

DIFFUSION TENSOR MAGNETIC RESONANCE IMAGING
AND FIBRE TRACTOGRAPHY IN THE BRAIN

Jacques-Donald Tournier

Radiology and Physics Unit
Institute of Child Health
University College London

Submitted for the Degree of Doctor of Philosophy
University of London

April 2003

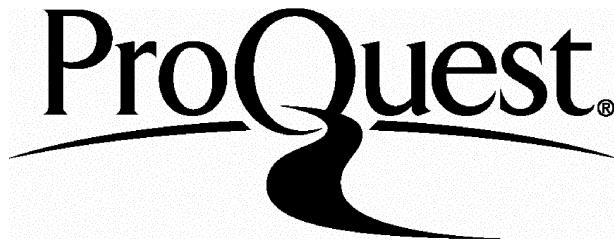
ProQuest Number: U642487

All rights reserved

INFORMATION TO ALL USERS

The quality of this reproduction is dependent upon the quality of the copy submitted.

In the unlikely event that the author did not send a complete manuscript and there are missing pages, these will be noted. Also, if material had to be removed, a note will indicate the deletion.



ProQuest U642487

Published by ProQuest LLC(2015). Copyright of the Dissertation is held by the Author.

All rights reserved.

This work is protected against unauthorized copying under Title 17, United States Code.
Microform Edition © ProQuest LLC.

ProQuest LLC
789 East Eisenhower Parkway
P.O. Box 1346
Ann Arbor, MI 48106-1346

ABSTRACT

Diffusion tensor magnetic resonance imaging allows detailed studies of the diffusion of water to be carried out non-invasively *in-vivo*. Brain white matter consists of coherently oriented fibre bundles, in which the diffusion coefficient is larger when measured along the direction of fibres rather than across them. Fibre-tracking techniques attempt to use this information to establish connectivity within the brain, by assuming that the direction of fastest diffusion corresponds to the orientation of the underlying fibre tract.

Fibre-tracking requires the acquisition of a large number of images with the diffusion weighting applied in different directions. A number of technical difficulties had to be overcome before such data could be acquired on our Siemens Vision 1.5T system. In particular, eddy-current problems needed to be addressed, and all image reconstruction was implemented offline to allow a more rapid acquisition of the raw data.

Fibre-tracking techniques are susceptible to noise, partial volume effects, diffusion anisotropy and other properties of the underlying structure. In this work, simulations were carried out to assess the reliability of a popular fibre-tracking algorithm when applied to artificial data sets designed to model *in-vivo* structures. Partial volume effects in particular were found to have a detrimental effect on the reliability of fibre tracking.

To address the issues raised above, a novel fibre-tracking technique is proposed, based on the evolution of a front. It is capable of branching and produces estimates of the likelihood of the connections established. The algorithm is designed to be more robust to noise and to deal with the problem of crossing fibres more adequately. Fibre tracts reconstructed using this technique are presented and discussed.

To illustrate a practical application of diffusion tensor MRI, the relationship between diffusion anisotropy and language deficits in patients with unilateral basal ganglia infarctions was investigated.

CONTENTS

1	Introduction	11
	References	14
2	Image acquisition	15
2.1	General principles	15
2.1.1	k -space	16
2.1.2	Properties of k -space	18
2.2	Echo-planar imaging	21
2.2.1	k -space sampling strategies	21
2.2.2	k -space trajectories	23
2.2.3	Multi-shot EPI	25
2.3	Artefacts in EPI	26
2.3.1	Bandwidth	26
2.3.2	Nyquist ghosting	27
2.3.3	Image distortions	31
2.3.4	Chemical shift artefacts	33
2.3.5	Limits on image resolution	34
	References	36
3	Diffusion in MRI	38
3.1	Introduction to diffusion	38
3.1.1	Molecular diffusion	38
3.1.2	The effect of diffusion on the NMR signal	39
3.1.3	Diffusion in biological systems	42
3.2	Diffusion-weighted imaging	45
3.2.1	The b -matrix	45
3.2.2	Diffusion-weighted sequences	49
3.2.3	Artefacts	52
3.3	Diffusion tensor imaging	56
3.3.1	The diffusion tensor	56
3.3.2	Diffusion encoding schemes	63
3.3.3	Tractography	64
3.3.4	Artefacts	68
3.3.5	Alternative models for diffusion	70

3.3.6	Unresolved issues in tractography	72
References		72
4	Imaging sequence implementation	78
4.1	Tetra-ortho sequence	79
4.2	20 direction sequence	81
4.2.1	Optimisation of diffusion gradient timing	82
4.2.2	Remaining eddy-current artefact reduction	84
4.2.3	Offline reconstruction	86
References		93
5	Simulations of diffusion tensor tractography	95
5.1	Introduction	95
5.2	Methods	97
5.2.1	Simulated data	97
5.2.2	Tracking algorithm	100
5.2.3	Simulations	100
5.2.4	Real Data	102
5.2.5	Measures of reliability of tracking	103
5.3	Results and discussion	104
5.3.1	Model A	104
5.3.2	Model B	109
5.3.3	Real Data	112
5.4	Conclusion	113
References		114
6	Front evolution tractography	116
6.1	Introduction	117
6.2	The FRET algorithm	118
6.2.1	Overview	118
6.2.2	The fibre orientation density function	119
6.2.3	Sampling orientation space	121
6.2.4	Front evolution	122
6.2.5	Surface reconstruction	124
6.2.6	Down-sampling the front	126
6.2.7	Propagation of the index of connectivity	127
6.2.8	Seed point selection	128
6.3	Methods	128
6.3.1	Data acquisition	128
6.3.2	Assessment of the FRET algorithm	128
6.3.3	Streamlines tracking algorithm	129
6.4	Results and discussion	129
6.4.1	Dependence on the signal to noise ratio	130
6.4.2	Dependence on initial seed region	135
6.4.3	Dependence on algorithm parameters	135
6.4.4	Dependence on partial volume effects	137

6.4.5 Comparison with streamlines	138
6.4.6 The index of connectivity	138
6.5 Conclusion	140
References	140
7 Unilateral basal ganglia infarction study	142
7.1 Introduction	143
7.2 Methods	146
7.2.1 Study population	146
7.2.2 Language assessment	146
7.2.3 Voxel-based morphometry	147
7.2.4 MR imaging and analysis	148
7.3 Results	151
7.3.1 Language performance	151
7.3.2 Conventional T_1 and T_2 weighted imaging	152
7.3.3 Voxel-based morphometry: structural imaging	152
7.3.4 Voxel-based morphometry: diffusion tensor imaging	156
7.4 Discussion	158
7.5 Conclusion	161
References	161
8 General discussion	164
8.1 Diffusion anisotropy	165
8.2 Diffusion tensor tractography	166
8.3 Remaining issues	168
References	169

LIST OF FIGURES

2.1	Conventional spin-echo pulse sequence	18
2.2	Image information in k -space	20
2.3	Blipped EPI read-out pulse sequence	22
2.4	Constant phase-encoding EPI read-out pulse sequence	23
2.5	Spiral EPI read-out pulse sequence	24
2.6	Interleaved EPI	26
2.7	Nyquist ghost artefact in EPI	27
2.8	Echo asymmetry caused by the time-reversal of alternate lines of k -space	28
2.9	Gradient pre-emphasis	29
2.10	Signal degradation due to the low-pass filter of the receiver system	30
2.11	Eddy-current-induced image distortions	32
2.12	Susceptibility-induced image distortion in EPI	33
2.13	Fat artefact in EPI	34
3.1	The Stejskal-Tanner pulsed-gradient spin-echo pulse sequence	40
3.2	Free and restricted diffusion	44
3.3	A pulsed-gradient sequence with cross-terms from the imaging gradients .	47
3.4	A pulsed-gradient sequence with minimal cross-terms	48
3.5	The diffusion-weighted stimulated echo sequence	51
3.6	The twice-refocused diffusion-weighted spin-echo sequence	52
3.7	Images acquired using 2D-FT, interleaved EPI and single-shot EPI	54
3.8	The diffusion ellipsoid	57
3.9	The shape of the diffusion tensor ellipsoid	58
3.10	Relating the diffusion-weighted signal to the diffusion tensor	60
3.11	Various maps available from the diffusion tensor	61
3.12	Voxel-linking tractography techniques	65
3.13	Streamlines tractography techniques	67
3.14	The effect of partial volume averaging	69
4.1	Tetrahedral-orthogonal pulse sequence	80
4.2	Maps generated using the tetrahedral-orthogonal sequence	81
4.3	Twice-refocused diffusion-weighted spin-echo sequence	82
4.4	Maps generated using the twice-refocused sequence	83
4.5	Eddy-current nulling using the twice-refocused sequence	84

4.6	Remaining eddy-current amplitude as a function of time constant	85
4.7	Minimising long term eddy-currents using a dummy loop	85
4.8	Overview of the image reconstruction process	87
4.9	Linear time sampling resulting in non-linear k -space sampling	89
4.10	The interlaced Fourier transform	91
5.1	Model A	97
5.2	Model B	98
5.3	The outcome measures	101
5.4	The <i>in vivo</i> fibre bundles used to test the tracking algorithm	102
5.5	Example tracks illustrating the effects of noise and interpolation	104
5.6	Model A: end-point x -offset as a function of step size	105
5.7	Model A: maximum departure as a function of step size	106
5.8	Model A: RMS departure as a function of step size	106
5.9	Model A: contours of RM_A as a function of SNR and step size	108
5.10	Model B: effect of interpolation on RM_B	109
5.11	Model B: effect of background anisotropy on RM_B	111
5.12	Real data: success rate as a function of number of acquisitions	112
6.1	Fibre ODF for various typical diffusion tensors	121
6.2	Increasing the sampling density along the most likely orientations	122
6.3	Overview of the front evolution process	123
6.4	Two possible approaches for front evolution	124
6.5	Bias introduced by discarding points that lie behind the front	125
6.6	Surface reconstruction	125
6.7	Downsampling the front	127
6.8	A track generated using FRET in the cortico-spinal tract	131
6.9	A similar track generated from a different data set	132
6.10	A track generated using FRET in the corpus callosum	133
6.11	A track generated using FRET from around the motor cortex	134
6.12	Dependence of FRET on the exact position of the initial seed point	135
6.13	A track generated in the cortico-spinal tract using a modified ODF	136
6.14	Effect of curvature threshold on tracks generated in the corpus callosum	137
6.15	Comparison of tracks generated using FRET and streamlines	139
7.1	Processing stages of voxel-based morphometry	147
7.2	Regions where small volume corrections were applied	150
7.3	Average language scores for both lesion groups	151
7.4	Individual language scores for patients in both lesion groups	152
7.5	Location of the lesion for each patient in the left-sided lesion group	153
7.6	Location of the lesion for each patient in the right-sided lesion group	153
7.7	Correlation with structural damage in left inferior frontal gyrus	154
7.8	Correlation with structural damage in left middle frontal gyrus	155
7.9	Correlation with structural damage in left insular cortex	155
7.10	Correlation with structural damage in left superior temporal gyrus	155
7.11	Correlation with structural damage in left deep frontal white matter	156
7.12	Correlation with fractional anisotropy in the left gyrus rectus	157

7.13 Correlation with fractional anisotropy in the left internal capsule	157
7.14 Correlation with fractional anisotropy in the left deep frontal white matter	157

LIST OF TABLES

3.1	Typical diffusion anisotropy values in various regions of the brain	62
4.1	Tetrahedral-orthogonal diffusion-encoding gradient directions	80
7.1	Description of the patients used in the study	146
7.2	Correlation between language function and structural damage	154
7.3	Correlation between language function and fractional anisotropy	156

ACKNOWLEDGEMENTS

I am hugely indebted to my two supervisors, Alan Connelly and Fernando Calamante for their continuous support throughout my PhD. They were always available to discuss any issue or problem that I faced, and always willing to give up their time to provide me with the advice and guidance that I needed. I would also like to express my gratitude to David Gadian for welcoming me in his lab, and giving me the opportunity to carry out this work. I am very grateful to all of them for their encouragement and friendship.

There are a number of people who have made a significant contribution to the work presented in this thesis. Through helpful discussions and friendly debates, they have provided me with new insights, and taught me more than they know. In particular, I am grateful to Martin King for never taking facts at face value and bringing my implicit assumptions to light. I am also grateful to the neurologists and neuroradiologists at Great Ormond Street Hospital who have taken the time to discuss and explain some of my stranger findings with me, in particular Kling Chong, Tim Cox, Dawn Saunders, Rod Scott and Vijeya Ganesan. I also wish to thank Alison Rowan, without whom some of the work presented here would not have been possible.

This work would also not have been possible without the financial and technical support of Siemens (Erlangen). In particular, I would like to thank David Porter for taking the time to discuss some of my technical problems, and providing me with helpful information.

I would also like to thank all the people that I have been lucky enough to work with over the past few years. I am grateful to Jo Perthen for her friendship, humour (intentional or otherwise) and for trying to improve my manners, albeit with limited success. I am also grateful to the MR research radiographers at Great Ormond Street Hospital, in particular Heather Dacie, Jane Ho and Lindy McWhirter for their optimism, friendship and sometimes dubious humour. I would also like to thank all the people from the Biophysics department, the UCL Medical Physics department, and the Wolfson Centre who have helped me and made my experience more enjoyable.

Finally, I would like to express my deepest gratitude to Joanna and to my family, who have supported me throughout, and always encouraged me to do my own thing.

PUBLICATIONS

Publications in peer-reviewed journals and international meetings during PhD thesis:

- [1] TOURNIER, J.-D., CALAMANTE, F., KING, M. D., GADIAN, D. G., AND CONNELLY, A. Limitations and requirements of diffusion tensor fiber tracking: an assessment using simulations. *Magnetic Resonance in Medicine* 47 (2002), 701–708.
- [2] TOURNIER, J.-D., CALAMANTE, F., GADIAN, D. G., AND CONNELLY, A. Diffusion-weighted MRI fibre-tracking using a front evolution algorithm. *NeuroImage* (2003), *accepted subject to revision*.
- [3] TOURNIER, J.-D., CALAMANTE, F., KING, M. D., GADIAN, D. G., AND CONNELLY, A. Limitations and requirements of diffusion tensor fibre-tracking: assessment using simulations. In *Proceedings of the ISMRM* (2001), vol. 9, p. 1521.
- [4] TOURNIER, J.-D., CALAMANTE, F., GADIAN, D. G., AND CONNELLY, A. A novel fibre-tracking technique: front evolution using a fibre orientation probability density function. In *Proceedings of the ISMRM* (2002), vol. 10, p. 538.
- [5] TOURNIER, J.-D., ROWAN, A., CALAMANTE, F., CONNELLY, A., BALDEWEG, T., VARGHA-KHADEM, F., AND GADIAN, D. G. Changes in grey and white matter structures associated with language in patients with acquired unilateral basal ganglia infarction revealed by structural and diffusion tensor MRI. In *Proceedings of the ISMRM* (2003), vol. 11, p. 401.

CHAPTER 1

INTRODUCTION

Magnetic resonance imaging (MRI) is a powerful imaging tool that can be sensitised to different properties of tissue, enabling images to be obtained with a wide range of contrasts. In addition to this versatility, MRI has the important advantage of being entirely non-invasive, with no known biologically hazardous effects. For these reasons, its use is becoming increasingly widespread, both in the clinical and research environments.

