (Iturria-Medina et al., 2007) and Q-ball driven fast marching tractography (FM) (Campbell et al., 2005).

In all algorithms, the obtained paths were smoothed for visualization purposes using a moving average filter, similarly to smoothing filters used before (Campbell et al., 2005; Parker et al., 2002; Staempfli et al., 2006; Zalesky, 2008). For the n^{th} point (x_n, y_n, z_n) along a curve, the filtered coordinates were $(\sum_{i=n-h}^{n+h} x_i, \sum_{i=n-h}^{n+h} y_i, \sum_{i=n-h}^{n+h} z_i)$, with $2h+1$ being an odd number, representing the width of the filter. A width of 9 was used here.

In-vivo images

Q-ball whole-brain scans of five healthy subjects (4 males, 1 female) were performed; local research ethics committee approval and informed consent were obtained. A single-shot, spin-echo, echo-planar, diffusion-weighted sequence was used (acquisition matrix 112×112 with in-plane resolution $2 \times 2 \text{ mm}^2$) in a Philips 3T Achieva clinical imaging system. A parallel imaging factor of 2 was employed. Eight $b=0 \text{ s/mm}^2$ images

were acquired and averaged. Diffusion weighting was applied in 61 evenly spaced directions (Cook et al., 2007) with $b=3000 \text{ s/mm}^2$ ($TE=72 \text{ ms}$, $TR=15292 \text{ ms}$). Fifty two interleaved slices were acquired with a thickness of 2 mm. Also, a standard 3D high resolution T1-weighted image (MPRAGE) was acquired for each subject. An acquisition matrix of 256×256 with in-plane resolution $1 \times 1 \text{ mm}^2$ was used for 160 slices of 1 mm thickness in sagittal orientation.

DW images were corrected for eddy current distortion and motion using a linear affine registration implemented in FSL (Smith et al., 2004). Brain was extracted using BET (Smith et al., 2004). The DW scan for each subject was registered to the respective MPRAGE using a non-linear registration method (Andersson et al., 2007), to account for the non-linear EPI distortions. Due to its similar contrast to the MPRAGE, the FA image was used for the registration. Probabilistic tissue segmentation masks were obtained from the MPRAGEs using SPM5 (Ashburner and Friston, 2005) and tissue masks were transformed to DW space using the inverse registration. The WM masks were used in the calculation of graph weights (Eq. (7.5)) to restrict propagation within white matter. Q-ball ODFs were estimated as in simulations. Up to three peaks were identified per ODF ($1 \leq N \leq 3$), so that a two or three tensor model was fitted in crossing regions and the diffusivity terms (Eq. (7.4)) of the graph weights were obtained.

Thalamic parcellation

MGT's ability to perform thalamic parcellation was tested. The cortex was divided into seven regions (Johansen-Berg et al., 2005) and the thalamus was segmented from the MPRAGE image using FSL's model-based subcortical segmentation (Patenaude et al., 2007). Thalamic and cortical masks were transformed to the DW space using the previously calculated non-linear registration and a nearest-neighbour interpolation. Each thalamic voxel was used as a seed for MGT and connection strength maps were calculated. Each seed was assigned the colour code of the most strongly connected cortical region. To allow propagation within the isotropic thalamus and cortex, both WM and GM masks were included in the weights ($P_{mat}(\mathbf{i}) = P_{WM}(\mathbf{i}) + P_{GM}(\mathbf{i})$ in Eq. (7.5)). To avoid multi-peaked ODFs that may erroneously arise from a noisy isotropic profile, the DTI model (Basser et al., 1994a) was fitted in all GM voxels and a diffusion ODF was obtained using Eq. (7.7). That allowed the algorithm to propagate along one orienta-

tion within the thalamus and reach WM. Path propagation was stopped whenever a cortical voxel was reached.

7.8 Results

Simulations

Fig. 7.4 presents the performance of different tractography algorithms on four crossing configurations with crossing angle 90° , 80° , 70° and 60° . A horizontal bundle *A* crossing with a diagonal bundle *B* were simulated; with the orientation of *B* changing. No noise was added initially, so that only the limitations of each algorithm are illustrated. The algorithms tested were the graph-based tractography (GT) (Fig. 7.4a), fast marching tractography (FM) (Fig. 7.4b), multigraph-based tractography with a $3 \times 3 \times 3$ neighbourhood (MGT₃) (Fig. 7.4c) and multigraph-based tractography with a $5 \times 5 \times 5$ neighbourhood (MGT₅) (Fig. 7.4d). In each case, the strongest paths arising from two end ROIs (indicated with green arrows, located within both bundles *A* and *B*) to the seed ROI (indicated with a black arrow, located in the horizontal bundle *A*) are plotted. The paths are colour-coded with the path strength (or connectivity in the case of FM), as this varies along their length. All the above algorithms are distributed so all voxels will be visited by a path of some strength (Fig. 7.1). However, since the seed is in the horizontal bundle, we would expect to see much higher path strengths within *A* rather than within the crossing bundle *B*.

This is not always the case for GT and FM, as suggested before (Fig. 7.1). As shown in Fig. 7.4a and Fig. 7.4b, the difference ΔM between the maximum path strengths in the two crossing bundles becomes quickly negligible, as crossing angle decreases. In the case of FM, the connectivity reduction after the crossing is due to the speed function utilized (Campbell et al., 2005). For both algorithms, ΔM is in the order of 6 – 8% for a 70° crossing. For a 60° crossing, GT fails to propagate within bundle *A*, while ΔM is roughly 0 for FM. Highly connected regions to the seed are, thus, falsely identified in these noise-free simulations.

On the contrary, both MGT implementations (Fig. 7.4c and Fig. 7.4d) succeed in propagating correctly within bundle *A* and in assigning to it much higher strengths

compared to *B*. The benefits from increasing the angular resolution are shown in Fig. 7.4d. Even if the non-relevant paths to the crossing bundle *B* are much smoother, the low paths strengths are retained.




Figure 7.4: Results of different tractography algorithms on four different crossing configurations (crossing angle varied from left to right from 90° to 60°). (a) Graph-based tractography (GT), (b) Fast marching tractography (FM), (c) Multigraph-based tractography with $3 \times 3 \times 3$ neighbourhood (MGT_3) and (d) Multigraph-based tractography with $5 \times 5 \times 5$ neighbourhood (MGT_5). In all cases, the same seed indicated by the black arrow in (a) has been used. Two target ROIs within the individual crossing bundles are chosen in each case, indicated by green arrows in (a). Paths are plotted between the seed and the target ROIs for each tractography algorithm. Paths are colour-coded with the path strengths for (a), (c) and (d) and with the connectivities for (b). Phantoms are shown as opaque volumes.

The behaviour of all algorithms on these crossing configurations was further tested against noise. The difference ΔM between the maximum path strength in bundle *A* and the maximum path strength in bundle *B*, expressed as a percentage of the strength in *A*, was studied across simulations. The seed ROI was kept the same as before. Tables 7.1 and 7.2 show the mean and standard deviation of ΔM across 100 Monte Carlo realizations of each phantom, for different tractography algorithms and for different noise levels (signal to noise ratio (SNR) of 40 and 20, respectively). The qualitative findings from the noise-free simulations are further verified and quantified. ΔM is close to zero or non-positive in many cases for FM and GT, especially for crossing angles smaller than or equal to 70° , indicating that propagation within crossing bundle *B* is favoured over propagation within *A*. On the contrary, when MGT is used, ΔM is always positive (on average) and at least 33%, in most cases. For 60° and SNR=40 (Table 7.1), MGT returns path strengths in bundle *A* 33% higher than in *B*, when both FM and GT wrongly favour propagation to *B* ($\Delta M = -1.9\%$ and -39% respectively). MGT returns relatively low ΔM values only in the case of SNR=20 and 60° crossing (Table 7.2). These low values are an artefact of the peak finding algorithm that fails, due to noise, to identify two peaks in some crossing voxels. In these voxels, MGT reduces to GT and favours the wrong propagation (as GT does). As a result, the mean ΔM values are decreased (Table 7.2) and the standard deviation is increased (Table 7.1), when a 60° crossing phantom is used and noise influences model selection.

Crossing Angle	GT	FM	MGT ₃	MGT ₅
90°	58.74% (8.16)	31.53% (3.28)	194.2% (10.29)	170.37% (9.5)
80°	47.68% (14.55)	18.52% (2.2)	146.67% (8.33)	84.76% (4)
70°	4.53% (2.79)	5.67% (1.92)	59.03% (4.24)	57.32% (5.22)
60°	-39.03% (6.61)	-1.91% (1.44)	33.15% (11.6)	37.17% (15.46)

Table 7.1: Mean (st. deviation) difference ΔM between path strengths in horizontal bundle *A* and crossing bundle *B*, for various algorithms and crossing angles, across 100 simulations at SNR=40.

It should be noticed that the standard deviation of ΔM is always smaller in FM compared to all other methods. This reflects the filtering that has been performed on

Crossing Angle	GT	FM	MGT ₃	MGT ₅
90°	46.05% (12.24)	25.02% (5.42)	161.31% (20.1)	133.26% (18.5)
80°	20.26% (9.5)	13.43% (3.48)	114.54% (14.2)	74.02% (8.29)
70°	0.003% (4.67)	2.09% (2.66)	45.85% (10.39)	43.06% (10.78)
60°	-31.84% (9.06)	-3.5% (3.01)	6.19% (10.7)	12.02% (8.7)

Table 7.2: Mean (st. deviation) difference ΔM between path strengths in horizontal bundle **A** and crossing bundle **B**, for various algorithms and crossing angles, across 100 simulations at SNR=20.

the FM connectivity values, which represent the minimum propagation speed along a path. A similar reduction of the standard deviation will occur if a minimum operation is used for all the graph-based methods, as shown in (Iturria-Medina et al., 2007).

In-vivo human data

The performance of MGT in human data was explored within the centrum semiovale, where callosal tracts cross with corona radiata and longitudinal tracts. For the ROIs chosen, erroneously elevated connection strengths were more evident between the corpus callosum and the corona radiata and the focus is therefore on these two tracts. Fig. 7.5 shows the diffusion ODFs estimated in this region. The individual ODF components, estimated in crossing regions using the multi-tensor model, are shown in blue. These have replaced the multi-peaked Q-ball ODFs. The single-peaked Q-ball ODFs, which were left unchanged, are shown in red.

Fig. 7.6 shows a comparison between GT and MGT, when a seed ROI in the body of the corpus callosum (CC) is used. Coronal maximum intensity projections of strength maps are illustrated. Graph-based tractography (Fig. 7.6a) exhibits similar behaviour as in simulations, with paths in the crossing corona radiata having relatively high strengths; almost equal in magnitude and in some cases higher than the strengths of lateral callosal paths (indicated by black arrows). Using MGT (Fig. 7.6b), propagation within lateral callosal tracts is favoured over the crossing corona radiata. Increasing the angular resolution of propagation (Fig. 7.6c), allows the algorithm to visit, through relatively strong paths, even more CC voxels. Despite the smoother path




Figure 7.5: Diffusion ODFs utilized during MGT execution. Regions from two coronal slices are shown, one at the level of the genu of the corpus callosum (top) and another at the corpus callosum body (bottom). Q-ball ODFs are depicted with red. Multiple single-peaked ODF components, estimated via the multi-tensor model, are shown in blue. All ODFs are sharpened (raised in the third power) and multiplied by the P_{WM} mask, so that gray matter ODFs are suppressed. The FA map is shown at the background.

transitions, strengths within the corona radiata are suppressed compared to strengths obtained using GT (Fig. 7.6a).

For a qualitative comparison, Q-ball probabilistic tractography, as implemented in CAMINO (Cook et al., 2006), was applied on the same seed. The probabilistic index of connectivity is presented in Fig. 7.6d. We can observe the systematic reduction




Figure 7.6: Connection strength maps obtained when seeding in the body of the corpus callosum. A maximum intensity projection along a coronal plane is shown. Results obtained using (a) GT, (b) MGT_3 and (c) MGT_5 . In (d) the index of connectivity obtained through Q-ball probabilistic tractography is shown, superimposed on the FA.

of connectivity with distance, which is not inherently present in either MGT or GT algorithms.

The aim of computing an index of connectivity is to perform anatomical inference; regions that are connected to the seed through anatomically plausible paths, should have the highest connectivity values. To explore this, the strongest paths arising from the same seed ROI in the body of the CC, with different tractography methods (Fig. 7.7), were plotted. For each method, a threshold was manually selected, so that only paths within the CC remained (propagation was stopped just after paths started to appear within the crossing corona radiata). We can observe that for GT (Fig. 7.7a) and FM (Fig. 7.7b), these strongest paths include mainly the medial portion of the corpus callosum. A larger CC portion is captured with MGT_3 (Fig. 7.7c), with some lateral CC

paths having high connection strengths. The increased angular resolution MGT₅ (Fig. 7.7d) allows a CC reconstruction that covers larger cortical areas in both hemispheres. Both medial and lateral CC paths are given larger connection strengths than the corona radiata paths.




Figure 7.7: Most strongly connected paths to a seed ROI in the corpus callosum obtained using (a) GT, (b) FM, (c) MGT₃ and (d) MGT₅. Paths are colour-coded with their strength along their length (or connectivity value for FM). A threshold has been set in each case, so that paths are stopped just after propagation started to appear within the corona radiata. The thresholds correspond to the following distribution percentiles for each case: (a) P_{97.5}, (b) P₉₇, (c) P₉₆ and (d) P₉₅. Paths are superimposed on fractional anisotropy images. The seed ROI is shown as a white rendered volume in the body of the corpus callosum.

Given the discrete nature of MGT in both angular and spatial fields, it would be challenging to test it against a highly curved tract. Fig. 7.8 shows the performance of MGT₅ when seeding within the cingulum bundle. The raw strength values are shown in Fig. 7.8a. A threshold was interactively chosen based on anatomical knowledge of cingulum's trajectory (Mori et al., 2005) and the top 2.5% strengths are shown in Fig.




Figure 7.8: Connection strength maps obtained when seeding in the cingulum. Maximum intensity projections along a sagittal plane are shown. Results are obtained using MGT_5 . (a) Raw strength values, (b) Thresholded strengths using the 97.5th percentile ($P_{97.5}=0.27$), superimposed on an anisotropy image.

7.8b. The course of the cingulum starting from the sub-genual region, propagating along the cingulate gyrus and ending up at the parahippocampal gyrus is illustrated.

MGT was further tested against its ability to perform thalamic parcellation. The cortex was divided into seven regions, shown colour-coded in Fig. 7.9a. Fig. 7.9b shows an axial slice of the MPRAGE image along with the thalamic parcellation obtained using the high angular resolution version of the proposed algorithm (MGT_5). There is a clear clustering of the thalamic voxels, with the more posterior and medial ones being more strongly connected to the temporal lobe and the more anterior ones being connected to the frontal lobe. This organization agrees reasonably with Oxford's thalamic atlas (Johansen-Berg et al., 2005), which has been registered from the MNI space to the subject's MPRAGE space (Fig. 7.9c). We can observe that the thalamic clusters follow from posterior to anterior the temporal, occipital, parietal, sensory, motor, premotor and frontal lobes. Fig. 7.10 presents the connectivity-based thalamic parcellation on different axial slices in the original DWI space. We can observe that the clustering is not so consistent for a few voxels in the medial thalamus. These voxels have a relatively low anisotropy, while some of them, being adjacent to the ventricles, might have been considered as thalamic voxels, due to errors in the segmentation and registration processes.

MGT was compared with GT in a group of subjects, to examine how reproducible the benefits from using the former are (Fig. 7.11). For each subject, an axial ROI covering the most inferior and medial part of the body of the corpus callosum was used as a seed for both GT and MGT_5 . Computation time was in the order of 10 and 30 seconds for GT and MGT_5 per subject, on a 2.8 GHz PC. Fig. 7.11 shows coronal maximum




Figure 7.9: Thalamic parcellation using cortical masks. (a) Cortex is divided in seven regions. (b) Thalamic parcellation using MGT₅. Each thalamic voxel is assigned the colour code of the cortical region, to which it is most strongly connected. Results are superimposed on the MPRAGE image. (c) Oxford's thalamic MNI atlas, registered on subject's MPRAGE.




Figure 7.10: Thalamic clusters obtained using MGT₅ on different axial slices. Inferior to superior slices are presented from left to right. Results are presented on the DWI space and are superimposed on the $b=0$ s/mm² image.

intensity projections of connection strengths for each subject. Strengths obtained from each algorithm and subject were thresholded at the 95th percentile of their distribution; thus, the top 5% connected voxels are illustrated. We can observe that GT assigns very high strengths to the medial callosal tracts, but the more lateral tracts are suppressed in favour of crossing tracts. This is not the case for MGT that reproduces a similar CC pattern across all subjects, comprised of both medial and lateral tracts. MGT performs better due to its ability to follow an appropriate, rather than any, fibre population in a crossing region.




Figure 7.11: Connection strength maps obtained using a seed ROI in the body of the corpus callosum and GT (left column) or MGT₅ (right column), for five different subjects. Strengths have been thresholded at the 95th percentile (P₉₅) in both cases and a coronal maximum intensity projection operation has been performed.

7.9 Discussion

A distributed tractography algorithm that finds the strongest path connecting each voxel to a seed, has been presented. Multigraph-based tractography (MGT) utilizes graph theory and the orientation information provided by Q-ball images to propagate through fibre crossings. The advantage of the new method is in assigning appropriate path strengths when a fibre crossing is encountered and reducing strengths of less anatomically plausible connections. Previous methods that utilize Q-ball ODFs (Campbell et al., 2005; Iturria-Medina et al., 2007) allow directionally unconstrained, except for a hard 90° limit, propagation to all crossing populations in crossing regions. This can lead to falsely elevated connection strengths and erroneous propagation within the wrong bundle, as shown in simulations (Fig. 7.4) and in real data (Figs. 7.6 and 7.11). Instead, a multigraph representation of the image is employed to account for multiple fibre populations coexisting in voxels of crossing regions. Given the trajectory so far, paths to the appropriate bundle dominate and false positive connections are reduced (Figs. 7.4 and 7.6). A model-based method is used to estimate, within voxel, different ODF components for different crossing populations, if any. The Q-ball ODFs are used to extract a model selection mask. The path propagation algorithm is modified and in crossing regions propagation to each crossing fibre component is explored individually. This is performed by utilizing the individual single-peaked ODF components rather than using a multi-peaked ODF as a whole.

Quantitative tractography algorithms provide indices of connectivity that can be utilized to aid inference on the underlying WM anatomy. However, due to limitations imposed by the tracking algorithm and/or the scanning protocol, this is not always feasible. A more indirect method commonly employed uses, apart from the connectivity information, inclusion, exclusion, or target ROIs (Hagmann et al., 2003), which require a-priori information. The aim of this method was to effectively improve the ability of performing inference without using any extra information, by reducing limitations arising from the tracking algorithm. Fig. 7.7d shows the large coverage of the CC that can be obtained using MGT by looking at the top connections to a seed placed in the CC body. Note that no other ROI is used and that for the same seed other methods suppress lateral callosal paths in favour of crossing corona radiata paths (Figs. 7.7a and 7.7b).

In the implementation presented in this chapter, diffusion ODFs govern path propagation. Fiber ODFs (fODFs) (Descoteaux et al., 2009; Tournier et al., 2007) can be used instead to compute arc weights; the latter, being a sharpened version of the diffusion ODFs (Descoteaux et al., 2009), provide more directionality and suppress more perpendicular diffusivity. Fig. 7.12 shows connection strength maps obtained when seeding in the corpus callosum, using MGT with ODFs and with fODFs respectively. The fODFs were computed using the deconvolution approach of (Descoteaux et al., 2009). In the case of multi-peaked fODFs, the multi-tensor model of Eq. (7.6) was fitted and single-peaked ODF components were obtained using Eq. (7.7). These were then deconvolved to give single-peaked fODF components. The extra directionality is exploited to increase the lateral CC portion that is assigned high path strengths.




Figure 7.12: Connection strength maps obtained using MGT₅ and (a) diffusion ODFs, (b) fiber orientation distribution functions (fODFs). Coronal maximum intensity projections are shown. A seed ROI is placed in the body of the corpus callosum. Strengths are thresholded at the 95th percentile in (a) and at the 90th percentile in (b). Thresholds were determined so that no propagation within the crossing corona radiata occurred in each case.

Comparison with other tractography methods

MGT has retained the advantages of graph-based tractography over streamline-based approaches (Behrens et al., 2003; Parker et al., 2003). Information from other imaging modalities on the voxel level can be integrated via a simple modification of the arc weights (Eq. (7.5)). Other information, such as prior anatomical knowledge, can be incorporated in a similar way. MGT is also relatively fast and avoids a systematic reduction of path strengths with distance from the seed, a common problem of probabilistic

tractography (Fig. 7.6d). Furthermore, due to the distributed nature of graph-based algorithms, fibre branching is inherently considered. The improved performance of MGT will also be beneficial to graph-based tractography applications on studying the human brain connectome (Iturria-Medina et al., 2008).

MGT effectively borrows the idea of multi-tensor (Behrens et al., 2007; Parker and Alexander, 2003) and Q-ball (Seunarine et al., 2007) probabilistic tractography. A different orientation uncertainty component is associated with each crossing population in these methods. Similarly, a different arc weight is assigned to each crossing population in MGT. However, path propagation is different. Streamline-based methods follow either of two alternatives within a crossing region: a) a crossing population is chosen and propagation is allowed only towards it (Parker and Alexander, 2003), b) all crossings populations that do not violate an arbitrarily set curvature threshold criterion are chosen and propagation is allowed towards all of them (Chao et al., 2008; Descoteaux et al., 2009). Given that acute angle crossings cannot always be differentiated from branchings, the former approach may underestimate connectivity in branching regions. This will be remedied by the latter approach, which, however for the same reason, may overestimate connectivity in crossing regions. The multigraph-based tractography algorithm lies in the middle of these two alternatives. All crossing/branching populations will be considered, but each with a different weight, depending on the trajectory so far. The weights reflect the prior belief that WM paths are smooth curves and do not change direction rapidly. In the presence of a crossing, propagation strengths to one population will dominate. In the presence of a branching, propagation strengths to all populations equally commensurate with the path will be of the same magnitude.

MGT was compared with two other distributed tractography techniques, one based on graph theory (Iturria-Medina et al., 2007) and another on front propagation (Campbell et al., 2005). In general, graph-based tractography has been shown to be more robust against noise compared to fast marching algorithms (Iturria-Medina et al., 2007). It also avoids the excessive branching that FM exhibits from the fastest route of front propagation, a limitation that has been discussed before (Jackowski et al., 2005) and is also evident in the simulations presented here (Fig. 7.4b). MGT keeps the advantages of GT and improves on that when a crossing is encountered. A detailed quantitative comparison with probabilistic tractography hasn't been performed, due to the

algorithms' different nature. Probabilistic tractography provides voxel visitation maps and computes the probability of connectivity, by repeating an experiment and counting the number of times an event has occurred. Distributed tractography techniques provide both paths and indices of connectivity along a path, which, however, encode a different type of information. The fuzzy nature of the acquired images is reflected in these indices, extending more naturally the binary framework of deterministic streamline tractography rather than converging to the results of probabilistic.