During the last 10 years, a particular form of image contrast, namely diffusion-weighted imaging (DWI) has become an established MRI technique. DWI is sensitised to the microscopic random motions of water molecules in the tissue, which are intimately linked with the physical arrangement of their immediate surroundings. Diffusion-weighted imaging thus offers a unique means of probing tissue microarchitecture at the cellular level. One of the most useful clinical applications of diffusion-weighted imaging is in the assessment of acute stroke patients: the region of ischaemia can be visualised within minutes of infarct onset, and hours before becoming apparent using more conventional MRI techniques [1].

In 1990, Moseley *et al* [2] showed that in the white matter and spinal cord of the cat, the diffusion coefficient measured using diffusion-weighted MRI was dependent on the direction of the measurement. This *anisotropy* was attributed to the coherent arrangement of fibres in those tissues causing a greater hindrance to the free movement of water molecules in directions perpendicular to the orientation of the fibres. In order to characterise this anisotropic behaviour, the diffusion *tensor* model was introduced in 1994 by Basser *et al* [3]. From this mathematical formalism, a number of rotationally invariant properties can be derived, in particular the *trace* (corresponding to the magnitude of the diffusion) and the *major eigenvector* (corresponding to the direction of fastest diffusion). Various measures of diffusion anisotropy have also been proposed, such as the *fractional anisotropy* (FA) [4], and are increasingly being used in interesting new applications (see Horsfield *et al* [5] for a recent review).

White matter consists of tightly packed bundles of neuronal axons that connect the

different cortical areas to each other and to the rest of the body. The coherent arrangement of these fibres results in pronounced diffusion anisotropy that can readily be observed using diffusion tensor imaging (DTI). The magnitude of this anisotropy depends on a number of factors, including axonal density and the degree of myelination [6, 7]. Maps of diffusion anisotropy may therefore be useful in assessing the integrity of white matter tissue, and have been used to investigate a number of white matter diseases [5].

DTI has also made it possible to estimate the orientation of the white matter fibres by equating it to the direction of fastest diffusion. Based on this assumption, a number of techniques have recently been proposed to map the path of white matter tracts by following the direction of fastest diffusion from point to point [8, 9]. Using such techniques, *in vivo* white matter tractography may be feasible, opening up the possibility of studying the connectivity of the brain *in vivo*.

The acquisition of diffusion tensor data is demanding on the scanner hardware, particularly the gradients. However, recent technological advances have allowed most modern MR scanners to meet those requirements. DTI is therefore becoming more common, and thus increasingly important.

This thesis therefore had four main objectives:

- to implement a diffusion tensor acquisition sequence for use in our institution
- to investigate the limitations and requirements of diffusion tensor tractography
- to develop a robust novel tractography algorithm
- to investigate potential applications for diffusion tensor imaging

The basics of MR image acquisition necessary for an understanding of diffusion-weighted MRI are given in chapter 2. The diffusion process and its effect on the NMR signal is described in chapter 3. This chapter also details MRI acquisition issues particular to DTI, and reviews some of the techniques currently available for the analysis of diffusion tensor data.

The first of the above objectives is addressed in chapter 4, where the image acquisition sequence implemented in our institution is described. This sequence was designed to have reduced sensitivity to eddy-current effects, while remaining time-efficient and providing images of optimal quality. The diffusion weighting was applied along 20 directions to minimise bias in the calculation of the diffusion tensor [10], and image reconstruction was performed offline to minimise the scan time.

The second objective is addressed in chapter 5, which describes numerical simulations performed using a well-established algorithm on synthetic data. Since there is currently no appropriate ‘gold standard’ for tractography, simulations on synthetic data where the

true outcome is known provide an alternative method of validating fibre-tracking algorithms. The results show the impact of noise, partial volume effects, and other parameters on the reliability of tracking, and provide a means to determine the minimum image quality (in terms of resolution and signal to noise ratio) required to track a given structure of interest successfully. The algorithm was also used on real data acquired on a volunteer to verify the relevance of the simulations to the *in vivo* case.

The results of the above simulations also provided a clearer focus to address the third issue: chapter 6 describes the concept behind and the implementation of a novel tractography algorithm, designed to address some of the limitations of current techniques, in particular susceptibility to noise and the lack of an indication of the likelihood of a connection. Results obtained using this algorithm are presented that are consistent with known anatomy. However, there are remaining limitations to the use of this algorithm, and these are discussed in the text.

Chapter 7 presents results of a study performed in collaboration with a neuropsychology PhD student, Alison Rowan. A group of patients with acquired unilateral basal ganglia infarctions was investigated using behavioural measures of speech and language, and a range of MRI techniques including diffusion tensor imaging. The role of the basal ganglia in the mediation of speech and language deficits in these patients is poorly understood. The aim of this study was to investigate potential additional abnormalities that may not readily be observed using conventional imaging. Using more advanced MR imaging and analysis methods, abnormalities remote from the site of the lesion were identified that showed a significant correlation with language performance. This study also illustrates the practical application of diffusion anisotropy maps.

Finally, chapter 8 discusses the issues raised in the preceding chapters, and highlights possible avenues for further research.

All the work presented in this thesis has been performed entirely by the author, except for chapter 7, which was performed in collaboration with a neuropsychology PhD student, Alison Rowan. The author's contribution to this chapter included the development and implementation of software for the image reconstruction (described in chapter 4) and for the generation of the parameter maps. This required a more user-friendly interface to the software to be written in order that the data processing could be collaborative. The author was also responsible for training Ms Rowan in the use of the software, and in the identification of corrupted or artefacted images. The latter required input from the author throughout the data processing period. The acquisition of the neuropsychology data and the subsequent correlation analyses were performed by Alison Rowan.

REFERENCES

- [1] BAIRD, A. E., AND WARACH, S. Magnetic resonance imaging of acute stroke. *Journal of Cerebral Blood Flow and Metabolism* 18 (1998), 583–609.
- [2] MOSELEY, M. E., COHEN, Y., KUCHARCZYK, J., MINTOROVITCH, J., ASQARI, H. S., WENDLAND, M. F., TSURUDA, J., AND NORMAN, D. Diffusion-weighted MR imaging of anisotropic water diffusion in cat central nervous system. *Radiology* 176 (1990), 439–445.
- [3] BASSER, P. J., MATTIELLO, J., AND LE BIHAN, D. MR diffusion tensor spectroscopy and imaging. *Biophysics Journal* 66 (1994), 259–267.
- [4] BASSER, P. J. Inferring microstructural features and the physiological state of tissues from diffusion-weighted images. *NMR in Biomedicine* 8 (1995), 333–344.
- [5] HORSFIELD, M. A., AND JONES, D. K. Application of diffusion-weighted and diffusion tensor MRI to white matter diseases – a review. *NMR in Biomedicine* 15 (2002), 570–577.
- [6] BEAULIEU, C., AND ALLEN, P. S. Determinants of anisotropic water diffusion in nerves. *Magnetic Resonance in Medicine* 31 (1994), 394–400.
- [7] BEAULIEU, C., AND ALLEN, P. S. Water diffusion in the giant axon of the squid: implications for diffusion-weighted MRI of the nervous system. *Magnetic Resonance in Medicine* 32 (1994), 579–583.
- [8] MORI, S., CRAIN, B. J., CHACKO, V. P., AND VAN ZIJL, P. C. M. Three-dimensional tracking of axonal projections in the brain by magnetic resonance imaging. *Annals of Neurology* 45 (1999), 265–269.
- [9] CONTURO, T. E., LORI, N. F., CULL, T. S., AKBUDAK, E., SNYDER, A. Z., SHIMONY, J. S., MCKINSTRY, R. C., BURTON, H., AND RAICHLE, M. E. Tracking neuronal fiber pathways in the living human brain. *Proceedings of the National Academy of Sciences* 96 (1999), 10422–10427.
- [10] PAPADAKIS, N. G., MURRILLS, C. D., HALL, L. D., HUANG, C. L.-H., AND CARPENTER, A. A. Minimal gradient encoding for robust estimation of diffusion anisotropy. *Magnetic Resonance Imaging* 18 (2000), 671–679.

CHAPTER 2

IMAGE ACQUISITION

CONTENTS

2.1	General principles	15
2.1.1	<i>k</i> -space	16
2.1.2	Properties of <i>k</i> -space	18
2.2	Echo-planar imaging	21
2.2.1	<i>k</i> -space sampling strategies	21
2.2.2	<i>k</i> -space trajectories	23
2.2.3	Multi-shot EPI	25
2.3	Artefacts in EPI	26
2.3.1	Bandwidth	26
2.3.2	Nyquist ghosting	27
2.3.3	Image distortions	31
2.3.4	Chemical shift artefacts	33
2.3.5	Limits on image resolution	34
References		36

2.1 GENERAL PRINCIPLES

Following a radio frequency (RF) excitation pulse, nuclear spins in a static magnetic field B_0 will precess about the field's axis at the Larmor frequency, $\omega = \gamma B_0$, where γ is the magnetogyric ratio, which relates the magnetic moment of the nucleus to its angular momentum [1]. Consider now the effect of applying a magnetic field gradient in addition to the main static B_0 field¹. Suppose a gradient $G_x = dB_z/dx$ is applied. The Larmor frequency is now:

$$\omega = \gamma(B_0 + G_x x) \quad (2.1)$$

¹By convention, B_0 is aligned along the z -axis

The resonant frequency of the spins is now position-dependent. Frequency analysis of the free induction decay (FID) in the presence of such a gradient shows a distribution of frequencies representing the spin density profile of the sample along the x -axis (ignoring relaxation effects). By acquiring a number of such profiles with the gradient oriented along a range of angles, enough information is acquired to deduce the internal distribution of spins within the sample, by means of the ‘back projection’ algorithm. This is the basis of the earliest method of MRI, known as projection reconstruction [2].

Although projection reconstruction is still used in a number of specific applications, there are drawbacks associated with its use [3]. In particular, the sampling of the data is not efficient: the centre of k -space is much more densely sampled than the outer regions of k -space (see next section for a description of k -space). In addition, signal that is off-resonance will be smeared around its true position, giving rise to characteristic ‘half-moon’ artefacts. Consequently, most techniques commonly in use today are based on Fourier analysis and the concept of k -space.

2.1.1 k -SPACE

The signal $S(\underline{x}, t)$ in the sample at position \underline{x} and time t can be represented by a complex number, having amplitude $\rho(\underline{x})$ and phase $\phi(\underline{x}, t)$:

$$S(\underline{x}, t) = \rho(\underline{x}) e^{-i\phi(\underline{x}, t)} \quad (2.2)$$

Ignoring the effects of relaxation, $\rho(\underline{x})$ represents the spin density of the sample as a function of position. The phase of the signal increases at a rate given by its angular velocity $\omega(\underline{x}, t)$:

$$\phi(\underline{x}, t) = \int_0^t \omega(\underline{x}, t') dt' \quad (2.3)$$

The angular velocity is given by the Larmor equation, and depends on the local magnetic field. In MRI, the field normally consists of a large static field B_0 and an applied gradient $\underline{G}(t)$. The Larmor equation then gives (see equation 2.1):

$$\omega(\underline{x}, t) = \gamma(B_0 + \underline{G}(t) \cdot \underline{x}) \quad (2.4)$$

The magnitude of B_0 is several orders of magnitude larger than the gradient-induced field, and represents a fast constant phase velocity that will be identical for all spins in the sample. Since it is relative phase differences that are of interest, the B_0 term is ignored, and the relative phase is given by:

$$\phi(\underline{x}, t) = \gamma \left[\int_0^t \underline{G}(t') dt' \right] \cdot \underline{x} = \underline{k}(t) \cdot \underline{x} \quad (2.5)$$

where we have defined the k -space vector $\underline{k}(t)$ as:

$$\underline{k}(t) = \gamma \int_0^t \underline{G}(t') dt' \quad (2.6)$$

If we now combine equations 2.2, 2.5 and 2.6, and add up the contributions to the signal from the whole sample, the signal can be expressed as a function of time:

$$S(t) = \int_V \rho(\underline{x}) e^{-i\underline{k}(t) \cdot \underline{x}} d^3 \underline{x} \quad (2.7)$$

Since $\underline{k}(t)$ is a function of time, this equation can be reformulated as a function of \underline{k} alone, such that the time variable does not need to be explicitly stated:

$$S(\underline{k}) = \int_V \rho(\underline{x}) e^{-i\underline{k} \cdot \underline{x}} d^3 \underline{x} \quad (2.8)$$

This is the Fourier transform equation, stating that $S(\underline{k})$ is the Fourier transform of the spin density distribution² $\rho(\underline{x})$ [4].

Conventional imaging sequences are based on the 2D Fourier transform. If a single slice is excited, for example that normal to the z -axis (by applying a magnetic field gradient along z during the RF excitation pulse), the equation for the signal $S(\underline{k})$ reduces to its 2D equivalent:

$$S(\underline{k}) = d \int_A \rho(\underline{x}) e^{-i\underline{k} \cdot \underline{x}} d^2 \underline{x} \quad (2.9)$$

where d is the slice thickness, and \underline{x} is now a 2D vector in the image plane. In this case the z dimension is already in image space, and a number of such slices can be acquired sequentially at various positions along z in a so-called multislice experiment. A typical 2D-FT spin-echo imaging sequence and its corresponding k -space sampling scheme are illustrated in figure 2.1. Following the slice-selective excitation, gradients are applied along the x and y axes, such that the position in k -space is shifted to a desired location. As described later, the 180° pulse reverses the position in k -space. A gradient is then applied along x while the signal is sampled. In this example, the x direction corresponds to the *frequency-encoding* or *read-out* axis, while the y direction corresponds to the *phase-encoding* axis. This process is repeated after a suitable relaxation time TR with an increment in the phase-encoding gradient until adequate coverage of k -space is achieved. The image of the spin density distribution is then calculated with a simple inverse 2D fast Fourier transform (FFT) [5]. It is also possible to apply equation 2.8 directly, by applying the excitation pulse without slice selection and using phase-encoding along z in addition to y , in a so-called 3D experiment. In this case, an extra FFT is also performed in the z

²Note that in a practical imaging situation, relaxation and other contrast mechanisms will also affect the magnitude of the signal, such that $\rho(\underline{x})$ will not depend solely on spin density

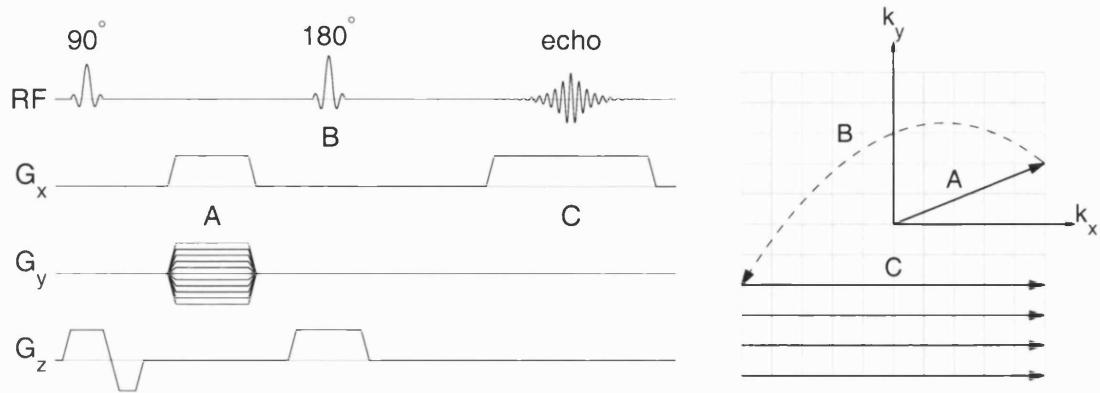


Figure 2.1: conventional spin-echo pulse sequence (left), and a schematic representation of its k -space trajectory (right). k_x corresponds to the read direction of k -space, while k_y corresponds to the phase-encoding direction. Multiple RF excitations are required to cover k -space. After each RF excitation, gradient pulses in the read (the ‘read dephasing gradient’) and phase-encoding direction (‘phase-encoding step gradient’) are applied to move the k -space position away from the centre (A). The 180° pulse (B) reverses the position to the start of the current line. Data are then collected along that line during the application of the read-out gradient pulse (C).

direction.