Nevertheless, a qualitative comparison was attempted (Fig. 7.6) between MGT and probabilistic tractography. To explore the abilities of graph-based tractography, thalamic parcellation was performed, using the path strengths as an index of connectivity (Fig. 7.9). The revealed thalamic organization (Fig. 7.9b) agrees well with Oxford's thalamic atlas (Johansen-Berg et al., 2005), obtained with probabilistic tractography (Fig. 7.9c). The fact that the atlas is an average of multiple subjects contributes to the observed differences. There are also differences due to the type of information each connectivity index encodes. Instead of simply using the path strength as an index of connectivity, other alternatives include the weakest link along a path (Campbell et al., 2005; Iturria-Medina et al., 2007) or the degree of agreement of the reconstructed trajectory with the underlying diffusion profiles (Jackowski et al., 2005; Parker et al., 2002). These indices are less curvature dependent and distance-independent and might influence the GM parcellation. Furthermore, some thalamic voxels exhibited high connection strengths with more than one cortical region. A fuzzier colour-coding scheme (Draganski et al., 2008) might therefore be more appropriate than the crisp one employed here. These different options remain to be explored.

Assumptions and future work

A mixture model of N prolate tensors has been utilized to obtain individual ODF components, when multiple ODF peaks are detected. When one peak is found, the ODF is used as it is. This choice accommodates other complex fibre configurations that may be present in a voxel, such as fibre fanning and fibre bending. In these cases, the ODF tends to have a single broad peak (Seunarine and Alexander, 2009).

Multi-peaked ODFs are indirectly decomposed to single-peaked components using a model-based approach. A multi-tensor model is fitted, with the orientations of the

tensors being fixed to the ODF peaks. This approach reduces the number of unknown parameters and the dimensionality of the non-linear model. An alternative has been presented in (Sotiropoulos et al., 2009), where a common single-fibre ODF kernel was estimated from the top anisotropic voxels in the brain and aligned with each ODF peak in each crossing voxel. This removed the need for fitting a multi-tensor model, which, however, allows the estimation of single-fibre ODFs that are more representative of the local data.

Furthermore, assuming that the ODF local maxima are good estimates of the underlying crossing orientations may not always be valid and may provide biased estimates (Zhan and Yang, 2006). The ODF decomposition approach described in (Schultz and Seidel, 2008) corrects for the direction bias and can be employed with MGT, instead of the mixture model used here.

To alleviate the effects of discrete path propagation, the angular resolution of the algorithm was enhanced by increasing the neighbourhood size. Smoother paths could be reconstructed and more voxels were strongly connected to the seed in the case of curved tracts (Figs. 7.4 and 7.7). However, an increase in the angular resolution comes with an increased curvature dependence of the arc weights. Fig. 7.13 shows how an arc weight between two adjacent voxels varies with the relative orientation difference of the respective ODF peaks. For a 3x3x3 neighbourhood, graph weights can tolerate up to 22° degrees of curvature between adjacent voxels, before they start decreasing from a value of one. For a 5x5x5 neighbourhood, this curvature value is reduced to 12°. This curvature dependence, present in both GT and MGT, is due to the imposed symmetry in the graph weights. A weight w_{ij} will be one, only if the solid angle cones from the two voxels contain the peaks of the respective ODFs. Increasing the neighbourhood size, reduces the solid angle and therefore reduces the chance of overlapping between the ODFs in i and j . Nevertheless, even in cases of highly curved tracts, such as the cingulum (Fig. 7.8), MGT assigned to it the top strength values that differentiated this bundle from the background. The curvature dependence can be removed, if no symmetry is imposed on the graph weights. A further refinement to the discrete path propagation can be the geometric correction scheme proposed in (Bodammer et al., 2009). This assigns different solid angles to different neighbours (i.e. face-to-face neighbours subtend a larger ω than corner-to-corner ones) and takes into account the geometry of

the acquisition grid.




Figure 7.13: Curvature dependence of symmetric arc weights. The difference between the ODFs orientation in two adjacent voxels is plotted versus the calculated arc weight, for two different neighbourhood systems. Black: 3x3x3 neighbourhood, solid angle $\omega = 4\pi/26$, Gray: 5x5x5 neighbourhood, solid angle $\omega = 4\pi/98$.

Similar to all quantitative tractography methods, including probabilistic approaches, there is an open question on how to identify the background connectivity and extract only the relevant paths. For visualisation purposes, a simple threshold was used to interactively choose the top connected voxels to the seed. Previous studies have similarly used user determined thresholds to identify anatomically plausible paths (Campbell et al., 2005; Iturria-Medina et al., 2007; Parker et al., 2002; Staempfli et al., 2006). More advanced methods will include the development of a null connectivity model to test a hypothesis against, similar to the recent work by (Clayden and Clark, 2009; Morris et al., 2008).

Summary

In summary, a distributed multigraph-based tractography algorithm for exploiting Q-ball imaging to propagate through fibre crossings has been presented. Limitations of previous distributed tracking methods that utilize Q-ball images were discussed and

a new approach was proposed. The new algorithm was compared both quantitatively and qualitatively with two other methods, using simulated data, as well in vivo data acquired from a group of healthy subjects. MGT propagates through complex fibre configurations, and assigns to paths strengths that are more representative of their anatomical plausibility. Furthermore, the proposed method was shown to have a potential on applications tackled so far only with streamline-based approaches, such as thalamic parcellation.

Appendix A: Multigraph-based tractography algorithm

The algorithm described here (MGT) uses an image multigraph G and a set of arc weights w as inputs, to compute the strongest path from a seed voxel s to all N fibre components of all image voxels i . The strengths $M_{i\dots s}^n$ of these strongest paths are also returned. For notational simplicity, $M_i^n \equiv M_{i\dots s}^n$ and $Path_i^n \equiv Path_{i\dots s}^n$. To build an image of path strengths, the largest strength M_i is also returned for each voxel i . A priority queue, implemented as a heap (Kruse and Ryba, 2000), is used as an auxiliary data structure. A function **Valid()** is used to identify in each propagation step, neighbours to which a transition will cause a curvature change of less than 90° ; propagation only to these neighbours is considered. The algorithm reduces to graph-based tractography (Iturria-Medina et al., 2007), if no crossings exist (i.e. $N = 1$ for all voxels). A $3 \times 3 \times 3$ neighbourhood is assumed, denoted by F_i^{neigh} ; slight modifications described in section 7.5 should be made for a larger neighbourhood system.

Algorithm 7.1 Multigraph-based Tractography

```

1: for all image voxels  $i$  do
2:   for each ODF component  $n$  of  $i$  do
3:      $M_i^n = \begin{cases} 0, & \text{if } i \neq s; \\ 1, & \text{if } i = s; \end{cases}$ 
4:      $Path_i^n = \text{Empty}$ ;
5:     Enqueue  $(i, n)$ ;
6:   while Queue  $\neq \text{Empty}$  do
7:     Dequeue element  $(i, n)$  with maximum  $M$  value;
8:      $V_i^{neigh} = \text{Valid}(i, n, F_i^{neigh})$ ;
9:     for all neighbours  $j \in V_i^{neigh}$  do
10:      for each ODF component  $k$  of  $j$  such that  $(j, k) \in \text{Queue}$  do
11:         $M_{trial} = M_i^n \cdot w_{inj_k}$ ;
12:        if  $M_{trial} > M_j^k$  then
13:           $M_j^k = M_{trial}$ ;
14:           $Path_j^k = \{Path_i^n \cup a(i, j)\}$ 
15:   for all image voxels  $i$  do
16:      $M_i = \max_n(M_i^n)$ ;

```

Function 7.2 Valid $\left(\mathbf{i}, n, \mathbf{F}_i^{neigh}\right)$

```
1:  $\mathbf{W} = \text{Empty};$ 
2: if  $\mathbf{i} = \mathbf{s}$  then
3:    $\mathbf{W} = \mathbf{F}_i^{neigh};$ 
4: else
5:    $\mathbf{r}_{last} = \text{last direction in } Path_i^n;$ 
6:   for all  $\mathbf{j} \in \mathbf{F}_i^{neigh}$  do
7:      $\text{dot} = \mathbf{r}_{last} \cdot \mathbf{r}_{ij};$ 
8:     if  $\text{dot} > 0$  then
9:        $\mathbf{W} = \{\mathbf{W} \cup \mathbf{j}\};$ 
10: return  $\mathbf{W};$ 
```

Appendix B: Fast-marching tractography algorithm

The fast marching algorithm (Sethian, 1996) utilizes a speed function v to propagate a front from a seed or a group of seeds \mathbf{s} . The implementation presented here, following the one presented in (Staempfli et al., 2006), computes a front arrival time T for all image voxels. Front speed values $v(\mathbf{i}, \mathbf{r}_{ij})$ are calculated between a voxel \mathbf{i} and each 3x3x3 neighbour \mathbf{j} , using the diffusion ODF at \mathbf{i} along the unit vector \mathbf{r}_{ij} (Campbell et al., 2005); \mathbf{r}_{ij} indicating the direction between the centres of \mathbf{i} and \mathbf{j} . During propagation, three sets of voxels are utilized, voxels already visited by the front, voxels lying just outside the front (the narrow band) and the remaining voxels, as has been shown in Fig. 4.5. Members of the narrow band are candidates for inclusion in the front, which propagates in each step towards the maximum speed direction. A priority queue is used as an auxiliary data structure to store the elements of the narrow band. An index $Front(\mathbf{i})$ indicates whether a voxel \mathbf{i} belongs to the front. An array $Previous(\mathbf{i})$ stores the voxel through which \mathbf{i} entered the front, while $v_{link}(\mathbf{i})$ stores the speed value for this transition.

Algorithm 7.3 Fast Marching Tractography

```

1:  $T(\mathbf{i}) = \infty$  and  $Front(\mathbf{i}) = 0 \ \forall$  voxel  $\mathbf{i}$ ;
2:  $Previous(\mathbf{i}) = undefined$  and  $v_{link}(\mathbf{i}) = 0 \ \forall$  voxel  $\mathbf{i}$ ;
3: for all seed voxels  $\mathbf{s}$  do
4:    $T(\mathbf{s}) = 0$  and  $Front(\mathbf{s}) = 1$ ;
5:    $v_{link}(\mathbf{s}) = \infty$ ;
6:   Enqueue  $\mathbf{s}$ ;
7: while Queue  $\neq$  Empty do
8:   Dequeue  $\mathbf{i}$  with minimum  $T$  value;
9:    $Front(\mathbf{i}) = 1$ ;
10:  for all neighbours  $\mathbf{j}$  of  $\mathbf{i}$  with  $Front(\mathbf{j}) = 0$  do
11:     $T_{trial} = T(\mathbf{i}) + |\mathbf{j} - \mathbf{i}| / v(\mathbf{i}, \mathbf{r}_{ij})$ ;
12:    if  $T_{trial} < T(\mathbf{j})$  then
13:       $T(\mathbf{j}) = T_{trial}$ ;
14:       $Previous(\mathbf{j}) = \mathbf{i}$ ;
15:       $v_{link}(\mathbf{j}) = v(\mathbf{i}, \mathbf{r}_{ij})$ ;
16:      if  $\mathbf{j} \notin$  Queue then
17:        Enqueue  $\mathbf{j}$ ;
    
```

Paths C can then be found from any voxel to the seed by going backwards, using the *Previous* array. The connectivity of a voxel will be the weakest link along its path to the seed, i.e. the minimum of the v_{link} values. In case that two consecutive steps in C cause a curvature change equal or greater than 90° , the connectivity of the voxels linked through C to the seed is set to zero.