2.1.2 PROPERTIES OF k -SPACE

During an imaging experiment, the NMR signal will be sampled at discrete time points. The Nyquist criterion states that in order to accurately represent the original signal, the interval δt between successive time samples must be less than half the period of the maximum angular frequency component ω_{max} contained within the signal [6]. In other words, the sampling rate $1/\delta t$ must be higher than $2(\omega_{max}/2\pi)$. Higher frequency components will be *aliased*, such that their original angular frequency ω will be interpreted as $\omega - 2\omega_{max}$.

This treatment can be applied to k -space, by recognising that the frequency measured in a MRI experiment is the difference between the reference frequency, as measured in a static B_0 field, and the true frequency of the spins in the presence of the imaging gradients [7]. It is then clear that the maximum frequency components in the measured signal emanate from the outermost edges of the sample, where the effect of the gradients is greatest. However, the frequency on one side of the sample will be negative with respect to the frequency on the other side. To differentiate between positive and negative frequencies, the phase of the signal must also be known, and this is achieved using quadrature detection [8]. Using the Nyquist criterion, the maximum spatial extent that can be correctly encoded in k -space is thus given by the k -space sampling rate, and is called the field of

view (FOV):

$$\text{FOV} = \frac{2\pi}{\delta k} \quad (2.10)$$

where the factor of 2 has been dropped since the FOV extends from $-x_{max}$ to $+x_{max}$. If the sample extends beyond the FOV, the signal emanating from that region will be *aliased* or *wrapped-around*, such that it will appear to originate from a region displaced by one FOV, to the other side of the image.

In the same way, the extent of k -space sampled, k_{max} , corresponds to the maximum spatial frequency that can be accurately represented in the image, and is related to the image resolution or pixel size, δx , as follows:

$$k_{max} = \frac{\pi}{\delta x} \quad (2.11)$$

The fact that the spin density $\rho(\underline{x})$ is a real quantity leads to another property of k -space, namely its conjugate symmetry. This is expressed mathematically as (see equation 2.8):

$$S(\underline{k}) = [S(-\underline{k})]^* \quad (2.12)$$

where z^* denotes the complex conjugate of z . This property can be exploited to shorten acquisition times, by acquiring fewer samples than otherwise required, and reconstructing the missing data using equation 2.12. This is the basis of partial Fourier acquisitions, described in section 2.2.2.

The centre of k -space contains information about the low spatial frequency components contained in the sample spin density distribution, whereas regions further away from the centre of k -space contain information about the higher spatial frequency components. An image formed from a k -space sample where the outer regions of k -space have been discarded thus contains very little fine detail or edge information (see figure 2.2), whereas the overall intensity for the different parts of the image is correct. On the other hand, an image formed from a k -space sample where the central region of k -space has been discarded contains only edge information, and the overall intensity for the different parts of the image is close to zero. Thus, edge information is contained in the outer regions of k -space, whereas the regional signal intensity information, or contrast, is contained near the centre of k -space. For this reason, the time at which the centre of k -space is acquired is of special significance, since it determines the overall contrast of the image.

It is also important to consider the effect of a 180° refocusing pulse on the position in k -space. A signal with phase ϕ will have its phase reversed to $-\phi$ by a 180° . In other words, the complex signal S will be transformed into its complex conjugate S^* . Using the conjugate symmetry of k -space, it can be seen that this corresponds to transforming

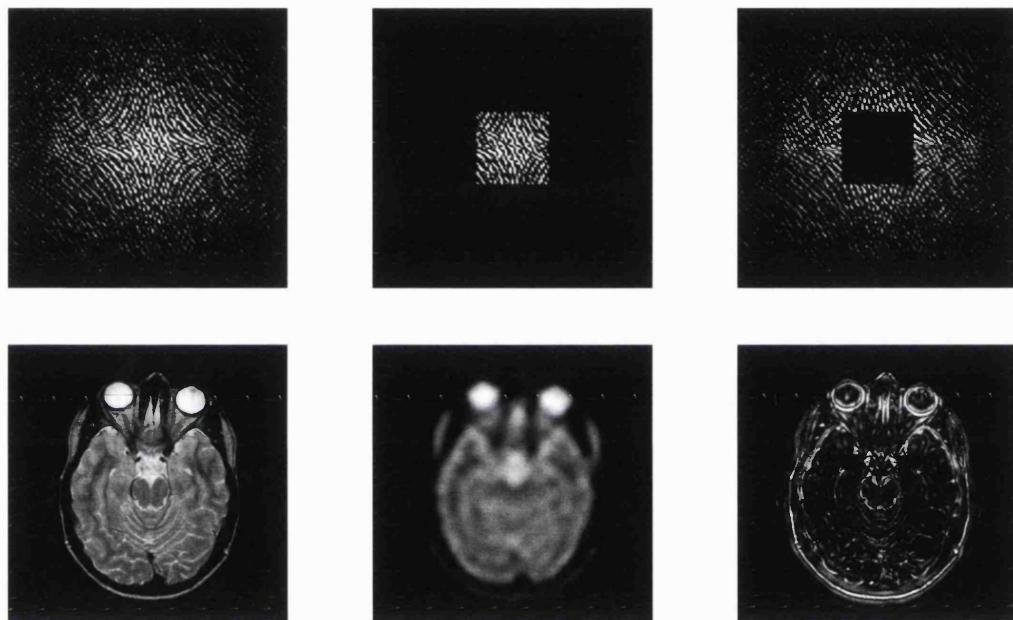


Figure 2.2: image information in k -space. Data acquired near the centre of k -space contain low spatial frequency information about the image, while data acquired near the periphery represent high spatial frequencies. Top: k -space data. Bottom: reconstructed image. Left: the original k -space data and the corresponding image. Middle: data outside of the centre of k -space have been set to zero, leaving only the low spatial frequency information. Note the loss of fine detail in the resulting image. Right: data near the centre of k -space have been set to zero, leaving only the high spatial frequency information. Fine detail in the resulting image are restored, but information about the overall contrast is lost.

the position in k -space from \underline{k} to $-\underline{k}$.

2.2 ECHO-PLANAR IMAGING

Equation 2.9 indicates that if data are acquired such that k -space is adequately sampled, then an image can be formed. 2D-FT techniques acquire one line of k -space per excitation, which makes them relatively slow. However, it is in theory possible to acquire all the data necessary for an image from a single excitation. This is the basis of echo-planar imaging (EPI) [9].

To acquire all of the k -space data required from a single excitation, the gradients must be manipulated such that an appropriate trajectory in k -space is sampled. To produce artefact-free images, such trajectories must sample k -space such that the Nyquist criterion is satisfied. However, the time available for sampling is limited, so the traversal of k -space must be as fast and efficient as possible. This will tend to place extreme demands on the gradient and receiver hardware [10]. Another important consideration is the complexity of the image reconstruction algorithm, which will depend on the characteristics of each trajectory. These factors will tend to impose sometimes conflicting requirements on EPI acquisition schemes, which are described in detail in this section.

2.2.1 k -SPACE SAMPLING STRATEGIES

There are two main approaches to k -space sampling: one aims to minimise hardware demands, while the other aims to minimise image reconstruction complexity. Although these two approaches are not in principle mutually exclusive, in practice sampling schemes that attempt to combine the benefits of both strategies will not be as fast and efficient as otherwise possible, and are rarely used in practice.

LINEAR SAMPLING IN k -SPACE

The simplest image reconstruction algorithm is the inverse 2D FFT. However, to be able to use such a simple approach, the data must be sampled on a regular Cartesian grid in k -space [11]. Sampling strategies that aim to minimise the complexity, and hence computation time of the reconstruction must sample the signal such that each data point is acquired on a Cartesian grid node in k -space. This is achieved using a k -space traversal such as that illustrated in figure 2.3. Although using a constant frequency-encoding gradient while sampling each line would result in a regular time sampling rate, such a scheme would incur a large time penalty from the gradient ramp-up dead time. In practice data must be acquired during the gradient ramps, such that the sampling rate cannot

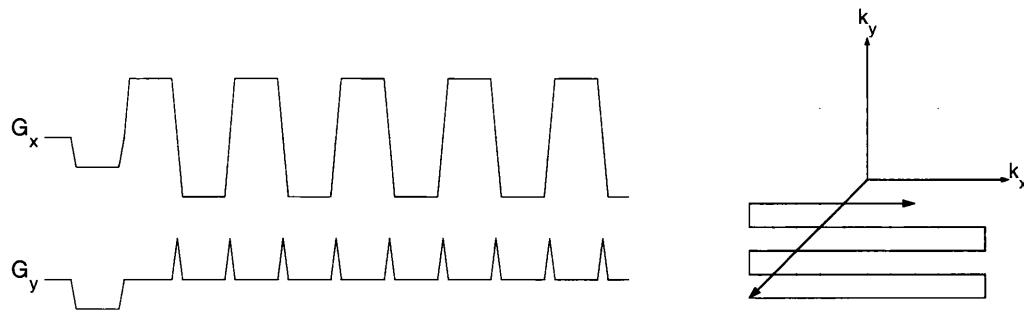


Figure 2.3: EPI read-out pulse sequence with a blipped phase-encoding gradient. Left: a schematic representation of the sequence. The oscillatory read gradient (G_x) generates multiple echoes, which are encoded differently by the successive phase-encoding blips (G_y). Right: a schematic representation of the corresponding k -space trajectory. The entire k -space can thus be scanned from a single RF excitation. In contrast to the conventional spin-echo sequence shown in figure 2.1, alternate lines of k -space are traversed in opposite directions.

be constant in time if the data are to be regularly sampled in k -space. The receiver must therefore be able to sample the signal at non-constant intervals, and must be accurately synchronised with the gradient waveform.

This type of EPI readout therefore requires specialised hardware, which most modern clinical scanners are equipped with. The receiver can be synchronised with the gradients using preparation scans (see section 2.3.2), which must be performed prior to the EPI sequence. Another drawback of this type of readout is that the receiver filter bandwidth cannot in general be adjusted while the sequence is running. This means that it must be set to accommodate the highest sampling rate present, such that the image SNR is lower than might ideally be achieved using a variable bandwidth.

LINEAR SAMPLING IN TIME

Acquiring data with a constant sampling rate is much less demanding on the receiver hardware. However, it then becomes impractical to sample data on a Cartesian grid in k -space, and image reconstruction will be significantly more complex than with linear sampling in k -space. The data will often need to be re-gridded onto a Cartesian grid, which is a very computer-intensive process. There may therefore be a significant time delay between running the sequence and obtaining the images. Since some scanners do not allow other sequences to run until the reconstruction has completed, this may introduce a significant amount of ‘dead time’ in the examination. In addition, since the Nyquist criterion needs to be satisfied at all points in k -space, the number of time samples required will be greater than the minimum number achieved with regular sampling.

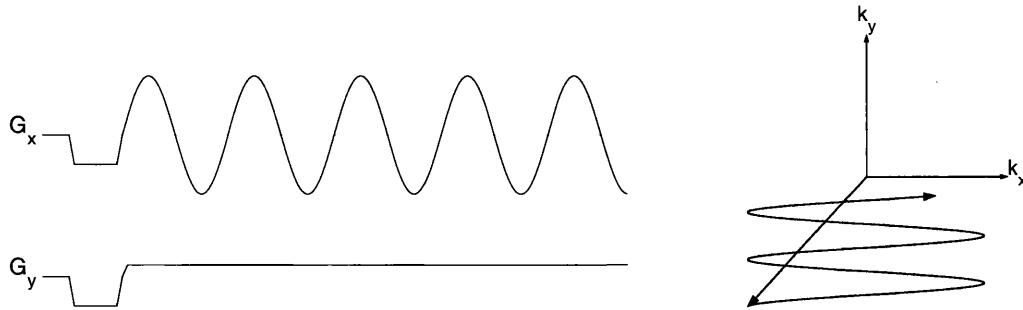


Figure 2.4: EPI read-out pulse sequence with a constant phase-encoding gradient. Left: a schematic representation of the sequence. The oscillatory read gradient (G_x) generates multiple echoes, which are encoded differently by the continuous application of the phase-encoding gradient (G_y). Right: a schematic representation of the corresponding k -space trajectory. The entire k -space can thus be scanned from a single RF excitation. As for the blipped EPI sequence (figure 2.3), alternate lines of k -space are traversed in opposite directions, although the trajectory is no longer collinear with the Cartesian grid.

2.2.2 k -SPACE TRAJECTORIES

The echo time in conventional spin-echo sequences is the same for each line of k -space. In EPI however, each line is acquired with a different echo time. Since the image contrast comes from the central portion of k -space, its echo time is used as the *effective* echo time of the EPI readout³.

CONSTANT PHASE-ENCODING EPI

EPI images can be acquired in the presence of a constant phase-encoding (CPE) gradient, such as that illustrated on figure 2.4. Since linear sampling in k -space is not possible using this type of trajectory, the reconstruction of the images tends to be computer-intensive. For this reason, however, there is little motivation to use non-linear sampling in time, and linear sampling in time tends to be used, which is less demanding on the receiver hardware.

BLIPPED EPI

The other common EPI readout trajectory is based on scanning through k -space as illustrated on figure 2.3. The rectilinear readout is achieved by applying a short duration phase-encoding gradient pulse, or ‘blip’, between every frequency-encoding gradient pulse, that increments k_y before each line is acquired along k_x . Blipped EPI uses linear sampling in

³In this work, the terms ‘echo time’ and ‘effective echo time’ are used interchangeably when referring to EPI sequences.

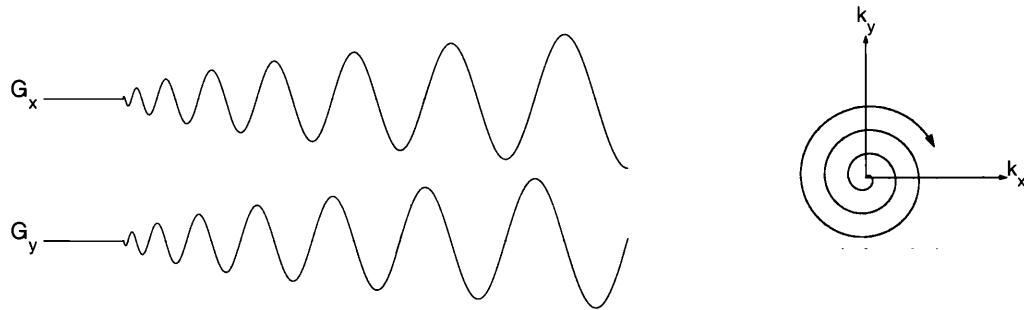


Figure 2.5: spiral EPI read-out pulse sequence. Left: a schematic representation of the sequence. The read and phase-encoding gradient oscillate with a 90° relative phase lag, generating the k -space trajectory shown schematically on the right. The entire k -space can thus be scanned from a single RF excitation. This sequence allows very short echo times, since the centre of k -space is sampled first. As for the constant phase-encoding EPI sequence (figure 2.4), the trajectory is not collinear with the Cartesian grid.

k -space, such that data samples are acquired on a Cartesian grid in k -space. Therefore, the main advantage of this trajectory is the simplicity of the image reconstruction.

SPIRAL EPI

This k -space trajectory involves starting from the centre of k -space and spiralling outwards to the edges (figure 2.5). This is done by oscillating the phase-encoding and frequency-encoding gradients with an increasing amplitude and a relative phase lag of 90°. This readout therefore allows the acquisition of images with a very short echo time (TE), since the centre of k -space is acquired first. The readout is intrinsically flow compensated [12], which is advantageous for MRA applications. The acquisition time is also shorter, since the corners of k -space are not acquired. This readout is less demanding on the gradient hardware, since the gradients do not need to be switched as rapidly. However, the scanner's pulse programming environment may not allow the application of the required gradient waveforms. In addition, spiral EPI is difficult to reconstruct, and artefacts manifest themselves differently than with blipped or CPE EPI, due to the fact that the bandwidth is not constant along either axes (see section 2.3.1).

PARTIAL FOURIER EPI

Due to the conjugate symmetry of k -space (see section 2.1.2), it is possible to acquire fewer data than otherwise required, and estimate those missing samples using equation 2.12. This can be used as an alternative method to shorten TE while reducing artefacts associated with a long readout, in particular T_2^* blurring (see section 2.3.5) and flow artefacts [12]. It does however lead to a reduction in SNR. Using this technique, acquisition

schemes can acquire as little as half of the data and still produce an image. However, in practice, a few extra lines near the centre of k -space are often acquired to correct for phase errors [13].

2.2.3 MULTI-SHOT EPI

The long readout period in EPI acquisition schemes introduces a number of problems, described in detail in section 2.3. Many of these problems can be alleviated using multi-shot EPI sequences, where different portions of k -space are acquired for each excitation, such that the combination of these different portions produces a complete k -space sample. These techniques allow the acquisition of higher resolution images with reduced artefacts. However, the time taken to acquire all the data necessary to reconstruct an image is markedly increased, since a suitable TR must be introduced to allow sufficient recovery of the signal before the next excitation. In addition, the use of these techniques to acquire diffusion-weighted images remains problematic due to the extreme motion sensitivity of the phase of the signal (see section 3.2.3).

There are a number of ways to divide up k -space in a manner suitable for multi-shot EPI. In general, care must be taken not to introduce discontinuities in the k -space signal that will cause image artefacts. Two commonly used multi-shot EPI schemes are described below.

MOSAIC EPI

In this scheme, k -space is divided into blocks to be acquired separately. It can easily be combined with a partial Fourier acquisition to further reduce scan time. Although this technique does reduce the readout period, the bandwidth per pixel along the phase-encoding axis remains similar to that of single-shot EPI techniques, such that bandwidth-related artefacts remain problematic (see section 2.3.1).

INTERLEAVED EPI

Interleaved EPI makes use of a larger phase-encoding gradient to acquire lines of k -space that are spaced further apart (figure 2.6). The next excitation can then be used to acquire the lines adjacent to those acquired previously, and so on. This results in a significant increase in the bandwidth per pixel along the phase-encoding direction, such that artefacts due to field inhomogeneities and chemical shift are much reduced (see section 2.3.3). For this reason, interleaved EPI is much more commonly used than mosaic EPI.

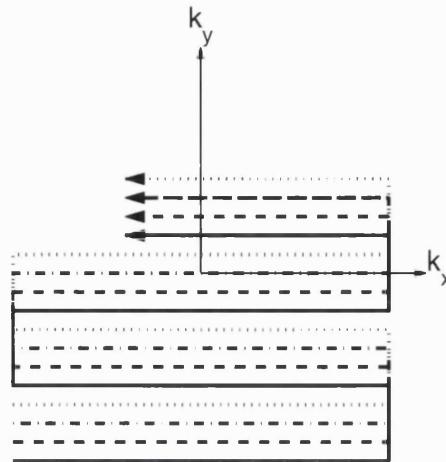


Figure 2.6: a schematic representation of interleaved EPI. Each linestyle represents the k -space trajectory of a single interleave, each acquired from a different RF excitation. Each acquisition starts from a slightly different position along k_y , such that the entire k -space is covered by the combination of all the interleaves. Within each interleave, alternate lines are spaced further apart than in single shot EPI, by applying a stronger phase-encoding gradient. This results in increased bandwidth per pixel and reduced artefacts due to field inhomogeneities and chemical shift.

2.3 ARTEFACTS IN EPI

This section concentrates on artefacts particular to the more commonly used blipped and CPE EPI readouts. Imaging artefacts using these types of k -space trajectories are due to the long readout period, the small bandwidth along the phase-encoding axis (see below), and the fact that alternate lines of k -space are acquired in opposite directions along the frequency-encoding axis. Image artefacts for other k -space trajectories have similar causes, but will manifest themselves differently depending on the details of the trajectory (see [12] for more details).

2.3.1 BANDWIDTH

Many of the artefacts commonly seen in EPI images are due to the different bandwidth per pixel (the frequency difference between adjacent voxels) along the phase-encoding axis from that along the frequency-encoding axis. If the dwell time (the time between adjacent samples) along the frequency-encoding axis is Δt , then the dwell time along the phase-encoding axis is approximately $N\Delta t$, where N is the number of samples along the frequency-encoding axis. Since the bandwidth per pixel is inversely proportional to the dwell time, it is much smaller along the phase-encoding axis. For example, a 128×128 EPI acquisition will typically have a $5 \mu\text{s}$ dwell time, corresponding to a 12 Hz

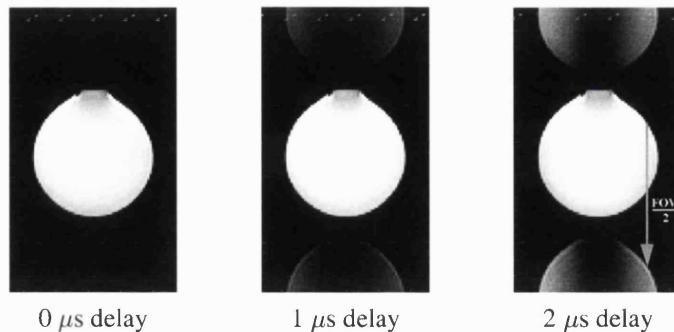


Figure 2.7: Nyquist ghost artefact in EPI. A ghost image is clearly seen displaced half the FOV from the main image. In this case, the source of the ghost is mainly a small misalignment (delay) between the gradient waveform and the analog-to-digital converter. As can be seen, the intensity of the ghost increases with increasing delay. (Adapted from F. Calamante's PhD thesis, 2000)

bandwidth per pixel along the phase-encoding axis. At 1.5 T, a 1 ppm frequency offset would then give rise to a 5 pixel shift along the phase-encoding axis, and only a 0.04 pixel shift along frequency-encoding axis. For this reason, small frequency offsets (due for example to field inhomogeneities or chemical shifts) will have a significant effect along the phase-encoding axis, and artefacts along the frequency-encoding axis can to all intents and purposes be ignored.

Bandwidth will also have an impact on the SNR. Due to the rapid traversal of k -space, the sampling rate needs to be very high. The bandwidth of the receiver therefore needs to be larger for EPI than it is for conventional sequences. This has a considerable impact on the SNR, since $\text{SNR} \propto 1/\sqrt{\text{BW}}$, where BW is the receiver bandwidth. A larger bandwidth will mean that noise from a broader range of frequencies will contribute to the measured signal, and translates to a reduced SNR for EPI images. However, EPI does provide a much higher SNR *per unit time* than conventional imaging.

2.3.2 NYQUIST GHOSTING

In EPI, alternate lines of k -space along the frequency-encoding axis are acquired by traversing k -space in opposite directions, and are therefore reversed with respect to time. This unique feature makes EPI particularly prone to an artefact commonly known as Nyquist or N/2 ghosting, which is due to inconsistencies in the signal between alternate lines of k -space. This introduces an extra frequency component at exactly half the sampling frequency, such that a ghost of the image is formed that is shifted from the original image by half the field of view along the phase-encoding axis (see figure 2.7). This kind of inconsistency can arise either from imperfect synchronisation of the receiver with the k -space trajectory, or from a modulation of the phase. If the data acquisition is delayed

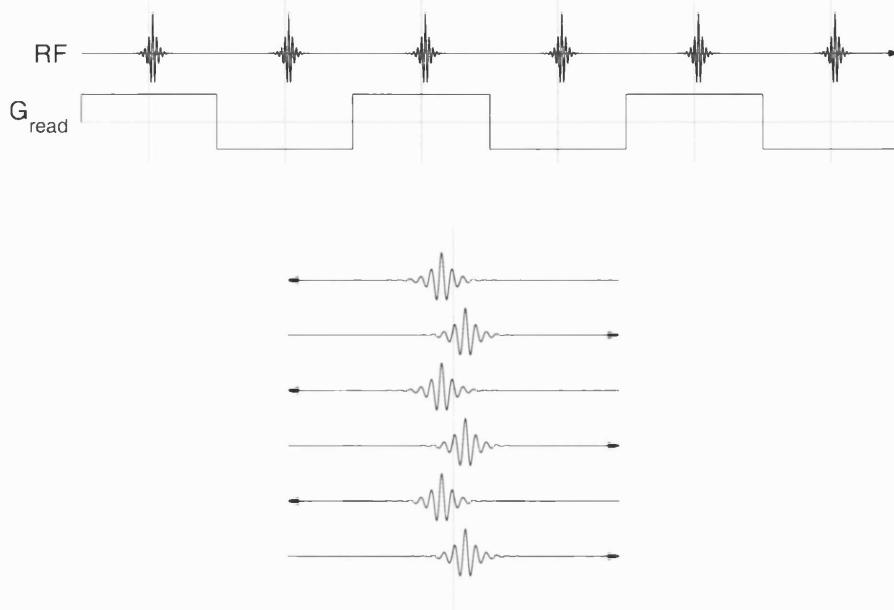


Figure 2.8: Echo asymmetry caused by the time-reversal of alternate lines of k -space. Top: schematic representation of the effect of a time delay between the gradient waveform and the data sampling. Bottom: after time-reversal of the even lines, the signal is modulated from line to line. The arrows indicate the time direction.

with respect to the gradient waveform, the centre of k -space for each line will be acquired later than expected. Since alternate lines are reversed with respect to time, this shift in time will translate to a shift in k -space that is reversed for alternate lines (figure 2.8). Phase modulation can be caused by eddy-currents, as described below.

EDDY-CURRENTS

When a magnetic field fluctuates, eddy-currents can be induced in any conducting structure exposed to the field. The fast gradient switching needed for EPI will therefore produce large eddy-currents that will oppose the intended change in the field. This will tend to introduce a delay between the application of a current through the gradient coil and the gradient field, as well as reduce the amplitude of the gradient field. For example, when using a sinusoidal gradient waveform, its frequency and shape will remain unchanged, but its phase and amplitude will be altered [14]. The resulting time delay between the application of the current in the coil and the gradient waveform can, if uncorrected, be a source of misalignment between the positions of the data samples in the time and k -space domains, and hence of Nyquist ghosting.

Eddy-currents may also contain a B_0 component that will alternate with the gradients. In this case, the resonant frequency will be different for alternate lines, resulting in a phase

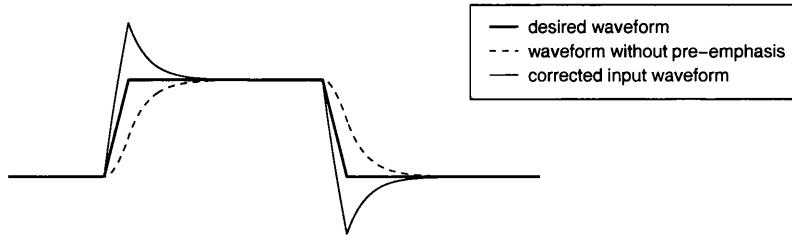


Figure 2.9: the desired gradient waveform (thick solid line) will be corrupted by unwanted eddy-currents, such that the sample experiences a different, delayed waveform (dashed line). With gradient pre-emphasis, the input waveform to the gradient amplifier is modified to alleviate these effects (thin solid line).

difference between these lines, and the appearance of a Nyquist ghost.

In order to reduce the magnitude of the eddy-currents induced by the gradients, a new gradient coil design was developed, that significantly reduced the magnitude of the gradient field induced outside the coil. These so-called *shielded gradients* consist of a second coil wound outside the primary gradient coil, such that the field generated by the first coil is exactly cancelled by the second coil. Although this results in reduced coil efficiency and increased inductance, this type of coil design has now become commonplace.

A technique known as *gradient pre-emphasis* is often used to alleviate the effect of eddy-currents on the main gradient field. The unwanted gradients induced by the eddy-currents, themselves produced by the applied gradient switching, can be estimated from the input gradient waveform using linear system theory [15]. Gradient pre-emphasis involves modifying the gradient waveform fed into the gradient coils in an effort to cancel the effects of unwanted gradients. For a typical trapezoidal gradient, this will in general involve ramping up more steeply, then overshooting and gradually declining to the intended gradient magnitude (see figure 2.9). Although pre-emphasis significantly reduces errors in the gradients due to eddy-currents, other approaches will be necessary to eliminate the remaining small timing errors between the gradients and the receiver.

RECEIVER

To prevent noise associated with frequencies higher than the receiver sampling rate from being aliased back and degrading the SNR of the signal, the receiver system must be equipped with a low pass filter. However, such a system will introduce a time delay between the incoming raw signal and the outgoing filtered signal (see figure 2.10). Oscillations will also be introduced in the signal following sudden intensity changes, but these fortunately have only minor consequences [14]. The delay introduced by the low pass filter is determined by the bandwidth of the receiver, and can thus be corrected straight-

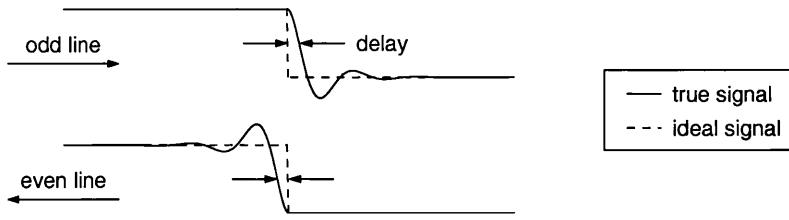


Figure 2.10: signal degradation due to the low-pass filter of the receiver system. The filter effectively convolves the incoming signal with a smoothing kernel, introducing both delay and ringing into the measured signal. The signal oscillations have only minor consequences, but if uncorrected, the delay will introduce Nyquist ghosting in EPI images due to the time-reversal of alternate lines.

forwardly.

The approach most often used to estimate and correct any delays between the gradients and the sampling involves the acquisition of a nyquist ghost correction scan. The central line of k -space is acquired repeatedly, using the same frequency-encoding gradient waveform as for the main EPI readout, but without any phase-encoding gradients. In the absence of any timing errors, all the resulting lines should be identical (in the absence of relaxation or off-resonance effects). If there are timing errors however, the set of lines acquired in one direction will be shifted in the opposite direction to the set of lines acquired in the other direction. The delay can then be estimated, and the receiver adjusted accordingly. This process is repeated until the delay falls within tolerance limits.

This adjustment is essential when non-linear time sampling is to be used, for example for blipped EPI. Without this correction, the data sampling rate would not be linear in k -space. This method is however not sufficient to entirely eliminate the Nyquist ghost artefact. The final adjustment scan can however be stored and later used to correct the remaining phase errors and any residual timing errors in post-processing, using the method outlined below.