Appendix C: Multiple prolate tensor model

The multi-tensor model utilized here, assumes mixtures of cylindrically-symmetric tensors only. These are tensors with their second and third eigenvalues equal ($\lambda_2 = \lambda_3$). A cylindrically symmetric tensor with principal eigenvector \mathbf{e}_1 and eigenvalues $\lambda_1, \lambda_2, \lambda_2$ can be expressed as (Alexander, 2006):

$$\mathcal{D} = (\lambda_1 - \lambda_2)\mathbf{e}_1\mathbf{e}_1^T + \lambda_2\mathbf{I}, \quad (7.8)$$

with \mathbf{I} being the identity matrix. Let $\mathbf{u} = \sqrt{\lambda_1 - \lambda_2}\mathbf{e}_1$. Then, the tensor can be described

by just 4 parameters:

$$\mathcal{D} = \mathbf{u}\mathbf{u}^T + \lambda_2 \mathbf{I} . \quad (7.9)$$

If the tensor orientation \mathbf{e}_1 is fixed and only its eigenvalues are of interest, the following parameterization guarantees the non-negativity of the estimated eigenvalues:

$$\mathcal{D} = \alpha^2 \mathbf{e}_1 \mathbf{e}_1^T + \beta^2 \mathbf{I} . \quad (7.10)$$

Once α and β are estimated, the tensor eigenvalues are $\alpha^2 + \beta^2, \beta^2, \beta^2$.

a) Two-tensor model: A mixture model of two tensors \mathcal{D}_A and \mathcal{D}_B is given by the following:

$$S_k = S_0 \left[f \exp(-b \mathbf{g}_k^T \mathcal{D}_A \mathbf{g}_k) + (1 - f) \exp(-b \mathbf{g}_k^T \mathcal{D}_B \mathbf{g}_k) \right] , \quad (7.11)$$

where the tensors are parameterized as shown before. To increase the robustness of the fit, the difference α^2 between the principal and secondary eigenvalues is assumed to be the same for both tensors. Therefore, the tensors are:

$$\mathcal{D}_A = \alpha^2 \mathbf{e}_{A1} \mathbf{e}_{A1}^T + \beta_A^2 \mathbf{I} \quad \text{and} \quad \mathcal{D}_B = \alpha^2 \mathbf{e}_{B1} \mathbf{e}_{B1}^T + \beta_B^2 \mathbf{I} , \quad (7.12)$$

with the orientations \mathbf{e}_{A1} and \mathbf{e}_{B1} being fixed to the respective Q-ball ODF peaks. This two-tensor model is non-linear and is fitted to the measurements using the Levenberg-Marquardt algorithm (Press et al., 1992). During each iteration of the algorithm, the partial derivatives $\partial S_k / \partial \zeta$ over all model parameters ζ need to be evaluated for the b value and \mathbf{g}_k direction vector of all k diffusion-sensitizing gradients. To ensure that the volume fraction $f \in [0, 1]$, it is replaced by $f = \sin^2 \theta$. For the initialization of Levenberg-Marquardt, the least squares estimates of the single tensor model are used (Alexander and Barker, 2005). The volume fractions are set to 0.5. The secondary eigenvalues of both tensors are set equal to the smallest eigenvalue obtained from the DTI model ($\lambda_{A2} = \lambda_{B2} = \lambda_3$). Finally, the principal eigenvalues are set equal ($\lambda_{A1} = \lambda_{B1}$), such that the average diffusivity along DTI's \mathbf{e}_1 is DTI's λ_1 . This gives the following

equation:

$$0.5 \cdot ADC_{A1} + 0.5 \cdot ADC_{B1} = \lambda_1 \implies$$

$$\lambda_{A1} = \lambda_{B1} = \frac{2\lambda_1 - \lambda_3[1 - (\mathbf{e}_1 \cdot \mathbf{e}_{A1})^2] - \lambda_3[1 - (\mathbf{e}_1 \cdot \mathbf{e}_{B1})^2]}{(\mathbf{e}_1 \cdot \mathbf{e}_{A1})^2 + (\mathbf{e}_1 \cdot \mathbf{e}_{B1})^2},$$

where the following expression has been used for ADC_{A1} , the apparent diffusion coefficient of tensor \mathcal{D}_A along the direction of \mathbf{e}_1 :

$$ADC_{A1} = \lambda_{A1}(\mathbf{e}_1 \cdot \mathbf{e}_{A1})^2 + \lambda_{A2}[(\mathbf{e}_1 \cdot \mathbf{e}_{A2})^2 + (\mathbf{e}_1 \cdot \mathbf{e}_{A3})^2] \implies$$

$$= \lambda_{A1}(\mathbf{e}_1 \cdot \mathbf{e}_{A1})^2 + \lambda_{A2}[1 - (\mathbf{e}_1 \cdot \mathbf{e}_{A1})^2].$$

Then, the initial values for the model parameters will be $\alpha^2 = \lambda_{A1} - \lambda_3$ and $\beta_A^2 = \beta_B^2 = \lambda_3$.

b) Three-tensor model: A mixture model of three tensors \mathcal{D}_A , \mathcal{D}_B and \mathcal{D}_C is given by the following:

$$S_k = S_0 \left[f_1 \exp(-b \mathbf{g}_k^T \mathcal{D}_A \mathbf{g}_k) + f_2 \exp(-b \mathbf{g}_k^T \mathcal{D}_B \mathbf{g}_k) + (1 - f_1 - f_2) \exp(-b \mathbf{g}_k^T \mathcal{D}_C \mathbf{g}_k) \right], \quad (7.13)$$

where the tensors are parameterized as shown in Eq. 7.12, i.e. with a common α^2 parameter. To ensure that the volume fractions $\in [0, 1]$, the following assignments are performed $f_1 = \sin^2 \theta_1$ and $f_2 = (1 - f_1) \sin^2 \theta_2 = (1 - \sin^2 \theta_1) \sin^2 \theta_2$. For the initialization of Levenberg-Marquardt, the least squares estimates of the single tensor model are used (Alexander and Barker, 2005). The volume fractions are set to 1/3. For the eigenvalues, it is assumed that all tensors are anisotropic with eigenvalues ratio $\lambda_{j1}/\lambda_{j2} = k > 1$, with k being the same for all tensors \mathcal{D}_j , $j = A, B, C$. All principal eigenvalues are initialized at the same value ($\lambda_{A1} = \lambda_{B1} = \lambda_{C1}$), such that the average diffusivity along DTI's \mathbf{e}_1 is λ_1 . This gives the equation:

$$(1/3) \cdot ADC_{A1} + (1/3) \cdot ADC_{B1} + (1/3) \cdot ADC_{C1} = \lambda_1 \implies$$

$$\lambda_{A1} = \lambda_{B1} = \lambda_{C1} = \frac{3k\lambda_1}{(k-1)(\mathbf{e}_1 \cdot \mathbf{e}_{A1})^2 + (k-1)(\mathbf{e}_1 \cdot \mathbf{e}_{B1})^2 + (k-1)(\mathbf{e}_1 \cdot \mathbf{e}_{C1})^2 + 3}.$$

The parameter k that controls the anisotropy of the initial tensor estimates is set to $k = 8$, representative of the difference between the principal and secondary DTI eigenvalues in the highly anisotropic corpus callosum voxels (Pierpaoli et al., 1996).

Part IV

Conclusions and Bibliography

8

Conclusions

8.1 Summary and conclusions

This thesis is concerned with diffusion-weighted (DW) MRI of the brain and white matter tractography. DW-MRI introduces contrast to images utilizing the random motions of water molecules within tissue. For tissues exhibiting structural coherence, such as brain white matter, molecules are hindered in a systematic way and diffusion becomes anisotropic; meaning that it is observed more towards a specific orientation. This is assumed to indicate the orientation of the underlying neuronal fibres. Tractography utilizes these orientation estimates to reconstruct the route of fibre bundles. Due to the relatively large resolution of DW-MRI (usually 2 mm) compared to the diameter of a single axon (mostly <10 μm), tractography *provides evidence* for bundles of thousands of axons.

Methodological improvements and contributions for performing tractography were presented in the research chapters. The problem of resolving crossing fibre configurations has been presented in the background sections along with existing methods. The question that was addressed in Chapter 5 was "How well can we reconstruct cross-

ing fibre configurations from conventional diffusion tensor images?". A regularization scheme was employed to impose continuity of the diffusion process between neighbouring voxels and, thus, continuity of fibre orientations. Relaxation labelling was adapted to iteratively choose amongst a set of fibre orientations, the ones that preserve smooth changes across neighbours. Furthermore, a hybrid minimization scheme was proposed to speed up the process. The spatial regularization improved the orientation estimates obtained from a two-tensor model correcting for a) noise artefacts and b) artefacts caused by the non-linear fitting procedure. Improvements at a global scale were evident in tractography results. Using images acquired at low b value and in less than 6 minutes, lateral paths of the corpus callosum were reconstructed using the regularized orientations, while Q-ball imaging failed on the same dataset. Apart from orientations, revised anisotropy and diffusivity information could be also obtained. This chapter showed that since diffusion is continuous in space, we do not need to rely only on local estimates to recover fibre orientations and reconstruct white matter structures.

Two global tractography methods were presented in the next chapters. The question that was addressed in Chapter 6 was "Given that images are fuzzy, can we tackle the tractography problem using fuzzy theory and also incorporate crossing fibre information?". Fuzzy connectedness tractography (FCT) was presented as a method that searches for the weakest link of the strongest path on the image grid, that connects any image voxel to a seed. This link determines the connectedness of the voxel. Compared to probabilistic indices of connectivity, fuzzy connectedness reflects orientational coherence along a path and does not drop systematically with the distance from the seed. It is a relative measure that was shown to be able to differentiate between paths belonging to the same bundles from non-coherent paths. Even highly curved tracts, such as the fornix and the cingulum that are challenging to reconstruct with streamlines, could be followed. Compared to other global tracking algorithms, fuzzy connectedness tractography (FCT) propagated through fibre crossings, by considering all the crossing orientations, but in a weighted manner that reduced the chance of false positives. As FCT can be combined with any DW-MRI reconstruction method that provides $N \geq 1$ fibre orientation estimates, results were shown using DTI, Q-ball and multi-tensor orientations. Finally, the relationship between FCT and graph-based tractography was discussed in this chapter; FCT was shown to be effectively a graph searching procedure.

Graph-based tractography is a relatively new framework that has shown a lot of potential. In Chapter 7, the question that was addressed was "What can we gain by incorporating fibre crossing information and increasing the angular resolution of propagation in graph-based tractography?". The development of multigraph-based tractography (MGT) was presented. Instead of fibre orientations as in FCT, orientation distribution functions were utilized in MGT, to partially map the uncertainty in the estimates. Elements of the previous chapters, including multi-tensor models and multiple voxel instances in crossing regions, were employed. A brain image was represented as a multigraph, meaning that multiple arcs existed between voxels that contained crossing configurations. An algorithm to search the graph exhaustively for the strongest path connecting each voxel to the seed was presented. MGT incorporated information from different imaging modalities, as both anatomical T_1 -weighted and diffusion-weighted images were used in the construction of the multigraph. A quantitative and qualitative comparison to existing methods showed that MGT can propagate robustly through fibre crossings. Its potential was illustrated by performing connectivity-based thalamic parcellation, a challenging task that involves propagation within grey matter. To the author's best knowledge, this is the first time that tractography-based parcellation of a nucleus is obtained through a non-streamline tracking method. Given these results and the flexibility of the new framework in reflecting different types of information into the connectivity indices, MGT was shown to be a promising method for brain connectivity analysis.