SOFTWARE CORRECTION

As mentioned above, repeated acquisition of the central line of k -space provides a means to estimate any timing errors or phase inconsistencies. In blipped EPI, the last scan acquired to adjust the receiver delay can also be used to correct all the images acquired subsequently. In CPE-EPI however, receiver delay adjustments are typically not performed, and it is common to acquire a Nyquist ghost correction scan for each EPI readout. Such a correction scan is usually acquired by scanning the central line of k -space 2 or more times, either after a low angle pulse immediately before the main sequence, immediately after the main excitation pulse, or after the EPI readout [16].

Although acquiring the central line of k -space once in both directions is in theory sufficient, it has been found in practice that applying the correction in the presence of off-resonance effects may re-introduce Nyquist ghosting [17]. A frequency offset will induce a steady phase evolution from line to line, which will wrongly be interpreted by the correction algorithm as a phase error to be eliminated. Acquiring an extra line in the correction scan means that the first and third lines can be averaged, cancelling out any phase evolution that may have occurred during that time and allowing the correction algorithm to correctly interpret any remaining phase differences [17].

2.3.3 IMAGE DISTORTIONS

It is common for images acquired using EPI to be distorted in some way, either locally or globally. Distortions are generally caused by imperfections in the magnetic field, such that the linear relationship between spatial location and frequency (equation 2.1) is no longer exact. Due to the low bandwidth per pixel along the phase-encoding axis, even a small difference in the field may shift the signal from the affected region by a few pixels along that axis. Imperfections in the field may arise from a number of causes, and the most common are outlined here.

EDDY-CURRENTS

As mentioned above, eddy-currents induced by fast gradient switching will often generate additional time-dependent gradients in the field, as well as a net B_0 component. Although gradient pre-emphasis (see section 2.3.2) can significantly alleviate the effect of these unwanted contributions, gradients due to long term eddy-currents may still be present during the EPI readout. The effect of these eddy-currents will depend on the orientation and magnitude of the gradient they induce, and the magnitude of their B_0 component [18]. In the following description, it is assumed that remaining eddy-current-induced gradients are constant during the EPI readout, and that frequency-encoding is applied along x , phase-encoding along y , and slice selection along z . The main field B_0 is also assumed to be aligned with the z -axis.

A B_0 shift will induce a global shift in frequency. Due to the low bandwidth per pixel along y , this will cause the image to be shifted uniformly along that axis. Similarly, a gradient along the z -axis can be viewed as a slice position dependent B_0 shift, such that the image will be shifted along y by an amount directly dependent on the slice position along z .

A background gradient along x will interfere with the frequency-encoding gradient, causing a gradient pulse applied in one direction to overshoot in k -space, and the op-

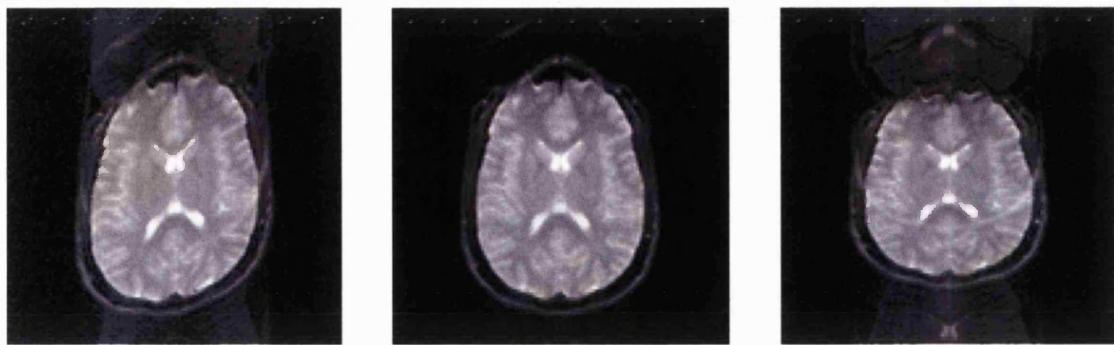


Figure 2.11: eddy-current-induced image distortions. Middle: an undistorted EPI image. Left: a sheared EPI image acquired in the presence of a small background gradient along the read-out direction. Right: a scaled EPI image acquired in the presence of a small background gradient along the phase-encoding axis.

posite gradient to undershoot. Thus, the position of the centre of k -space will steadily accumulate a small offset along x , resulting in an image that is sheared, as illustrated in figure 2.11.

A background gradient along y will alter the magnitude of the phase-encoding gradient. This will change the k -space sampling interval along y , Δk_y , such that the field of view along y is different. The image thus appears stretched or compressed along y , depending on the sign of the gradient, as illustrated in figure 2.11. This is commonly referred to as image scaling.

Eddy-current-induced image distortions can be minimised by techniques such as active shielding and gradient pre-emphasis, as described in section 2.3.2. However, the large prolonged gradients required for diffusion imaging tend to produce large eddy-currents that cannot be fully compensated for using these methods. The effects of eddy-currents can be further alleviated by careful sequence design and image processing techniques, described in more detail in section 3.2.3.

SUSCEPTIBILITY ARTEFACTS

At the interface between two substances with different magnetic susceptibilities, there will be a local discontinuity in the field, which will induce local field gradients. Spins in such areas will therefore resonate at a different frequency, such that their position in image space is shifted along the phase-encoding axis. The magnitude of that shift is given by $\Delta y = \Delta B_0/G_{pe}$, where G_{pe} is the magnitude of the phase-encoding gradient. This gives rise to strong image distortions, especially at air-tissue and bone-tissue interfaces, such as the frontal sinus (figure 2.12).

The local field gradients can also contribute to the phase-encoding gradient, and alter

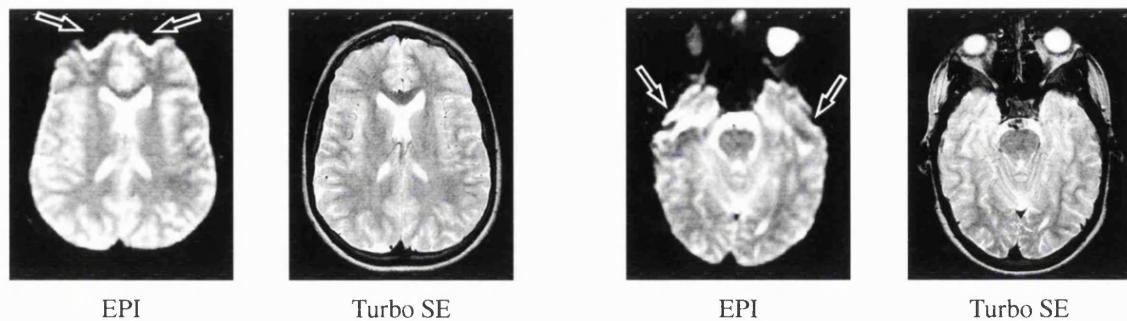


Figure 2.12: image distortion in EPI ('susceptibility artefact'). The figure shows typical echo-planar images (128x128 matrix) with their corresponding Turbo spin-echo images. Image distortions can be observed in the images acquired using EPI (arrows), as compared to the Turbo spin-echo images. There is also a redistribution of the signal, with regions of increased signal intensity, and regions of reduced signal. (Adapted from F. Calamante's PhD thesis, 2000)

its magnitude locally, resulting in a scaling of the affected area. Where the gradient is decreased, the signal that would otherwise originate from several pixels may be compressed into a single pixel, resulting in a dramatic increase in its signal intensity. Conversely, where the gradient is increased, the signal that would otherwise originate from a single pixel may be stretched over several pixels, resulting in a drop in signal intensity.

A number of methods have been proposed to correct susceptibility-induced image distortions [19]. Two images can be acquired to estimate the pixel shift, one with the phase-encoding applied along the x -axis, the other with the phase-encoding applied along the y -axis [20]. A similar approach is to acquire two images with the phase-encoding applied in opposite directions [21]. Another approach involves the acquisition of a field map, from which the pixel shift can be directly estimated [22]. A more recently proposed technique involves measuring the point spread function for each voxel, and deconvolving it from the measured signal [23]. All these approaches require the acquisition of additional data to estimate the pixel shift, and are therefore used primarily either to image regions of the brain that are particularly prone to susceptibility-induced artefacts, or in applications such as functional MRI, where the same correction can be applied to a series of images.

2.3.4 CHEMICAL SHIFT ARTEFACTS

Spins that possess a large chemical shift, and therefore resonate at a slightly different frequency, will appear shifted along the phase-encoding axis in EPI images. This is again due to the low bandwidth per pixel in that direction. At 1.5 T, the lipid signal has a chemical shift with respect to water of around 200 Hz [24], while the bandwidth for a

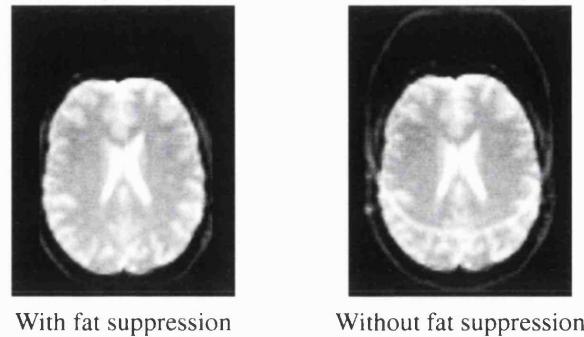


Figure 2.13: fat artefact in EPI. The figure shows an example of an echo-planar image acquired with (left) and without (right) fat suppression. In the absence of fat suppression, the lipid signal is visible and can be seen to be displaced along the phase-encoding direction (vertically) due to its different chemical shift. (Adapted from F. Calamante's PhD thesis, 2000)

typical EPI acquisition is of the order of 10 Hz per pixel in the phase-encoded direction. Subcutaneous fat is therefore often seen as a ring of signal shifted by approximately 20 pixels from its correct anatomical position in the image (see figure 2.13). Since the brain itself does not contain any MR visible lipids, the signal from fat is of little interest, and the subcutaneous fat signal must be suppressed. This is done by exploiting the large chemical shift between the lipid and water signals, either to selectively excite the lipid signal before spoiling it, or to selectively excite the water signal [12].

2.3.5 LIMITS ON IMAGE RESOLUTION

T_2^* DECAY

The time during which data can be acquired is limited by T_2^* decay. Since the speed of traversal of k -space is limited by the maximum gradient amplitude and switching rate, there is a limit to the extent of k -space that can be adequately covered. For this reason it is relatively rare to find EPI images acquired at resolutions higher than 128×128 . Although it is possible to acquire images with larger matrix sizes, their quality may be considerably degraded. In particular, geometric distortions will be accentuated, blurring will be increased, and the SNR will drop significantly.

T_2^* BLURRING

During the course of the acquisition, the signal decay due to T_2^* along the frequency-encoding direction is negligible, whereas that along the phase-encoding direction is considerable. This T_2^* modulation of the signal has profound consequences for the point

spread function (PSF).

The PSF describes the appearance in the final image of a single point in the sample. Consider a one-dimensional sample that consists of a single Dirac delta function, $S(x) = \delta(x)$. Its k -space representation is given by its Fourier transform, $S(k) = 1$. However, the extent of k -space sampled during a MRI experiment is limited, such that the signal is effectively weighted by a top hat function that sets the signal outside the acquisition window to zero. The k -space representation of a single point in the sample then becomes:

$$S_{acq}(k) = 1 \times \Pi(k/k_{max}) \quad \text{where } \Pi(w) = \begin{cases} 1 & \text{for } |w| < 1 \\ 0 & \text{otherwise} \end{cases} \quad (2.13)$$

The contribution in the final image of a single point in the sample is given by the inverse Fourier transform of the previous expression, and corresponds to the PSF:

$$PSF(x) = \frac{k_{max}}{2\pi} \text{sinc}(k_{max} x) \quad (2.14)$$

This function can be used to derive the appearance in the image $R_{acq}(x)$ of an arbitrary sample $R(x)$, by using the convolution theorem [25]:

$$R_{acq}(x) = R(x) * PSF(x) \quad (2.15)$$

This treatment applies when the effects of relaxation can be ignored. In EPI however, data in k -space are also weighted by T_2^* decay, depending on the time of acquisition of each sample. For blipped or CPE EPI, each phase-encoded line in k -space will be modulated by the T_2^* envelope, denoted here by $M(k_y)$. In this case, the PSF along the phase-encoding axis is given by the convolution of the inverse Fourier transform of the T_2^* envelope, $M(y)$, with the original PSF:

$$PSF_{EPI}(y) = PSF(y) * M(y) \quad (2.16)$$

Since the effects of relaxation along the frequency-encoding axis can be ignored, the PSF along that direction is unaffected. The PSF in EPI is therefore degraded and widened along the phase-encoding direction by the T_2^* modulation of the signal during the acquisition, which causes blurring and a loss of resolution along that axis [14]. In other words, T_2^* effects tend to attenuate the high frequency components of the signal along the phase-encoding axis, leading to a loss of resolution along that axis in the image. Furthermore, if the characteristic width of $M(k_y)$ is reduced, as may happen for example if the acquisition is lengthened, this would correspond to a even wider PSF and introduce additional blurring.

REFERENCES

- [1] CALLAGHAN, P. T. *Introductory nuclear magnetic resonance*. In [26], 1991, chapter 2.
- [2] CALLAGHAN, P. T. *Principles of imaging*. In [26], 1991, chapter 1.
- [3] CALLAGHAN, P. T. *High resolution k-space imaging*. In [26], 1991, chapter 4.
- [4] BRACEWELL, R. N. *Groundwork*, third ed. In [27], 2000, chapter 2.
- [5] BRACEWELL, R. N. *The discrete Fourier transform and the FFT*, third ed. In [27], 2000, chapter 11.
- [6] BRACEWELL, R. N. *Sampling and series*, third ed. In [27], 2000, chapter 10.
- [7] CALLAGHAN, P. T. *The influence of magnetic field gradients*. In [26], 1991, chapter 3.
- [8] GADIAN, D. G. *Instrument design and operation*. In [28], 1995, chapter 7.
- [9] SCHMITT, F., STEHLING, M. K., AND TURNER, R. *The historical development of echo-planar magnetic resonance imaging*. In [29], 1995, chapter 1.
- [10] BOWTELL, R., AND SCHMITT, F. *Echo-planar imaging hardware*. In Schmitt et al. [29], 1995, chapter 3.
- [11] COHEN, M. *Theory of echo-planar imaging*. In Schmitt et al. [29], 1995, chapter 2.
- [12] WIELOPOLSKI, P. A., SCHMITT, F., AND STEHLING, M. K. *Echo-planar imaging pulse sequences*. In Schmitt et al. [29], 1995, chapter 4.
- [13] MARGOSIAN, P., SCHMITT, F., AND PURDY, D. Faster MR imaging : imaging with half the data. *Health Care Instrumentation 1* (1986), 195–197.
- [14] FISCHER, H., AND LADEBECK, R. *Echo-planar imaging image artefacts*. In Schmitt et al. [29], 1995, chapter 6.
- [15] JEHENSON, P., WESTPHAL, M., AND SCHUFF, N. Analytical method for the compensation of eddy-current effects induced by pulsed magnetic field gradients in NMR systems. *Journal of Magnetic Resonance 90* (1990), 264–278.
- [16] SCHMITT, F., AND WIELOPOLSKI, P. A. *Echo-planar image reconstruction*. In Schmitt et al. [29], 1995, chapter 5.
- [17] HEID, O. Robust EPI phase correction. In *Proceedings of the ISMRM* (1997), vol. 5, p. 2014.
- [18] HASELGROVE, J. C., AND MOORE, J. R. Correction for distortion of echo-planar images used to calculate the apparent diffusion coefficient. *Magnetic Resonance in Medicine 36* (1996), 960–964.
- [19] ZENG, H., AND CONSTABLE, R. T. Image distortion correction in epi: comparison of field mapping with point spread function mapping. *Magnetic Resonance in Medicine 48* (2002), 137–146.
- [20] CORDES, D., ARFANAKIS, K., HAUGHTON, V., AND MEYERAND, M. E. Geometric distortion correction in EPI using two images with orthogonal phase-encoding directions. In *Proceedings of the ISMRM* (2000), vol. 8, p. 1712.
- [21] CHIOU, J. Y., AND NALCIOGLU, O. A simple method to correct off-resonance related distortion in echo planar imaging. In *Proceedings of the ISMRM* (2000), vol. 8, p. 1711.
- [22] JEZZARD, P., AND BALABAN, R. Correction for geometric distortion in echo planar images from B_0 field variations. *Magnetic Resonance in Medicine 34* (1995), 65–73.
- [23] ROBSON, M. D., GORE, J. C., AND CONSTABLE, R. T. Measurement of the point spread function in MRI using constant time imaging. *Magnetic Resonance in Medicine 38* (1997), 733–740.
- [24] GADIAN, D. G. *Pulse sequences*. In [28], 1995, chapter 8.
- [25] BRACEWELL, R. N. *The basic theorems*, third ed. In [27], 2000, chapter 6.

- [26] CALLAGHAN, P. T. *Principles of nuclear magnetic resonance microscopy*. Oxford University Press, Oxford, UK, 1991.
- [27] BRACEWELL, R. N. *The Fourier transform and its applications*, third ed. McGraw-Hill International Editions, Singapore, 2000.
- [28] GADIAN, D. G. *NMR and its applications to living systems*. Oxford University Press, Oxford, UK, 1995.
- [29] SCHMITT, F., STEHLING, M. K., AND TURNER, R., Eds. *Echo-planar imaging: theory, technique and application*. Springer-Verlag, 1995.

CHAPTER 3

DIFFUSION IN MRI

CONTENTS

3.1	Introduction to diffusion	38
3.1.1	Molecular diffusion	38
3.1.2	The effect of diffusion on the NMR signal	39
3.1.3	Diffusion in biological systems	42
3.2	Diffusion-weighted imaging	45
3.2.1	The b -matrix	45
3.2.2	Diffusion-weighted sequences	49
3.2.3	Artefacts	52
3.3	Diffusion tensor imaging	56
3.3.1	The diffusion tensor	56
3.3.2	Diffusion encoding schemes	63
3.3.3	Tractography	64
3.3.4	Artefacts	68
3.3.5	Alternative models for diffusion	70
3.3.6	Unresolved issues in tractography	72
References		72

3.1 INTRODUCTION TO DIFFUSION

3.1.1 MOLECULAR DIFFUSION

Diffusion refers to the thermally driven random motion of particles, commonly known as Brownian motion. It is classically described by Fick's law [1], in terms of a macroscopic flux density \underline{J} driven by a concentration gradient ∇C :

$$\underline{J} = -D \nabla C \tag{3.1}$$

where C is the concentration and D is the diffusion coefficient, typically measured in mm^2/s . Conservation of mass [2] dictates that:

$$\frac{\partial C}{\partial t} = -\nabla \cdot \underline{J} \quad (3.2)$$

which can be combined with Fick's law to give the diffusion equation:

$$\frac{\partial C}{\partial t} = D \nabla^2 C \quad (3.3)$$

This equation can be solved for various simple boundary conditions. In particular, the solution for particles in an unbound one-dimensional medium with uniform diffusivity, satisfying the initial condition $C(x, 0) = \delta(x)$ is¹:

$$C(x, t) = \frac{1}{\sqrt{4\pi Dt}} \exp\left(\frac{-x^2}{4Dt}\right) \quad (3.4)$$

Einstein [3] showed that the diffusion coefficient relates the variance in the spin displacement to the diffusion time. The variance $\langle x^2 \rangle$ in the spin displacement can be derived for the one-dimensional free diffusion case from equation 3.4:

$$\langle x^2 \rangle = 2Dt \quad (3.5)$$

This can also be expressed in terms of the conditional probability distribution $P(x_t | x_0, t)$, the probability of finding a molecule at position x_t after a time t , given that it was originally at position x_0 . In this case, $x = x_t - x_0$ and $\langle x^2 \rangle$ corresponds to the variance of $P(x_t | x_0, t)$.

Equation 3.4 suggests that the diffusivity D can be inferred by monitoring the concentration profile over time using tracers. Such techniques have been applied successfully in biological systems, such as the brain [4, 5]. However, because of the invasive nature of exogenous tracers, these cannot be used *in-vivo* with humans. Nevertheless, NMR does allow non-invasive diffusion measurements to be performed *in-vivo* by using water as an endogenous tracer, as described below.

3.1.2 THE EFFECT OF DIFFUSION ON THE NMR SIGNAL

The effect of diffusion on the NMR signal can be understood from a simple pulsed gradient spin-echo (PGSE) experiment [6] (figure 3.1). We define G as the gradient strength, δ as the gradient duration, and Δ as the time interval between the onset of the two gradient

¹The equivalent expression for the three-dimensional case is: $C(\underline{x}, t) = \left(\frac{1}{\sqrt{4\pi Dt}}\right)^3 \exp\left(\frac{-\underline{x}^2}{4Dt}\right)$

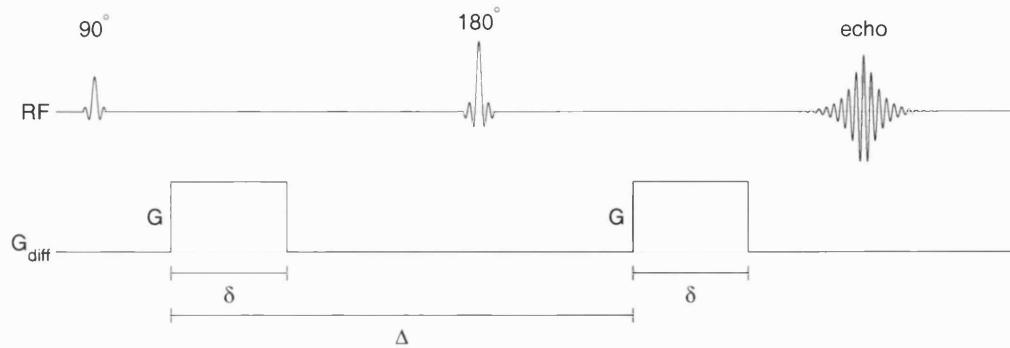


Figure 3.1: a schematic representation of the Stejskal-Tanner pulsed-gradient spin-echo (PGSE) diffusion-weighted pulse sequence. The first diffusion-encoding gradient pulse induces a position-dependent phase shift that is undone by the second gradient pulse. Static spins will therefore experience no net phase shift. However, for spins that moved along the direction of the diffusion-encoding gradient between the application of the two pulses, the phase shift induced by the first gradient pulse will not be cancelled out exactly by the second pulse. An ensemble of diffusing spins will therefore carry a distribution of phase shifts, leading to signal loss.

pulses. It is assumed that $\delta \ll \Delta$, such that the motion of the spins during the short duration of the gradient pulse is negligible. This assumption is known as the *narrow pulse approximation* [7]. The first gradient induces a phase shift ϕ_1 that depends on the spin position (see equation 2.5). Without loss of generality, we can assume that the gradient pulses are applied along the z -axis. In this case:

$$\phi_1 = \gamma \int_0^\delta \underline{G}(t) \cdot \underline{x} dt = \gamma G \delta z_1 \quad (3.6)$$

where γ is the magnetogyric ratio and z_1 is the spin position along the z -axis. The 180° pulse reverses ϕ_1 to $-\phi_1$. The second pulse then produces a phase shift ϕ_2 :

$$\phi_2 = \gamma \int_\Delta^{\Delta+\delta} \underline{G}(t) \cdot \underline{x} dt = \gamma G \delta z_2 \quad (3.7)$$

where z_2 is the spin position during the second pulse. The net dephasing $\Delta\phi$ is therefore:

$$\Delta\phi = \phi_2 - \phi_1 = \gamma G \delta (z_1 - z_2) \quad (3.8)$$

For static spins, $z_1 = z_2$ and the second pulse simply refocuses the effect of the first pulse. However, a system of moving spins will experience a net dephasing that will depend on their motion during the time interval Δ . The total magnetisation is given by the

vector sum of the magnetic moments of the individual nuclei. Using complex notation:

$$M = M_0 \sum_j e^{i\Delta\phi_j} \quad (3.9)$$

where j refers to each individual spin within the voxel, and M_0 is the total magnetisation of the voxel in the absence of dephasing. For bulk or *coherent* motion such as flow, the phase of each spin will be shifted by the same amount, leaving the magnitude of the magnetisation unchanged, but altering its phase. For incoherent motion however, there will be a distribution of phase shifts, and the resulting destructive interference will lead to a reduction in the magnitude of the signal. It is this signal drop that forms the basis of diffusion-weighted imaging.

The phase shift $\Delta\phi_j$ of each individual spin depends on its shift in position between the two gradient pulses. The motion of the diffusing spins can be described by $P(z_2|z_1, \Delta)$, the conditional probability of finding a spin at position z_2 after a time Δ , given that it was originally at position z_1 . Therefore, for a sample containing a large number of spins, equation 3.9 can be rewritten as:

$$\frac{M}{M_0} = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} P(z_2 | z_1, \Delta) \exp[i\gamma G \delta(z_1 - z_2)] dz_1 dz_2 \quad (3.10)$$

We can see from equation 3.4 that for *free diffusion* along the z -axis, the conditional probability is given by:

$$P(z_2 | z_1, \Delta) = \frac{1}{\sqrt{4\pi D\Delta}} \exp\left(\frac{-(z_1 - z_2)^2}{4D\Delta}\right) \quad (3.11)$$

Equations 3.10 and 3.11 can then be combined and simplified to give:

$$M = M_0 \exp[-\gamma^2 G^2 \delta^2 D \Delta] \quad (3.12)$$

The effects of gradient strength and timing are often conveniently grouped together in the so-called b -factor [8], and equation 3.12 can be further simplified to:

$$M = M_0 e^{-bD} \quad \text{where} \quad b = \gamma^2 G^2 \delta^2 \Delta \quad (3.13)$$

Equation 3.13 relates the measured signal attenuation to the diffusivity, and is the basis for the measurement of diffusion using NMR. However, the narrow-pulse approximation is not in general valid, due to limited gradient strengths. In this case, diffusion occurring during the application of the gradients needs to be taken into account (see section 3.2.1). Additional gradients, including imaging pulses and background residual gradients, may

also have a significant effect on the measured values (see section 3.2.1). Moreover, the diffusion process may be hindered or restricted, as discussed in the next section.

3.1.3 DIFFUSION IN BIOLOGICAL SYSTEMS

Biological tissues possess complex, heterogeneous microstructure, where the transport properties of spins can only be approximated by the classical free diffusion approach [9]. In particular, the presence of impermeable or semi-permeable barriers to diffusion, exchange between the various compartments, and the orientation dependence of the diffusion process all need to be taken into account. As a result, the conditional probability distribution function mentioned above (equation 3.4) is predicted to be non-Gaussian, and the measured diffusion coefficient will be dependent on the experimental conditions [10]. Consequently, the diffusion parameter measured in biological systems is termed the *apparent* diffusion coefficient (ADC).

MULTIPLE COMPARTMENTS

Tissue water is contained within various compartments, notably the intra/extracellular environments, and is in a state of continuous exchange between these compartments [11]. The observed signal attenuation will depend on the diffusion coefficient in each compartment, the rate of exchange between these compartments, their relative volume fractions, and specific imaging parameters, notably the diffusion time [12].

The treatment of diffusion in such systems is simplified by considering the two extreme cases for the rate of exchange. In the *slow exchange* regime, spins can be assumed to remain within their original compartments during the course of the experiment. In other words, the average diffusion distance is much smaller than the average compartment size. In this case, the signals from each compartment simply add together, such that the total signal decays multi-exponentially [11]. For example, in a two-compartment system, the signal is given by:

$$M(b) = M_0(f_1 e^{-b\text{ADC}_1} + f_2 e^{-b\text{ADC}_2}) \quad (3.14)$$

where f_1 and f_2 are the volume fractions for each compartment (such that $f_1 + f_2 = 1$) and ADC_1 and ADC_2 are the corresponding apparent diffusion coefficients.

On the other hand, in the *fast exchange* regime, the spins will be completely randomised between the compartments during the course of the experiment. In other words, the average diffusion distance is much greater than the average compartment size. The observed ADC is then averaged over all the compartments, and the signal will decay mono-exponentially. For the two-compartment system introduced above, the observed

ADC would be:

$$\text{ADC} = f_1 \text{ADC}_1 + f_2 \text{ADC}_2 \quad (3.15)$$

However, diffusion in biological tissues is more likely to occur in an intermediate regime rather than either the fast or slow exchange limits.

In addition, the different T_2 relaxation times within each compartment [13, 14] need to be taken into account. The intracellular T_2 is expected to be different from the extracellular T_2 due to the different water mobilities in the two compartments [11], introducing a dependence of the effective volume fractions on the echo time.

RESTRICTED OR HINDERED DIFFUSION

The presence of barriers to the free movements of molecules within the sample will introduce a dependence of the measured ADC on the diffusion time [15]. When the diffusion time is short, the distance over which the molecules will have diffused will also be short, such that they are unlikely to encounter the cell membrane (see figure 3.2). In this case, the molecules behave as though they were diffusing freely, and the presence of boundaries has no detectable effect on the measured diffusion coefficient. However, if longer diffusion times are used, a larger proportion of the molecules will strike the boundaries, and their average displacement will be lower than expected from the free diffusion model, leading to an underestimation of the diffusion coefficient.

In the simple case of a fully restricted system with impermeable membranes, the mean square displacement $\langle r^2 \rangle$ will deviate from its expected linear relationship with diffusion time (equation 3.5), and level off as the diffusion distance reaches the size of the restricting compartment [16]. The effect of restriction on the ADC will thus become significant as the diffusion distance reaches the characteristic length of the restricted compartment. If the membranes are permeable, the mean square displacement will increase at a reduced rate with respect to the freely diffusing case, causing the ADC to decrease towards an asymptotic value, dependent on the permeability of the membrane.

The effect of restricted or hindered diffusion on the ADC will depend on the shape and size of the restricted compartment, the permeability of the barriers, and the type of NMR experiment performed. There is thus no simple way to interpret ADC measurements obtained on biological tissues. However, this effect may be used to determine compartmental geometry, by measuring the signal attenuation as a function of the diffusion weighting [17]. This technique, commonly known as q -space imaging, is unfortunately very time-consuming and difficult to undertake in practice due to its hardware demands (see section 3.3.5).