8.2 Future perspectives

A common issue for all quantitative tractography methods is identifying statistically significant connections to a seed. With both frameworks presented in this thesis, an index of connectivity is computed for every image voxel. Probabilistic streamline tractography exhibits a similar behaviour, even if not all voxels are necessarily assigned a connectivity value. To identify which of these values have a physical meaning, a proper null hypothesis needs to be developed in order to test for significant anatomical connections. An effort towards this direction has been presented in (Morris et al., 2008) for probabilistic tractography and similar studies are needed for other quantitative trac-

tography methods.

A field that has recently drawn attention is integration of information from multiple imaging modalities to study brain connectivity. Approaches such as fuzzy, graph-based and multigraph-based tractography presented in this thesis allow such integration at the voxel level. Recent studies have attempted multi-modal information incorporation at the region level. In (Hagmann et al., 2008), functional connectivity, the temporal correlation of resting-state fMRI time series (Friston et al., 1993), is correlated with indices of structural connectivity obtained from streamline tractography. In (van den Heuvel et al., 2009), the structural connections between nodes of functional networks present during rest (Greicius et al., 2003) are identified. Frameworks that consider simultaneously the different types of information are needed, since structural and functional connectivity complement each other.

Towards this direction, whole-brain connectivity matrices as the ones presented in (Iturria-Medina et al., 2008), allow the computation of indices representing both direct and indirect structural connections (Sporns, 2002), in accordance with functional indices that depict direct or indirect functional connections. The improved performance of multigraph-based tractography is expected to improve the connectivity matrices of (Iturria-Medina et al., 2008) and thus subsequently calculated measures. It is also important to define appropriately the structural connectivity between regions of interest, in terms of the connectivity values of the comprising voxels. Different alternatives are proposed in (Iturria-Medina et al., 2008) and (Kreher et al., 2008b), however it is not yet clear which of them represents best the underlying anatomical connectivity. Further exploration is needed to answer this question.

There are also open questions in the utilization of the fibre and diffusion ODFs in performing tractography. Considering them as in the multigraph-based framework or in the curve propagation framework of (Descoteaux et al., 2009), maps some of the uncertainty of the fibre orientation estimates. Given though the deterministic computation of these functions, data uncertainty is not fully considered and reflected in the results. Incorporating the ODF uncertainty in tractography frameworks could have a dramatic effect on the interpretability of connectivity maps (Behrens and Jbabdi, 2009), since a distribution of connectivity indices could be obtained for each voxel. MCMC approaches to estimate this uncertainty have been suggested in (Fonteiijn et al., 2007;

Kaden et al., 2007), however their routine application is hindered by their computational requirements. More effective approaches can be based on the residual bootstrap (Haroon et al., 2009) or on exact Bayesian inference (Sotiropoulos et al., 2010), an approach recently suggested by our group.

Apart from crossing fibres, there are other complex configurations that are more difficult to resolve and are still a challenging problem. These include fibre fanning, bending and kissing (Seunarine and Alexander, 2009). Modelling these configurations at a sub-voxel scale seems to be a potential solution, as recently showed in (Nedjati-Gilani and Alexander, 2009). Furthermore, instead of using only local, within-voxel data, exploiting neighbourhood information can assist in the reconstruction of these configurations. The idea of spatial continuity of orientations that was utilized in this thesis was also employed in (Savadjiev et al., 2008) to disambiguate between fibre crossings, fibre fannings and single fibre configurations.

More studies are also needed on tractography validation. Given that it is impossible to validate the results within each living subject against a gold standard, prior anatomical knowledge is usually invoked as evidence of successful reconstruction. The simplest validation method is through computer simulations. Some groups have developed biologically-inspired phantoms for which the ground-truth is known (Campbell et al., 2005; Poupon et al., 2008). Recent studies have performed in-vitro (Dyrby et al., 2007) and in-vivo (Dauguet et al., 2007) tractography on animal brains and evaluate the overlap between tractography results and paths identified by injected axonal tracers. A qualitative comparison between tractography in living humans and white matter tracts dissected from post-mortem brains has been presented in (Lawes et al., 2008). Despite these attempts, quantitative studies that compare different tracking methods are missing.

The development of ultra-high field MRI scanners will enhance the potential of DW-MRI. The higher signal to noise ratio (SNR) and resolution that can be (theoretically at least) achieved, are expected to improve data quality and possibly increase the range of diffusion imaging applications. In (Mukherjee et al., 2008) preliminary results on 7T Q-ball imaging are presented. The higher SNR allowed the reduction of the number of DW directions and thus scanning time, while keeping the same resolving power in fibre crossings reconstruction. Such improvements may allow more routine

utilization of Q-space/Q-ball protocols. Along with continuing methodological developments, they can increase the impact of diffusion-weighted MRI and tractography on tackling neuroscience questions (Hagmann et al., 2008), but also on clinical applications (Ciccarelli et al., 2008).

Bibliography

- A. L. Alexander, K. M. Hasan, M. Lazar, J. S. Tsuruda, and D. L. Parker. Analysis of partial volume effects in diffusion-tensor MRI. *Magnetic Resonance in Medicine*, 45(5):770–80, 2001.
- D. C. Alexander. Multiple-fiber reconstruction algorithms for diffusion MRI. *Annals of the New York Academy of Sciences*, 1064:113–33, 2005a.
- D. C. Alexander. Maximum entropy spherical deconvolution for diffusion MRI. *Information Processing in Medical Imaging (IPMI)*, 19:76–87, 2005b.
- D. C. Alexander. An introduction to computational diffusion MRI: the diffusion tensor and beyond. In J. Weickert and H. Hagen, editors, *Visualization and Processing of Tensor Fields*, pages 83–106. Springer, Berlin Heidelberg, 2006.
- D. C. Alexander and G. J. Barker. Optimal imaging parameters for fiber-orientation estimation in diffusion MRI. *Neuroimage*, 27(2):357–67, 2005.
- D. C. Alexander, G. J. Barker, and S. R. Arridge. Detection and modeling of non-Gaussian apparent diffusion coefficient profiles in human brain data. *Magnetic Resonance in Medicine*, 48(2):331–340, 2002.
- A. W. Anderson. Measurement of fiber orientation distributions using high angular resolution diffusion imaging. *Magnetic Resonance in Medicine*, 54(5):1194–206, 2005.
- J. Andersson, M. Jenkinson, and S. Smith. Non-linear registration, aka spatial normalisation. Technical Report TR07JA2, Centre for Functional Magnetic Resonance Imaging of the Brain, Department of Clinical Neurology, Oxford University, 2007.
- J. Ashburner and K. J. Friston. Unified segmentation. *Neuroimage*, 26(3):839–51, 2005.
- Y. Assaf and P. J. Basser. Composite hindered and restricted model of diffusion (CHARMED) MR imaging of the human brain. *Neuroimage*, 27(1):48–58, 2005.
- Y. Assaf, D. Ben-Bashat, J. Chapman, S. Peled, I. E. Biton, M. Kafri, Y. Segev, T. Hendler, A. D. Korczyn, M. Graif, and Y. Cohen. High b-value q-space analyzed diffusion-weighted MRI: application to multiple sclerosis. *Magnetic Resonance in Medicine*, 47(1):115–26, 2002.
- Y. Assaf, R. Z. Freidlin, G. K. Rohde, and P. J. Basser. New modeling and experimental framework to characterize hindered and restricted water diffusion in brain white matter. *Magnetic Resonance in Medicine*, 52(5):965–78, 2004.
- P. J. Basser. Inferring microstructural features and the physiological state of tissues from diffusion-weighted images. *NMR in Biomedicine*, 8:333–44, 1995.
- P. J. Basser. Relationships between diffusion tensor and q-space MRI. *Magnetic Resonance in Medicine*, 47(2):392–7, 2002.
- P. J. Basser and D. K. Jones. Diffusion-tensor MRI: theory, experimental design and data analysis - a technical review. *NMR in Biomedicine*, 15(7-8):456–67, 2002.

- P. J. Basser, J. Mattiello, and D. LeBihan. Estimation of the effective self-diffusion tensor from the NMR spin echo. *Journal of Magnetic Resonance Series b*, 103:247–54, 1994a.
- P. J. Basser, J. Mattiello, and D. LeBihan. MR diffusion tensor spectroscopy and imaging. *Biophysical Journal*, 66(1):259–67, 1994b.
- P. J. Basser, S. Pajevic, C. Pierpaoli, J. Duda, and A. Aldroubi. In vivo fiber tractography using DT-MRI data. *Magnetic Resonance in Medicine*, 44(4):625–32, 2000.
- P. G. Batchelor, D. L. G. Hill, F. Calamante, and D. Atkinson. Study of connectivity in the brain using the full diffusion tensor from MRI. In *Information Processing in Medical Imaging (IPMI)*, pages 121–133, 2001.
- C. Beaulieu. The basis of anisotropic water diffusion in the nervous system - a technical review. *NMR in Biomedicine*, 15(7-8):435–55, 2002.
- T. E. Behrens, M. W. Woolrich, M. Jenkinson, H. Johansen-Berg, R. G. Nunes, S. Clare, P. M. Matthews, J. M. Brady, and S. M. Smith. Characterization and propagation of uncertainty in diffusion-weighted MR imaging. *Magnetic Resonance in Medicine*, 50(5):1077–88, 2003.
- T. E. Behrens, H. J. Berg, S. Jbabdi, M. F. Rushworth, and M. W. Woolrich. Probabilistic diffusion tractography with multiple fibre orientations: What can we gain? *Neuroimage*, 34(1):144–55, 2007.
- T. E. J. Behrens and S. Jbabdi. Mr diffusion tractography. In H. Johansen-Berg and T.E. Behrens, editors, *Diffusion MRI: From quantitative measurement to in vivo neuroanatomy*, pages 333–351. Elsevier, 2009.
- J. I. Berman, S. Chung, P. Mukherjee, C. P. Hess, E. T. Han, and R. G. Henry. Probabilistic streamline Q-ball tractography using the residual bootstrap. *Neuroimage*, 39(1):215–222, Jan 2008.
- N. C. Bodammer, J. Kaufmann, M. Kanowski, and C. Tempelmann. Monte Carlo-based diffusion tensor tractography with a geometrically corrected voxel-centre connecting method. *Physics in Medicine and Biology*, 54(4):1009–1033, 2009.
- K. Brodmann. *Vergleichende localisationslehre der grosshirnrinde in ihren principien dargestellt auf grund des zellenbaues* [Translated by L. J. Garey, *Brodmann's localisation in the Cerebral Cortex*, 1994]. Barth, 1909.
- G. Buzsaki. *Rhythms of the brain*. Oxford University Press, 2006.
- P. T. Callaghan. The measurement of motion using spin echoes. In *Principles of nuclear magnetic resonance microscopy*, pages 328–370. Oxford University Press, 1991.
- J. S. Campbell, K. Siddiqi, V. V. Rymer, A. F. Sadikot, and G. B. Pike. Flow-based fiber tracking with diffusion tensor and Q-ball data: validation and comparison to principal diffusion direction techniques. *Neuroimage*, 27(4):725–36, 2005.
- B. M. Carvalho, C. J. Gau, G. T. Herman, and T. Y. Kong. Algorithms for fuzzy segmentation. *Pattern Analysis & Applications*, 2:73–81, 1999.
- M. Catani, R. J. Howard, S. Pajevic, and D. K. Jones. Virtual in vivo interactive dissection of white matter fasciculi in the human brain. *Neuroimage*, 17(1):77–94, 2002.

- Y. P. Chao, J. H. Chen, K. H. Cho, C. H. Yeh, K. H. Chou, and C. P. Lin. A multiple streamline approach to high angular resolution diffusion tractography. *Medical Engineering and Physics*, 30(8):989–996, 2008.
- K. H. Cho, C. H. Yeh, J. D. Tournier, Y. P. Chao, J. H. Chen, and C. P. Lin. Evaluation of the accuracy and angular resolution of Q-ball imaging. *Neuroimage*, 42(1):262–271, 2008.
- O. Ciccarelli, M. Catani, H. Johansen-Berg, C. Clark, and A. Thompson. Diffusion-based tractography in neurological disorders: concepts, applications, and future developments. *Lancet Neurology*, 7(8):715–727, 2008.
- C. A. Clark and D. Le Bihan. Water diffusion compartmentation and anisotropy at high b values in the human brain. *Magnetic Resonance in Medicine*, 44(6):852–859, Dec 2000.
- C. A. Clark, T. R. Barrick, M. M. Murphy, and B. Anthony Bell. White matter fiber tracking in patients with space-occupying lesions of the brain: a new technique for neurosurgical planning? *Neuroimage*, 20(3):1601–1608, 2003.
- J. D. Clayden and C. A. Clark. Model-based streamline rejection for probabilistic tractography. In *Proceedings of the International Society for Magnetic Resonance in Medicine (ISMRM) Meeting*, page 1440, Honolulu, USA, 2009.
- T. E. Conturo, N. F. Lori, T. S. Cull, E. Akbudak, A. Z. Snyder, J. S. Shimony, R. C. McKinstry, H. Burton, and M. E. Raichle. Tracking neuronal fiber pathways in the living human brain. *Proceedings of the National Academy of Sciences*, 96(18):10422–7, 1999.
- P. A. Cook, M. Symms, P. A. Boulby, and D. C. Alexander. Optimal acquisition orders of diffusion-weighted MRI measurements. *Journal of Magnetic Resonance Imaging*, 25(5):1051–8, 2007.
- P.A. Cook, D.C. Alexander, and G.J.M. Parker. Modelling noise-induced fibre-orientation error in diffusion-tensor MRI. In *IEEE International Symposium on Biomedical Imaging (ISBI)*, pages 332–335, April 2004.
- P.A. Cook, Y. Bai, S. Nedjati-Gilani, K. K. Seunarine, M. G. Hall, G. J. Parker, and D. C. Alexander. Camino: Open-source diffusion-MRI reconstruction and processing. In *Proceedings of the International Society for Magnetic Resonance in Medicine (ISMRM) Meeting*, page 2759, Seattle, USA, 2006.
- T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein. *Introduction to algorithms*. MIT press, 2002.
- O. Coulon, D. C. Alexander, and S. Arridge. Diffusion tensor magnetic resonance image regularization. *Medical Image Analysis*, 8(1):47–67, 2004.
- J. Dauguet, S. Peled, V. Berezovskii, T. Delzescaux, S. K. Warfield, R. Born, and C.-F. Westin. Comparison of fiber tracts derived from in-vivo DTI tractography with 3D histological neural tract tracer reconstruction on a macaque brain. *Neuroimage*, 37(2): 530–538, 2007.
- J. Dejerine and M. Dejerine-Klumpke. *Anatomie des centres nerveux. 1st volume*. Rueff & Co, 1895.

- F. Dell'Acqua, G. Rizzo, P. Scifo, R. A. Clarke, G. Scotti, and F. Fazio. A model-based deconvolution approach to solve fiber crossing in diffusion-weighted MR imaging. *IEEE Trans Biomed Eng*, 54(3):462–472, Mar 2007.
- M. Descoteaux. *High angular resolution diffusion MRI: from local estimation to segmentation and tractography*. PhD thesis, University of Nice-Sophia Antipolis, 2008.
- M. Descoteaux, E. Angelino, S. Fitzgibbons, and R. Deriche. Apparent diffusion coefficients from high angular resolution diffusion imaging: estimation and applications. *Magnetic Resonance in Medicine*, 56(2):395–410, Aug 2006.
- M. Descoteaux, E. Angelino, S. Fitzgibbons, and R. Deriche. Regularized, fast, and robust analytical Q-ball imaging. *Magnetic Resonance in Medicine*, 58(3):497–510, 2007.
- M. Descoteaux, R. Deriche, T. R. Knosche, and A. Anwender. Deterministic and probabilistic tractography based on complex fibre orientation distributions. *IEEE Transactions on Medical Imaging*, 28(2):269–86, 2009.
- E. W. Dijkstra. A note on two problems in connexion with graphs. *Numerische Mathematik*, 1:269–271, 1959.
- Z. Ding, J. C. Gore, and A. W. Anderson. Reduction of noise in diffusion tensor images using anisotropic smoothing. *Magnetic Resonance in Medicine*, 53:485–490, 2005.
- B. Draganski, F. Kherif, S. Klöppel, P. A. Cook, D. C. Alexander, G. J. M. Parker, R. Deichmann, J. Ashburner, and R. S. J. Frackowiak. Evidence for segregated and integrative connectivity patterns in the human basal ganglia. *Journal of Neuroscience*, 28(28):7143–7152, Jul 2008.
- T. B. Dyrby, L. V. Sogaard, G. J. Parker, D. C. Alexander, N. M. Lind, W. F. C. Baare, A. Hay-Schmidt, N. Eriksen, B. Pakkenberg, O. B. Paulson, and J. Jelsing. Validation of in vitro probabilistic tractography. *Neuroimage*, 37(4):1267–1277, 2007.
- P. Fillard, C. Poupon, and J.-F. Mangin. A novel global tractography algorithm based on an adaptive spin glass model. In *Proceedings of the Medical Image Computing and Computer Assisted Intervention (MICCAI) Conference*, pages 927–934, London, UK, 2009.
- P. T. Fletcher, R. Tao, W. K. Jeong, and R. T. Whitaker. A volumetric approach to quantifying region-to-region white matter connectivity in diffusion tensor MRI. *Information Processing in Medical Imaging (IPMI)*, 20:346–58, 2007.
- H. M. J. Fonteijn, F. A. J. Verstraten, and D. G. Norris. Probabilistic inference on Q-ball imaging data. *IEEE Transactions on Medical Imaging*, 26:1515–1524, 2007.
- L. R. Frank. Anisotropy in high angular resolution diffusion-weighted MRI. *Magnetic Resonance in Medicine*, 45(6):935–9, 2001.
- L. R. Frank. Characterization of anisotropy in high angular resolution diffusion-weighted MRI. *Magnetic Resonance in Medicine*, 47(6):1083–99, 2002.
- O. Friman, G. Farneback, and C. F. Westin. A bayesian approach for stochastic white matter tractography. *IEEE Transactions on Medical Imaging*, 25(8):965–78, 2006.

- K. J. Friston, C. D. Frith, P. F. Liddle, and R. S. Frackowiak. Functional connectivity: the principal-component analysis of large (PET) data sets. *Journal of Cerebral Blood Flow and Metabolism*, 13(1):5–14, Jan 1993.
- G. Gerig, O. Kubler, R. Kikinis, and F. A. Jolesz. Nonlinear anisotropic filtering of MRI data. *IEEE Transactions on Medical Imaging*, 11(2):221–232, 1992.
- N. Geschwind. Disconnexion syndromes in animals and man. I. *Brain*, 88(2):237–294, 1965a.
- N. Geschwind. Disconnexion syndromes in animals and man. II. *Brain*, 88(3):585–644, 1965b.
- H. Gray. *Anatomy of the human body*. Lea & Febiger, 1918.
- M. D. Greicius, B. Krasnow, A. L. Reiss, and V. Menon. Functional connectivity in the resting brain: a network analysis of the default mode hypothesis. *Proceedings of the National Academy of Sciences*, 100(1):253–258, 2003.
- H. Gudbjartsson and S. Patz. The Rician distribution of noisy MRI data. *Magnetic Resonance in Medicine*, 34(6):910–4, 1995.
- E. M. Haacke, R. W. Brown, M. R. Thompson, and R. Venkatesan. *Magnetic Resonance Imaging: Physical principles and sequence design*. John Wiley & Sons, 1999.
- N. S. Hageman, A. W. Toga, K. L. Narr, and D. W. Shattuck. A diffusion tensor imaging tractography algorithm based on Navier-Stokes fluid mechanics. *IEEE Transactions on Medical Imaging*, 28(3):348–60, 2009.
- P. Hagmann, J. P. Thiran, L. Jonasson, P. Vandergheynst, S. Clarke, P. Maeder, and R. Meuli. DTI mapping of human brain connectivity: statistical fibre tracking and virtual dissection. *Neuroimage*, 19(3):545–54, 2003.
- P. Hagmann, L. Cammoun, X. Gigandet, R. Meuli, C. J. Honey, V. J. Wedeen, and O. Sporns. Mapping the structural core of human cerebral cortex. *PLoS Biology*, 6(7):1479–1473, 2008.
- D. E. Haines. *Neuroanatomy. An Atlas of Structures, Sections, and Systems*. Lippincott Williams & Wilkins, 2004.
- H. A. Haroon, D. M. Morris, K. V. Embleton, D. C. Alexander, and G. J. Parker. Using the model-based residual bootstrap to quantify uncertainty in fiber orientations from Q-ball analysis. *IEEE Transactions on Medical Imaging*, 28(4):535–50, 2009.
- K. M. Hasan, P. J. Basser, D. L. Parker, and A. L. Alexander. Analytical computation of the eigenvalues and eigenvectors in DT-MRI. *Journal of Magnetic Resonance*, 152(1):41–7, 2001.
- D. M. Healy, H. Hendriks, and P. T. Kim. Spherical deconvolution. *Journal of Multivariate Analysis*, 67:1–22, 1998.
- C. P. Hess, P. Mukherjee, E. T. Han, D. Xu, and D. B. Vigneron. Q-ball reconstruction of multimodal fiber orientations using the spherical harmonic basis. *Magnetic Resonance in Medicine*, 56(1):104–17, 2006.