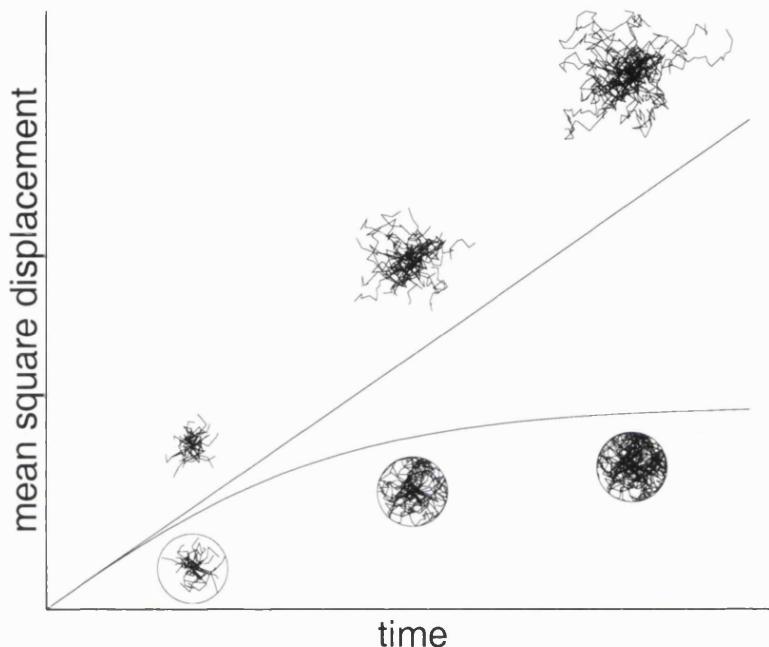


Figure 3.2: free and restricted diffusion. With free diffusion, for example water molecules in a bottle, the mean square displacement, $\langle r^2 \rangle$, increases linearly with time, according to Einstein's equation (equation 3.5). The slope then gives the diffusion coefficient, D . If diffusion molecules are restricted to a closed space such as an impermeable cell, the apparent diffusion coefficient measured by MRI depends strongly on the diffusion time. For short diffusion times, most molecules do not have time to reach the boundary, and the diffusion behaves as if it were essentially free. As the diffusion time increases, more molecules are reflected back at the boundary, and the average displacement reaches an asymptote, such that the measured ‘apparent’ diffusion coefficient decreases to zero. Diffusion in the presence of semi-permeable barriers lies between these two extremes: the mean square displacement will increase non-linearly as a function of time, and the apparent diffusion coefficient will decrease to an asymptotic value.

DIFFUSION ANISOTROPY

In a 3-dimensional system, the mobility of the water molecules may be orientationally dependent. Such a situation may occur in the presence of coherently oriented obstructions to diffusion, as in skeletal muscle [18] and brain white matter [19]. In these tissues, water mobility is highest along the orientation of the fibres, due to the higher degree of hindrance encountered by water molecules when moving across the fibres, rather than along their axis. This behaviour is known as *anisotropic* diffusion. The ADC measured then depends on the orientation of the sample with respect to the direction of the diffusion measurement. However, structures that exhibit anisotropic diffusion at the microscopic level may appear isotropic at the macroscopic level if the orientational coherence of the tissue only applies over short (microscopic) distances, resulting in a ‘powder average’ effect that is difficult to interpret (see section 3.3.4).

Since a single ADC value cannot adequately describe such an anisotropic system, the concept of an *apparent diffusion tensor*, \mathbf{D} , often referred to simply as the diffusion tensor, was introduced [20]. Equation 3.13 can be rewritten using this tensor formalism [21]:

$$M = M_0 e^{-\mathbf{b} \cdot \mathbf{D}} \quad (3.16)$$

where \mathbf{b} is the b -matrix, describing the direction and strength of the diffusion weighting used in the measurement, and the *tensor dot product* is introduced, defined as:

$$\mathbf{A} : \mathbf{B} = \sum_{i=1}^3 \sum_{j=1}^3 A_{ij} B_{ij} \quad (3.17)$$

A more detailed description of the b -matrix is given in section 3.2.1. For a full description of diffusion in anisotropic systems, the reader is referred to section 3.3.1.

3.2 DIFFUSION-WEIGHTED IMAGING

3.2.1 THE b -MATRIX

The Bloch equations [22], modified for diffusion [23], can be used to calculate the signal attenuation for a general gradient pulse sequence. The general expression for the spin-echo signal attenuation at the echo time TE in an anisotropic medium is then given by:

$$\ln \left[\frac{M(\text{TE})}{M(0)} \right] = - \int_0^{\text{TE}} \underline{k}(t)^T \mathbf{D} \underline{k}(t) dt \quad (3.18)$$

where \mathbf{D} is the apparent diffusion tensor and $\underline{k}(t)$ is the k -space vector as introduced in equation 2.6, namely:

$$\underline{k}(t) = \gamma \int_0^t \underline{G}(t') dt' \quad (3.19)$$

Note that the k -space vector is reversed after the application of a 180° pulse (see section 2.1.2). This must be taken into account, for example by reversing the sign of all gradients following a refocusing pulse.

It can then be shown that:

$$\underline{k}(t)^T \mathbf{D} \underline{k}(t) = [\underline{k}(t) \cdot \underline{k}(t)^T] : \mathbf{D} \quad (3.20)$$

where the colon is used to denote the tensor dot product, as defined in equation 3.17. Since the diffusion tensor \mathbf{D} is assumed to be time-independent, the b -matrix introduced in equation 3.16 can then be written:

$$\mathbf{b} = \int_0^{TE} \underline{k}(t) \cdot \underline{k}(t)^T dt \quad (3.21)$$

The diffusion-encoding gradients are commonly much larger and longer than the imaging gradients, and the contribution of the latter to the b -matrix is often ignored. However, it has been shown that the effect of these imaging gradients may not be negligible, and may introduce significant errors in diffusion measurements. These effects and ways to alleviate them are discussed in more detail in the next sub-section.

It is also interesting to note that most diffusion-weighted imaging sequences apply the diffusion-encoding gradients at their maximum amplitude along a single direction per excitation (with the exception of trace-weighted sequences [24, 25]). Thus, assuming that imaging gradients are negligible, and that the diffusion gradients are applied along a single direction, the k -space vector can be written as:

$$\underline{k}(t) = \gamma G \left[\int_0^t m(t') dt' \right] \hat{\underline{u}} \quad (3.22)$$

where G is the peak amplitude of the diffusion gradient, $\hat{\underline{u}}$ is the unit vector pointing along the direction of application of the diffusion gradient, and $m(t)$ is a function describing the gradient amplitude modulation. Using this notation, the b -matrix can be written:

$$\mathbf{b} = \gamma^2 G^2 F_{seq} [\hat{\underline{u}} \cdot \hat{\underline{u}}^T] \quad \text{where} \quad F_{seq} = \int_0^{TE} \left[\int_0^t m(t') dt' \right]^2 dt \quad (3.23)$$

Using this formulation, all the directionality information is contained within the $\hat{\underline{u}} \cdot \hat{\underline{u}}^T$ term, and the diffusion sensitivity information, or b -value, is given as $\gamma^2 G^2 F_{seq}$. Thus,

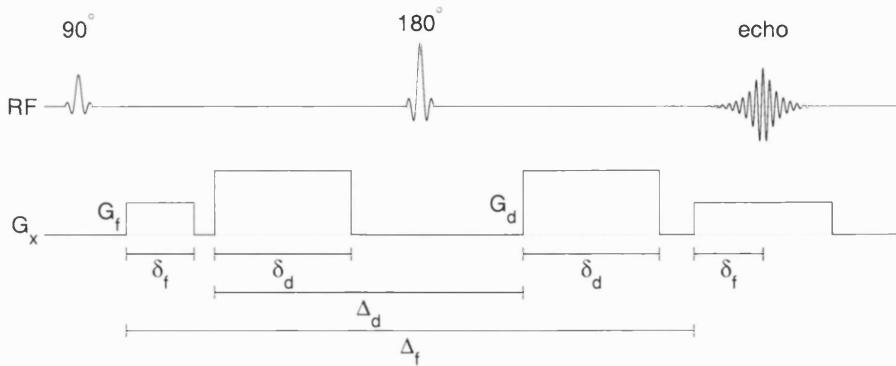


Figure 3.3: a schematic pulsed-gradient spin-echo pulse sequence that does not minimise cross-terms. The read-dephasing gradient pulse is applied before the first diffusion-encoding gradient pulse, leading to additional diffusion encoding arising from its interaction with the diffusion-encoding gradient pulses. This contribution can be significant and lead to overestimation of the apparent diffusion coefficient (see text for details).

when designing a single-direction diffusion-weighted sequence, it is only necessary to consider the effect of the gradient timings on F_{seq} . For example, in the case of the simple narrow-pulse bipolar sequence described in section 3.1.2, it can be shown that $F_{seq} = \delta^2 \Delta$, and that the b -value is therefore given by $\gamma^2 G^2 \delta^2 \Delta$. Note that in practice, the diffusion-encoding gradient pulse duration δ will not be negligible, and the narrow-pulse approximation cannot be assumed. See section 3.2.2 for a description of the more general pulsed gradient spin-echo (PGSE) sequence.

CONTRIBUTION FROM THE IMAGING GRADIENTS

Imaging gradients used in typical PGSE sequences will also contribute to the b -matrix. If these are not taken into account, the ADC measured may be significantly overestimated [26]. Consider a simple PGSE sequence such as that shown in figure 3.3. By solving equation 3.21 along the read direction (the x -axis), the corresponding b -matrix element can be shown to consist of three different contributions [8]:

$$b_{xx} = b_i + b_d + b_{ct} \quad (3.24)$$

where b_i is the contribution from the imaging gradients alone, b_d is the contribution from the diffusion gradients alone, and b_{ct} is the contribution arising from a *cross term* between the imaging and diffusion gradients:

$$\begin{aligned} b_i &= \gamma^2 G_f^2 \delta_f^2 (\Delta_f - \delta_f/3) \\ b_d &= \gamma^2 G_d^2 \delta_d^2 (\Delta_d - \delta_d/3) \end{aligned}$$

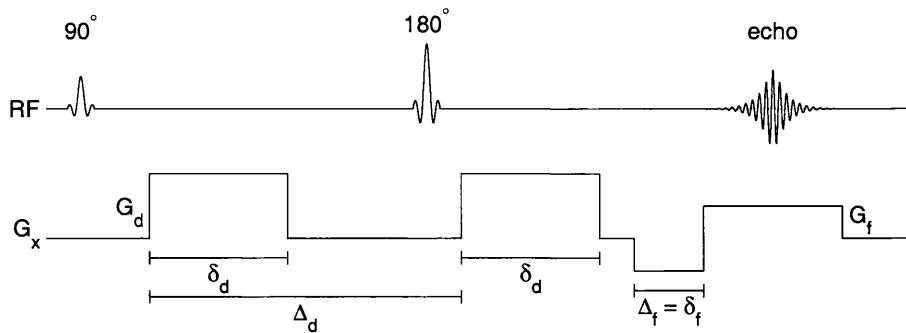


Figure 3.4: a schematic pulsed-gradient spin-echo sequence designed to minimise cross-terms. In contrast to the sequence shown in figure 3.3, the diffusion gradient pulses have been refocused prior to the application of the read-dephasing gradient pulse. There is thus no extra diffusion attenuation arising from interactions between the diffusion-encoding and imaging gradients. Although the imaging gradients do cause additional attenuation, this is independent of the diffusion gradients and can be absorbed into the M_0 term.

$$b_{ct} = 2\gamma^2 G_d G_f \delta_d \delta_f \Delta_d$$

where the parameters G_f , G_d , Δ_f , Δ_d , δ_f , and δ_d are as shown in figure 3.3.

The extra contributions from the imaging gradients and the cross terms arising from their interaction with the diffusion gradients may be significant. Consider the simple PGSE sequence shown in figure 3.3, with the following parameters: $G_d = 20$ mT/m, $\Delta_d/\delta_d = 45/20$ ms, $G_f = 10$ mT/m, $\Delta_f/\delta_f = 83/10$ ms. If only the diffusion pulses are taken into account, the calculated b_{xx} term is given by $b_d = 439$ s/mm². However, when all gradient pulses are accounted for, equation 3.24 gives $b_{xx} = 754$ s/mm². The contribution from the imaging gradients alone (in this case, $b_i = 57.0$ s/mm²) is independent of the diffusion gradients, and can therefore be absorbed into the M_0 term without affecting the measured ADC. However, the contribution from the cross-terms ($b_{ct} = 258$ s/mm²) is linear in G_d : it will be present only when the diffusion gradients are applied. If it is not taken into account, the resulting underestimation of the b -value will lead to an equivalent overestimation of the measured ADC.

It is possible to reduce the effects of imaging gradients on the b -matrix by careful pulse sequence design. Cross terms can be eliminated by ensuring that the diffusion gradient pulses are refocused before the application of the imaging gradients. In the example above, the read-dephasing gradient can be moved to a position immediately before the read-out pulse, as shown in figure 3.4. The b_{xx} term for that sequence would then be 444 s/mm², a value much closer to that desired.

Some imaging gradient pulses need to be applied during the diffusion encoding. Such gradients should be refocused as quickly as possible in order to minimise their effect. The

180° slice-selection pulse thus has little effect since it is intrinsically refocused.

The EPI read-out train also contributes to the b -matrix. Mattiello *et al* [27] have shown that since each read-out pulse is immediately refocused by the next pulse, their combined effect is negligible. However, the phase-encoding gradients are not immediately refocused and can have a measurable effect, particularly for low values of the diffusion-weighting.

3.2.2 DIFFUSION-WEIGHTED SEQUENCES

The measurement of diffusion will in general require the acquisition of at least two images with different diffusion weightings, such that the ADC can be estimated from a two point fit. It has been shown that for an optimal estimate, the difference between the two b -values should be of the order of 1/ADC, with the lower b -value being as close to zero as possible [28]. In most practical situations, a number of diffusion-weighted images will be acquired with a high b -value sequence, and at least one additional image will be acquired using an otherwise identical sequence, with little or no diffusion sensitivity (often referred to as the $b=0$ image). In this way, the diffusion-weighted images can be referenced to the $b=0$ image to create ADC maps that have no spin density or T_2 contribution.

A typical diffusion-weighted imaging sequence will be sensitive to the motion of spins on the order of microns, and images are therefore very prone to motion-related artefacts (see section 3.2.3). For this reason, single-shot techniques such as EPI are typically used for the read-out part of the sequence.

In general, the diffusion weighting will be applied in a number of directions to account for the effects of diffusion anisotropy. Average ADC or *trace* maps can be generated by acquiring 3 diffusion-encoded images with the gradient applied in orthogonal directions, in addition to the $b=0$ image [29]. To estimate the diffusion tensor, at least 6 diffusion-weighted images need to be acquired in addition to the $b=0$ image, and schemes suitable to acquire such data will be described in more detail in section 3.3. In this section, a number of pulse sequences suitable for the acquisition of diffusion-weighted images are described.

THE CONSTANT FIELD GRADIENT SPIN-ECHO SEQUENCE

Diffusion measurements were first performed by applying a constant diffusion-encoding gradient during a spin-echo experiment [2]. From equation 3.23, the b -value for this sequence can be shown to be $\gamma^2 G^2 TE^3 / 12$. Although this technique has allowed the very precise measurement of the diffusion coefficient in a variety of simple liquids, it is very rarely used in modern practice, due to the presence of the diffusion gradient during the application of the RF pulses and during signal collection [30].