- R. K. Hobbie. *Intermediate Physics for Medicine and Biology*. Springer-Verlag, New York, 1997.
- T. Hosey, G. Williams, and R. Ansorge. Inference of multiple fiber orientations in high angular resolution diffusion imaging. *Magnetic Resonance in Medicine*, 54(6):1480–9, 2005.
- H. Huang, J. Zhang, P. C. M. van Zijl, and S. Mori. Analysis of noise effects on DTI-based tractography using the brute-force and multi-ROI approach. *Magnetic Resonance in Medicine*, 52(3):559–565, 2004.
- Y. Iturria-Medina, E. J. Canales-Rodriguez, L. Melie-Garcia, P. A. Valdes-Hernandez, E. Martinez-Montes, Y. Aleman-Gomez, and J. M. Sanchez-Bornot. Characterizing brain anatomical connections using diffusion weighted MRI and graph theory. *Neuroimage*, 36:645–60, 2007.
- Y. Iturria-Medina, R. C. Sotero, E. J. Canales-Rodriguez, Y. Aleman-Gomez, and L. Melie-Garcia. Studying the human brain anatomical network via diffusion-weighted MRI and graph theory. *Neuroimage*, 40:1064–76, 2008.
- M. Jackowski, C. Y. Kao, M. Qiu, R. T. Constable, and L. H. Staib. White matter tractography by anisotropic wavefront evolution and diffusion tensor imaging. *Medical Image Analysis*, 9:427–40, 2005.
- K. M. Jansons and D. C. Alexander. Persistent angular structure: new insights from diffusion magnetic resonance imaging data. *Inverse Problems*, 19:1031–1046, 2003.
- S. Jbabdi, M. W. Woolrich, J. L. Andersson, and T. E. Behrens. A Bayesian framework for global tractography. *Neuroimage*, 37:116–29, 2007.
- J. H. Jensen, J. A. Helpert, A. Ramani, H. Lu, and K. Kaczynski. Diffusional kurtosis imaging: the quantification of non-gaussian water diffusion by means of magnetic resonance imaging. *Magnetic Resonance in Medicine*, 53:1432–40, 2005.
- B. Jian, B. C. Vemuri, E. Ozarslan, P. R. Carney, and T. H. Mareci. A novel tensor distribution model for the diffusion-weighted MR signal. *Neuroimage*, 37(1):164–176, 2007.
- H. Johansen-Berg, T. E. Behrens, E. Sillery, O. Ciccarelli, A. J. Thompson, S. M. Smith, and P. M. Matthews. Functional-anatomical validation and individual variation of diffusion tractography-based segmentation of the human thalamus. *Cerebral Cortex*, 15:31–9, 2005.
- D. K. Jones. Determining and visualizing uncertainty in estimates of fiber orientation from diffusion tensor MRI. *Magnetic Resonance in Medicine*, 49:7–12, 2003.
- D. K. Jones. The effect of gradient sampling schemes on measures derived from diffusion tensor MRI: a Monte Carlo study. *Magnetic Resonance in Medicine*, 51:807–15, 2004.
- D. K. Jones. Tractography gone wild: probabilistic fibre tracking using the wild bootstrap with diffusion tensor MRI. *IEEE Trans Med Imaging*, 27:1268–74, 2008.

- D. K. Jones. Gaussian modeling of the diffusion signal. In H. Johansen-Berg and T.E. Behrens, editors, *Diffusion MRI: From quantitative measurement to in vivo neuroanatomy*, pages 37–54. Elsevier, 2009.
- D. K. Jones and C. Pierpaoli. Confidence mapping in diffusion tensor magnetic resonance imaging tractography using a bootstrap approach. *Magnetic Resonance in Medicine*, 53:1143–9, 2005.
- D. K. Jones, M. A. Horsfield, and A. Simmons. Optimal strategies for measuring diffusion in anisotropic systems by magnetic resonance imaging. *Magnetic Resonance in Medicine*, 42:515–25, 1999.
- E. Kaden, T. R. Knösche, and A. Anwender. Parametric spherical deconvolution: inferring anatomical connectivity using diffusion MR imaging. *Neuroimage*, 37(2):474–488, 2007.
- E. R. Kandel, J. H. Schwartz, and T. M. Jessel. *Principles of Neural Science*. McGraw-Hill, 2000.
- N. Kang, J. Zhang, E. S. Carlson, and D. Gembris. White matter fiber tractography via anisotropic diffusion simulation in the human brain. *IEEE Transactions on Medical Imaging*, 24:1127–37, 2005.
- B. W. Kernighan and D. M. Ritchie. *The C programming language*. Prentice Hall, 1988.
- P. B. Kingsley. Introduction to diffusion tensor imaging mathematics: Part I. Tensors, rotations, and eigenvectors. *Concepts in Magnetic Resonance Part A*, 28A:101–122, 2006a.
- P. B. Kingsley. Introduction to diffusion tensor imaging mathematics: Part III. Tensor calculation, noise, simulations, and optimization. *Concepts in Magnetic Resonance Part A*, 28A:155–179, 2006b.
- C. Koch and G. Laurent. Complexity and the nervous system. *Science*, 284:96–98, 1999.
- M. A. Koch, D. G. Norris, and M. Hund-Georgiadis. An investigation of functional and anatomical connectivity using magnetic resonance imaging. *Neuroimage*, 16(1): 241–250, 2002.
- B. W. Kreher, J. F. Schneider, I. Mader, E. Martin, J. Hennig, and K. A. Il’yasov. Multi-tensor approach for analysis and tracking of complex fiber configurations. *Magnetic Resonance in Medicine*, 54:1216–25, 2005.
- B. W. Kreher, I. Mader, and V. G. Kiselev. Gibbs tracking: a novel approach for the reconstruction of neuronal pathways. *Magnetic Resonance in Medicine*, 60(4):953–963, 2008a.
- B. W. Kreher, S. Schnell, I. Mader, K. A. Il’yasov, J. Hennig, V. G. Kiselev, and D. Saur. Connecting and merging fibres: pathway extraction by combining probability maps. *Neuroimage*, 43(1):81–89, 2008b.
- R. L. Kruse and A. J. Ryba. *Data Structures and Program Design in C++*. Prentice Hall, 2000.

- P. C. Lauterbur. Image formation by induced local interactions: Examples employing nuclear magnetic resonance. *Nature*, (242):190–191, 1973.
- J. H. LaVail and M. M. LaVail. Retrograde axonal transport in the central nervous system. *Science*, 176(42):1416–1417, 1972.
- I. N. C. Lawes, T. R. Barrick, V. Murugam, N. Spierings, D. R. Evans, M. Song, and C. A. Clark. Atlas-based segmentation of white matter tracts of the human brain using diffusion tensor tractography and comparison with classical dissection. *Neuroimage*, 39(1):62–79, 2008.
- M. Lazar and A. L. Alexander. An error analysis of white matter tractography methods: synthetic diffusion tensor field simulations. *Neuroimage*, 20(2):1140–1153, 2003.
- M. Lazar and A. L. Alexander. Bootstrap white matter tractography (BOOT-TRAC). *Neuroimage*, 24:524–32, 2005.
- M. Lazar, D. M. Weinstein, J. S. Tsuruda, K. M. Hasan, K. Arfanakis, M. E. Meyerand, B. Badie, H. A. Rowley, V. Haughton, A. Field, and A. L. Alexander. White matter tractography using diffusion tensor deflection. *Human Brain Mapping*, 18:306–21, 2003.
- M. Lazar, J. H. Jensen, L. Xuan, and J. A. Helpert. Estimation of the orientation distribution function from diffusional kurtosis imaging. *Magnetic Resonance in Medicine*, 60(4):774–781, Oct 2008.
- D. Le Bihan, E. Breton, D. Lallemand, P. Grenier, E. Cabanis, and M. Laval-Jeantet. MR imaging of intravoxel incoherent motions: application to diffusion and perfusion in neurologic disorders. *Radiology*, 161(2):401–407, Nov 1986.
- A. D. Leow, S. Zhu, L. Zhan, K. McMahon, G. I. de Zubicaray, M. Meredith, M. J. Wright, A. W. Toga, and P. M. Thompson. The tensor distribution function. *Magnetic Resonance in Medicine*, 61(1):205–214, 2009.
- S. Lifshits, A. Tamir, and Y. Assaf. Combinatorial fiber-tracking of the human brain. *Neuroimage*, 48(3):532–540, 2009.
- H. Lu, J. H. Jensen, A. Ramani, and J. A. Helpert. Three-dimensional characterization of non-gaussian water diffusion in humans using diffusion kurtosis imaging. *NMR in Biomedicine*, 19:236–47, 2006.
- E. Ludwig and J. Klingler. *Atlas cerebri humani: The Inner Structure of the Brain Demonstrated on the Basis of Macroscopical Preparations*. Little, Brown & Company, 1956.
- P. Mansfield and A. A. Maudsley. Medical imaging by NMR. *British Journal of Radiology*, 50(591):188–194, 1977.
- P. Mitra and B. Halperin. Effects of finite gradient-pulse widths in pulsed-field-gradient diffusion measurements. *Journal of Magnetic Resonance, Series A*, 113:94–101, 1995.
- S. Mori, B. J. Crain, V. P. Chacko, and P. C. van Zijl. Three-dimensional tracking of axonal projections in the brain by magnetic resonance imaging. *Annals of Neurology*, 45:265–9, 1999.

- S. Mori, S. Wakana, P. C. Van Zijl, and L.M. Nagae-Poetscher. *MRI Atlas of Human White Matter*. Elsevier, 2005.
- D. M. Morris, K. V. Embleton, and G. J. Parker. Probabilistic fibre tracking: differentiation of connections from chance events. *Neuroimage*, 42:1329–39, 2008.
- P. Mukherjee, C. P. Hess, D. Xu, E. T. Han, D. A. Kelley, and D. B. Vigneron. Development and initial evaluation of 7-T Q-ball imaging of the human brain. *Magnetic Resonance Imaging*, 26(2):171–180, 2008.
- S. Nedjati-Gilani and D. C. Alexander. Models for fanning and bending sub-voxel structures in diffusion MRI. In *Proceedings of the DMFC MICCAI 2009 Workshop*, London, UK, 2009.
- L. G. Nyul, A. X. Falcao, and J. K. Udupa. Fuzzy-connected 3D image segmentation at interactive speeds. *Graphical Models*, 64:259–281, 2003.
- R. L. O’Gorman and D. K. Jones. Just how much data need to be collected for reliable bootstrap DT-MRI? *Magnetic Resonance in Medicine*, 56(4):884–890, 2006.
- E. Ozarslan and T. H. Mareci. Generalized diffusion tensor imaging and analytical relationships between diffusion tensor imaging and high angular resolution diffusion imaging. *Magnetic Resonance in Medicine*, 50:955–65, 2003.
- E. Ozarslan, T. M. Shepherd, B. C. Vemuri, S. J. Blackband, and T. H. Mareci. Resolution of complex tissue microarchitecture using the diffusion orientation transform (DOT). *Neuroimage*, 31:1086–103, 2006.
- S. Pajevic and P. J. Basser. Parametric and non-parametric statistical analysis of DT-MRI data. *Journal of Magnetic Resonance*, 161:1–14, 2003.
- S. Pajevic and C. Pierpaoli. Color schemes to represent the orientation of anisotropic tissues from diffusion tensor data: application to white matter fiber tract mapping in the human brain. *Magnetic Resonance in Medicine*, 42(3):526–540, 1999.
- G. J. Parker and D. C. Alexander. Probabilistic Monte Carlo based mapping of cerebral connections utilising whole-brain crossing fibre information. *Information Processing in Medical Imaging (IPMI)*, 18:684–95, 2003.
- G. J. Parker and D. C. Alexander. Probabilistic anatomical connectivity derived from the microscopic persistent angular structure of cerebral tissue. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 360:893–902, 2005.
- G. J. Parker, J. A. Schnabel, M. R. Symms, D. J. Werring, and G. J. Barker. Nonlinear smoothing for reduction of systematic and random errors in diffusion tensor imaging. *Journal of Magnetic Resonance Imaging*, 11:702–710, 2000.
- G. J. Parker, C. A. Wheeler-Kingshott, and G. J. Barker. Estimating distributed anatomical connectivity using fast marching methods and diffusion tensor imaging. *IEEE Transactions on Medical Imaging*, 21:505–12, 2002.
- G. J. Parker, H. A. Haroon, and C. A. Wheeler-Kingshott. A framework for a streamline-based probabilistic index of connectivity (PICO) using a structural interpretation of MRI diffusion measurements. *Journal of Magnetic Resonance Imaging*, 18:242–54, 2003.

- B. Patenaude, S. Smith, D. Kennedy, and M. Jenkinson. FIRST - FMRIB's integrated registration and segmentation tool. In *Proceedings of the Organization for Human Brain Mapping (OHBM) Meeting*, page 857, Chicago, USA, 2007.
- S. Peled, O. Friman, F. Jolesz, and C. F. Westin. Geometrically constrained two-tensor model for crossing tracts in DWI. *Magnetic Resonance Imaging*, 24:1263–70, 2006.
- P. Perona and J. Malik. Scale-space and edge detection using anisotropic diffusion. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12(7):629–639, 1990.
- C. Pierpaoli and P. J. Basser. Toward a quantitative assessment of diffusion anisotropy. *Magnetic Resonance in Medicine*, 36:893–906, 1996.
- C. Pierpaoli, P. Jezzard, P. J. Basser, A. Barnett, and G. Di Chiro. Diffusion tensor MR imaging of the human brain. *Radiology*, 201:637–48, 1996.
- D. Pollard. *Fourier Transforms*, pages 202–203. Cambridge University Press, 2001.
- C. Poupon, C. A. Clark, V. Frouin, J. Regis, I. Bloch, D. Le Bihan, and J. Mangin. Regularization of diffusion-based direction maps for the tracking of brain white matter fascicles. *Neuroimage*, 12:184–95, 2000.
- C. Poupon, B. Rieul, I. Kezele, M. Perrin, F. Poupon, and J. F. Mangin. New diffusion phantoms dedicated to the study and validation of high-angular-resolution diffusion imaging (HARDI) models. *Magnetic Resonance in Medicine*, 60(6):1276–1283, 2008.
- W. Press, B. Flannery, S. Teukolsky, and W. Vetterling. *Numerical Recipes in C : The Art of Scientific Computing*. Cambridge University Press, 2nd edition, 1992.
- A. Ramirez-Manzanares and M. Rivera. Basis tensor decomposition for restoring intra-voxel structure and stochastic walks for inferring brain connectivity in DT-MRI. *International Journal of Computer Vision*, 69(1):77– 92, 2006.
- A. Rosenfeld, R. A. Hummel, and S. W. Zucker. Scene labeling by relaxation operations. *IEEE Transactions on Systems, Man and Cybernetics*, 6:420–433, 1976.
- P. K. Saha, J. K. Udupa, and D. Odhner. Scale-based fuzzy connected image segmentation: Theory, algorithms, and validation. *Computer Vision and Image Understanding*, 77:145–174, 2000.
- K. E. Sakaie. How fast can PAS go? In *Proceedings of the International Society for Magnetic Resonance in Medicine (ISMRM) Meeting*, page 3323, Toronto, Canada, 2008.
- R. Salvador, A. Pena, D. K. Menon, T. A. Carpenter, J. D. Pickard, and E. T. Bullmore. Formal characterization and extension of the linearized diffusion tensor model. *Human Brain Mapping*, 24(2):144–155, 2005.
- P. Savadjiev, J. S. Campbell, G. B. Pike, and K. Siddiqi. 3D curve inference for diffusion MRI regularization and fibre tractography. *Medical Image Analysis*, 10:799–813, 2006.
- P. Savadjiev, J. S. W. Campbell, M. Descoteaux, R. Deriche, G. B. Pike, and K. Siddiqi. Labeling of ambiguous subvoxel fibre bundle configurations in high angular resolution diffusion MRI. *Neuroimage*, 41(1):58–68, 2008.

- T. Schultz and H. P. Seidel. Estimating crossing fibers: a tensor decomposition approach. *IEEE Transactions on Visualization and Computer Graphics*, 14:1635–42, 2008.
- J. A. Sethian. A fast marching level set method for monotonically advancing fronts. *Proceedings of the National Academy of Sciences*, 93:1591–5, 1996.
- K. K. Seunarine and D. C. Alexander. Multiple fibers: Beyond the diffusion tensor. In H. Johansen-Berg and T.E. Behrens, editors, *Diffusion MRI: From quantitative measurement to in vivo neuroanatomy*, pages 55–72. Elsevier, 2009.
- K. K. Seunarine, P. A. Cook, M. G. Hall, K. V. Embleton, G. J. M. Parker, and D. C. Alexander. Exploiting peak anisotropy for tracking through complex structures. In *Proceedings of the IEEE International Conference on Computer Vision Workshop on MM-BIA*, 2007.
- S. Skare, T. Li, B. Nordell, and M. Ingvar. Noise considerations in the determination of diffusion tensor anisotropy. *Magnetic Resonance Imaging*, 18(6):659–669, 2000.
- S. M. Smith, M. Jenkinson, M. W. Woolrich, C. F. Beckmann, T. E. Behrens, H. Johansen-Berg, P. R. Bannister, M. De Luca, I. Drobnjak, D. E. Flitney, R. K. Niazy, J. Saunders, J. Vickers, Y. Zhang, N. De Stefano, J. M. Brady, and P. M. Matthews. Advances in functional and structural MR image analysis and implementation as FSL. *Neuroimage*, 23 Suppl 1:S208–19, 2004.
- S. N. Sotiropoulos, C. R. Tench, P. S. Morgan, and L. Bai. Robust graph-based tracking through crossing fibre configurations. In *Proceedings of the IEEE International Symposium on Biomedical Imaging (ISBI)*, pages 1394–1397, Boston, USA, 2009.
- S. N. Sotiropoulos, D. E. Jones, L. Bai, and T. Kypraios. Exact and analytic Bayesian inference for orientation distribution functions. In *Proceedings of the IEEE International Symposium on Biomedical Imaging (ISBI)*, 2010.
- O. Sporns. Graph theory methods for the analysis of neuronal connectivity patterns. In R. Kotter, editor, *Neuroscience databases: a practical guide*, pages 171–185. Kluwer Academic Publishers, 2002.
- O. Sporns, G. Tononi, and R. Kötter. The human connectome: A structural description of the human brain. *PLoS Computational Biology*, 1(4):245–251, 2005.
- P. Staempfli, T. Jaermann, G. R. Crelier, S. Kollias, A. Valavanis, and P. Boesiger. Resolving fiber crossing using advanced fast marching tractography based on diffusion tensor imaging. *Neuroimage*, 30:110–20, 2006.
- E. O. Stejskal and J. E. Tanner. Spin diffusion measurements: Spin echoes in the presence of a time-dependent field gradient. *Journal of Chemical Physics*, 42(1):288–292, 1965.
- C. R. Tench. *Diffusion tensor MRI tractography and its application to multiple sclerosis*. PhD thesis, University of Nottingham, 2003.
- C. R. Tench, P. S. Morgan, L. D. Blumhardt, and C. Constantinescu. Improved white matter fiber tracking using stochastic labeling. *Magnetic Resonance in Medicine*, 48: 677–83, 2002.

- J. D. Tournier, F. Calamante, D. G. Gadian, and A. Connelly. Diffusion-weighted magnetic resonance imaging fibre tracking using a front evolution algorithm. *Neuroimage*, 20:276–88, 2003.
- J. D. Tournier, F. Calamante, D. G. Gadian, and A. Connelly. Direct estimation of the fiber orientation density function from diffusion-weighted MRI data using spherical deconvolution. *Neuroimage*, 23:1176–85, 2004.
- J. D. Tournier, F. Calamante, and A. Connelly. Robust determination of the fibre orientation distribution in diffusion MRI: non-negativity constrained super-resolved spherical deconvolution. *Neuroimage*, 35:1459–72, 2007.
- D. Tschumperle and R. Deriche. Orthonormal vector sets regularization with PDEs and applications. *International Journal of Computer Vision*, 50:237–252, 2002.
- D. S. Tuch. Q-ball imaging. *Magnetic Resonance in Medicine*, 52:1358–72, 2004.
- D. S. Tuch, T. G. Reese, M. R. Wiegell, N. Makris, J. W. Belliveau, and V. J. Wedeen. High angular resolution diffusion imaging reveals intravoxel white matter fiber heterogeneity. *Magnetic Resonance in Medicine*, 48:577–82, 2002.
- R. Turner, D. Le Bihan, J. Maier, R. Vavrek, L. K. Hedges, and J. Pekar. Echo-planar imaging of intravoxel incoherent motion. *Radiology*, 177:407–14, 1990.
- J. K. Udupa and P. K. Saha. Fuzzy connectedness and image segmentation. *Proceedings of the IEEE*, 91(10):1649–1669, 2003.
- J. K. Udupa and S. Samarasekera. Fuzzy connectedness and object definition: Theory, algorithms, and applications in image segmentation. *Graphical Models and Image Processing*, 58:246–261, 1996.
- M. P. van den Heuvel, R. C. Mandl, R. S. Kahn, and H. E. Hulshoff Pol. Functionally linked resting-state networks reflect the underlying structural connectivity architecture of the human brain. *Human Brain Mapping*, 30:3127 – 3141, 2009.
- V. J. Wedeen, P. Hagmann, W. Y. Tseng, T. G. Reese, and R. M. Weisskoff. Mapping complex tissue architecture with diffusion spectrum magnetic resonance imaging. *Magnetic Resonance in Medicine*, 54:1377–86, 2005.
- V. J. Wedeen, R. P. Wang, J. D. Schmahmann, T. Benner, W. Y. I. Tseng, G. Dai, D. N. Pandya, P. Hagmann, H. D’Arceuil, and A. J. de Crespigny. Diffusion spectrum magnetic resonance imaging (DSI) tractography of crossing fibers. *Neuroimage*, 41(4):1267–1277, 2008.
- C.-F. Westin, S. E. Maier, H. Mamata, A. Nabavi, F. A. Jolesz, and R. Kikinis. Processing and visualization for diffusion tensor MRI. *Medical Image Analysis*, 6(2):93–108, 2002.
- B. Whitcher, D. S. Tuch, J. J. Wisco, A. G. Sorensen, and L. Wang. Using the wild bootstrap to quantify uncertainty in diffusion tensor imaging. *Human Brain Mapping*, 29(3):346–362, 2008.
- M. R. Wiegell, H. B. Larsson, and V. J. Wedeen. Fiber crossing in human brain depicted with diffusion tensor MR imaging. *Radiology*, 217:897–903, 2000.

- T. H. Williams, N. Gluhbegovic, and J. Y. Jew. *The Human Brain: Dissections of the Real Brain*. University of Iowa, 1997.
- T. A. Woolsey, J. Hanaway, and M. H. Gado. *The Brain Atlas. A visual guide to the human central nervous system*. Wiley, 2008.
- Y. C. Wu and A. L. Alexander. Hybrid diffusion imaging. *Neuroimage*, 36(3):617–629, 2007.
- A. Zalesky. DT-MRI fiber tracking: a shortest paths approach. *IEEE Transactions on Medical Imaging*, 27:1458–71, 2008.
- W. Zhan and Y. Yang. How accurately can the diffusion profiles indicate multiple fiber orientations? a study on general fiber crossings in diffusion MRI. *Journal of Magnetic Resonance*, 183:193–202, 2006.
- F. Zhang, E. R. Hancock, C. Goodlett, and G. Gerig. Probabilistic white matter fiber tracking using particle filtering and von Mises-Fisher sampling. *Medical Image Analysis*, 13:5–18, 2009.
- D. Zhou, I. L. Dryden, A. Koloydenko, and L. Bai. A Bayesian method with reparameterization for diffusion tensor imaging. In *Proceedings of the SPIE Medical Imaging Conference*, volume 6914, pages 2J–2J11, San Diego, USA, 2008.
- D. Zwillinger. *CRC Standard Mathematical Tables and Formulae*. Chapman & Hall/CRC Press, 31st edition, 2003.