THE PULSED GRADIENT SPIN-ECHO SEQUENCE

The pulsed gradient spin-echo (PGSE) sequence, first demonstrated by Stejskal and Tanner in 1965 [6], is the most common way of acquiring diffusion-weighted images. It consists of a pair of gradient pulses placed either side of the 180° refocusing pulse, as shown in figure 3.1. Assuming negligible gradient ramps, the gradient pattern, in the same form as equation 3.23, is given by:

$$m(t') = \begin{cases} 1 & 0 < t' < \delta \\ -1 & \Delta < t' < \Delta + \delta \\ 0 & \text{elsewhere} \end{cases} \quad (3.25)$$

Note that the sign of the second gradient pulse has been reversed to account for the 180° refocusing pulse. To obtain F_{seq} for this sequence (see equation 3.23), the integral of $m(t')$ needs to be evaluated:

$$\int_0^t m(t') dt' = \begin{cases} t & 0 < t < \delta \\ \delta & \delta < t < \Delta \\ -s & -\delta < s < 0 \end{cases} \quad (3.26)$$

where the variable $s = t - (\Delta + \delta)$ has been introduced in order to simplify the following steps. This then needs to be squared:

$$\left[\int_0^t m(t') dt' \right]^2 = \begin{cases} t^2 & 0 < t < \delta \\ \delta^2 & \delta < t < \Delta \\ s^2 & -\delta < s < 0 \end{cases} \quad (3.27)$$

and finally integrated a second time to give F_{seq} :

$$\begin{aligned} F_{seq} &= \int_0^{TE} \left[\int_0^t m(t') dt' \right]^2 dt \\ &= \left[\frac{t^3}{3} \right]_0^\delta + \left[\delta^2 t \right]_\delta^\Delta + \left[\frac{s^3}{3} \right]_{-\delta}^0 \\ &= \delta^2(\Delta - \delta/3) \end{aligned} \quad (3.28)$$

The b -value for this diffusion encoding is therefore $\gamma^2 G^2 \delta^2 (\Delta - \delta/3)$. From equation 3.13, we can see that this is the same value as would be obtained using a narrow-pulse experiment with a diffusion time of $\Delta - \delta/3$. The PGSE sequence therefore has an *effective* diffusion time of $\Delta - \delta/3$.

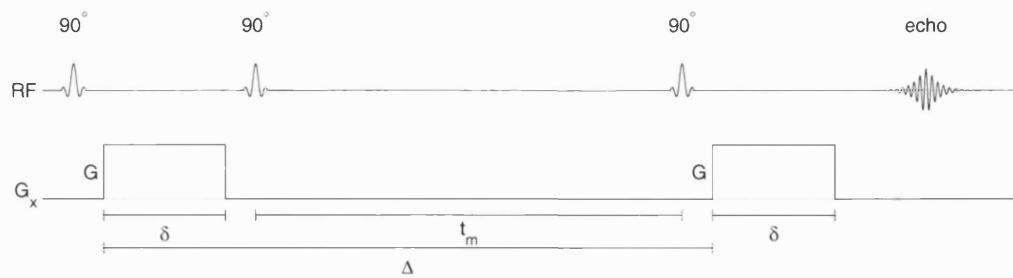


Figure 3.5: a schematic representation of the diffusion-weighted stimulated echo sequence. The diffusion encoding is applied in a similar way to the PGSE sequence (see figure 3.1). The main difference between the two sequences is that the magnetisation is stored along the z -axis between the application of the second and third 90° RF pulses. During this time, the signal decays with T_1 , which is usually much longer than T_2 , such that diffusion measurements at long diffusion times are possible. However, the stimulated echo sequence carries the inherent penalty that only half the signal is preserved.

THE STIMULATED ECHO SEQUENCE

Bipolar diffusion-encoding gradient pulses can easily be introduced into a stimulated echo sequence, as illustrated in figure 3.5. During the mixing time t_m , the signal decays with T_1 , which is longer than T_2 , such that longer diffusion times can be achieved than with the PGSE sequence. This allows large values of Δ , and hence large b -values to be achieved. However, half the signal is inherently lost using a stimulated echo sequence [31], such that it only becomes advantageous over the PGSE sequence for long echo times or in tissues with a short T_2 .

THE TWICE-REFOCUSED SEQUENCE

The use of large, rapidly switched, prolonged gradients in PGSE diffusion-weighted sequences makes them prone to eddy-current-induced artefacts (described in section 3.2.3). By introducing an extra 180° refocusing pulse, the long diffusion gradient pulses can each be split into two shorter pulses of opposite polarity, as illustrated in figure 3.6. This doubles the number of gradient amplitude transitions, and shortens the time lapse between successive transitions, such that eddy-currents set up by one gradient transition are more quickly countered by the next transition. Eddy-current-induced artefacts are thus significantly reduced, while the diffusion-encoding efficiency remains similar to that of the PGSE sequence [32]. This type of sequence can be further optimised by fine-tuning the duration of the gradient pulses, as described in detail in section 4.2.1.

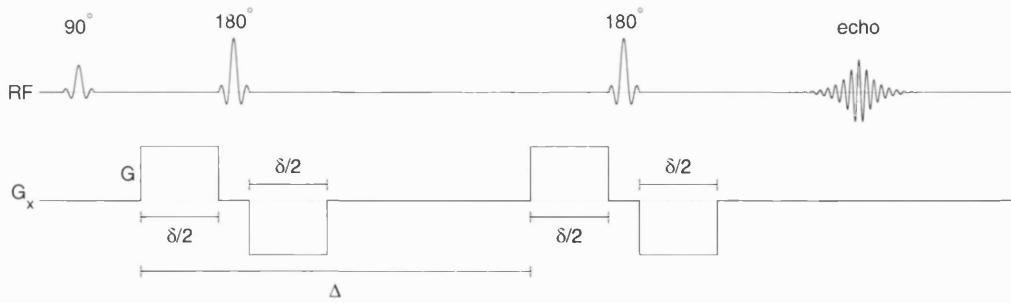


Figure 3.6: a schematic representation of the twice-refocused diffusion-weighted spin-echo sequence. In this sequence, each diffusion-encoding gradient pulse of a standard PGSE sequence (see figure 3.1) has been split into two pulses separated by a 180° refocusing pulse. The use of shorter pulses alleviates eddy-current-related problems while retaining a similar diffusion-encoding efficiency to the PGSE sequence.

TRACE-WEIGHTED SEQUENCES

It is possible to design a gradient pattern such that the resulting b -matrix contains identical diagonal elements, and zero off-diagonal elements (in other words, \mathbf{b} is a multiple of the identity matrix). The signal intensity measured using such a scheme would then be directly dependent on the trace of the diffusion tensor (see section 3.3.1). The simplest such scheme consists of applying a bipolar gradient pulse pair along each of three orthogonal axes in succession. This encoding unfortunately results in only a small degree of diffusion weighting. However, more complicated gradient patterns can be designed to achieve higher degrees of trace-weighted diffusion encoding, such as that proposed by Mori and van Zijl [24] and Wong *et al* [25].

3.2.3 ARTEFACTS

Diffusion-weighted images are typically acquired using an EPI read-out, and are therefore prone to artefacts specific to that technique, described in section 2.3. However, the extreme motion sensitivity of diffusion-weighted sequences introduces a number of further problems, described below. Also, since very large gradients are typically used to achieve the high b -values necessary, eddy-currents need to be taken into consideration.

BULK MOTION

The effect on the signal of bulk, or *coherent* motion arises from its effect on the phase of the individual spins. If a bulk translation occurs during the diffusion time, every spin in the sample will acquire the same additional phase shift, resulting in a phase shift in the overall signal. This process does not however introduce any extra dephasing, such that

the magnitude of the signal is unchanged. This is in contrast to the case of *incoherent* motion, whereby each spin will acquire a different phase shift, depending on its motion during the diffusion time. In this case, destructive interference will cause a reduction in the magnitude of the signal, which is the basis of diffusion imaging. Thus, using a single shot technique, a small bulk translation has no effect on the signal intensity in the image.

On the other hand, a small bulk rotation will affect the image intensity in single-shot diffusion-weighted imaging if the rotation axis and the direction of application of the diffusion gradient both have in-plane components. The effect of a bulk rotation is to introduce a linear phase shift in the sample along a direction perpendicular both to the axis of rotation and to the applied diffusion gradient, shifting the centre of k -space away from its initial position [33]. If this shift remains in the k_x - k_y plane, the effect on the image will be to introduce a linear phase shift that will leave the magnitude of the signal unaffected. However, if the axis of rotation and the diffusion gradient both have non-zero components in the imaging (x - y) plane, then the centre of k -space will be shifted along k_z away from the acquisition plane, resulting in signal loss.

The k -space phase and position errors introduced by bulk translation and rotation have a far more considerable effect when using multi-shot techniques, such as segmented EPI or 2D-FT: any inconsistency between the different segments of k -space will introduce severe ghosting in the image (see figure 3.7). To alleviate this problem, the relative shifts in phase and position induced by motion between the acquisition of the different segments must be estimated and corrected.

The navigator echo technique was introduced for this purpose [34, 33]. In its simplest form, it consists of a single line through the centre of k -space, acquired immediately before or after the acquisition of the corresponding k -space segment, after the diffusion-encoding has been applied. The phase shift present in the navigator echo is then identical to that present in the corresponding k -space segment. The relative phase shifts between the different navigator echoes can be estimated and used to correct the relative phase shifts between the corresponding k -space segments before they are recombined to form an artefact-free image.

In addition to phase errors, this simple one-dimensional navigator echo approach can also correct for shifts in the position of the centre of k -space along the axis of acquisition of the navigator echo. Thus, it can correct artefacts due to translation, as well as a limited subset of rotations. A wider range of rotations can be corrected using a two-dimensional navigator echo approach, where the central portion of k -space is acquired instead of just a single line through it [35]. However, it is still not possible to correct for the set of rotations that cause the centre of k -space to shift out of the plane of k -space acquired (for example, a rotation about the frequency-encoding axis with the diffusion gradient applied

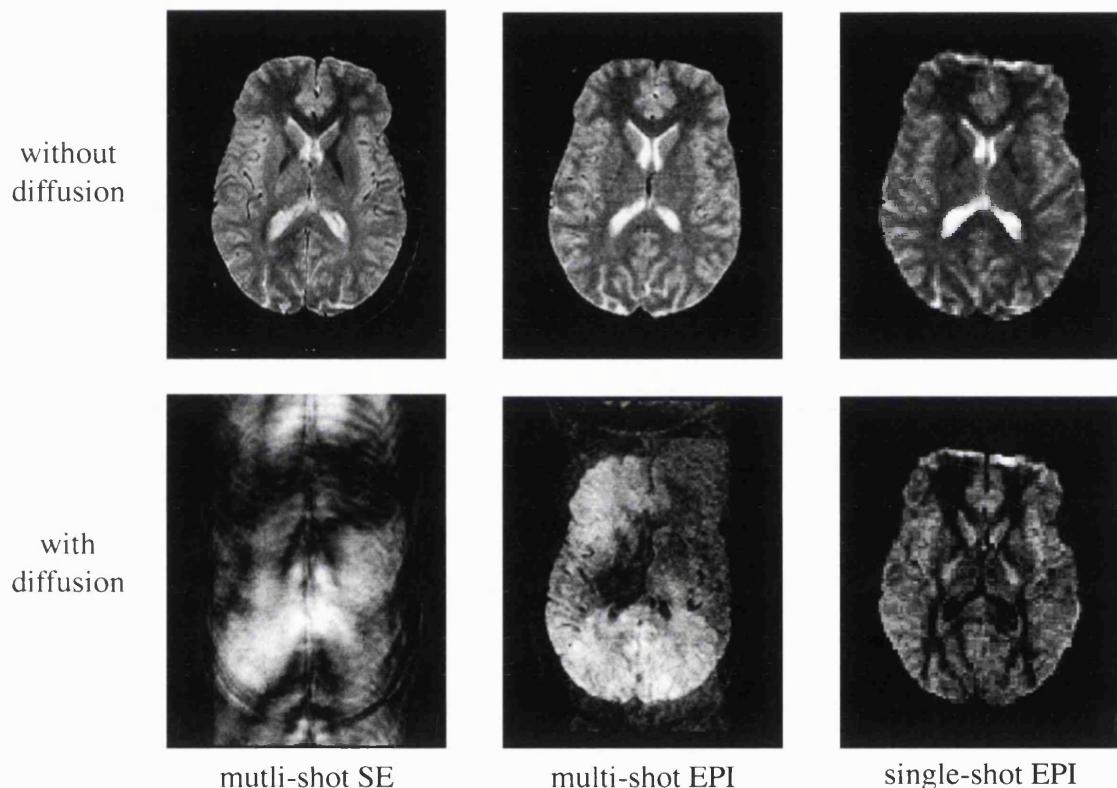


Figure 3.7: images acquired using three different diffusion-weighted sequences: conventional multi-shot 2D-FT spin-echo sequence (left column), multi-shot interleaved segmented EPI (16 segments) (middle column), and single-shot EPI (right column). The top row shows images acquired with a b -value = 0 s/mm^2 (unweighted images), while the bottom row shows images acquired with a b -value $\simeq 600 \text{ s/mm}^2$. Motion induces phase errors in the multi-shot sequences, producing ghosting in the image. (Adapted from F. Calamante's PhD thesis, 2000)

along the phase-encoding axis), and k -space segments that are corrupted in this way need to be re-acquired [36]. Although it may be possible to correct for these rotations using a 3-dimensional navigator echo technique, such an approach is impractical.

PULSATILE MOTION

Diffusion-weighted sequences are also sensitive to the motion induced by pulsations in the vasculature. If such a pulsation occurs during or between the diffusion-encoding gradients, spins will experience an extra phase shift. However, in contrast to the bulk motion case, these shifts may not be coherent on the scale of a voxel, leading to additional spin dephasing and signal attenuation. Since these are local phase shifts, they cannot be corrected using navigated echoes. This manifests itself in diffusion-weighted imaging as signal dropout in the affected area.

This artefact can be alleviated by using cardiac gating, and synchronising the acquisition of the images with the stable phase of the cardiac cycle [37]. For single shot techniques, this simply lengthens the total scan time, since the number of images per cardiac cycle is limited by the duration of the appropriate cardiac phase. When using multi-shot techniques, it is also important to minimise phase inconsistencies between the different segments of each image. These segments therefore need to be acquired at the same point in the cardiac cycle, whilst also avoiding the unstable phase of the cardiac cycle.

EDDY CURRENTS

The diffusion sensitivity is commonly induced using large prolonged gradients that are ramped as fast as possible in order to keep the echo time at a minimum. These sequences are therefore particularly sensitive to eddy-current artefacts (see section 2.3.3). Moreover, a number of diffusion-weighted images are typically combined together to form, for example, an ADC map. Since each of these images is acquired using a different gradient orientation, they are affected by different eddy-currents, and are distorted in different ways (see section 2.3.3) [38]. To produce artefact-free maps, these distortions need to be minimised or corrected, such that the images match up correctly [39].

In addition to the techniques previously covered in section 2.3.2, various methods have been proposed to reduce these effects in DWI. Careful design of the gradient pattern may significantly reduce distortions. For example, two negative ‘dummy’ gradient pulses can be added to a PGSE sequence, one before the excitation pulse, the other after the read-out period, such that they introduce additional eddy-currents that counter the effects of those set up by the positive diffusion gradients, without otherwise modifying the imaging sequence [40]. Alternatively, the durations of the gradient pulses in a twice-

refocused spin-echo sequence (described on page 51) can be tuned to effectively eliminate the effects of eddy-currents with a particular time constant (see section 4.2.1 for a detailed description).

It is also possible to correct some of the phase errors induced by eddy-currents by acquiring additional data without phase-encoding [38, 41]. Each line of k -space in the non-phase-encoded data contains information about the evolution of the phase during data collection that can be used to correct the phase-encoded data. However, the disadvantage of these techniques in single-shot EPI is that the duration of the experiment is increased two- or three-fold. Another common technique uses affine registration to retrospectively re-align the images with a reference image [42, 43], which has the advantage that it does not require the acquisition of additional data.

3.3 DIFFUSION TENSOR IMAGING

As previously mentioned, the ADC depends on the direction in which it is being measured. This anisotropy is due for the most part to the presence of barriers to diffusion, such as cell membranes, organised such that diffusion is more hindered in certain directions and less in others. It is thus particularly apparent in tissues with a coherently oriented microarchitecture, such as white matter tracts. To describe such an orientation dependent system, the diffusion tensor model was introduced by Basser and Le Bihan in 1994 [20].

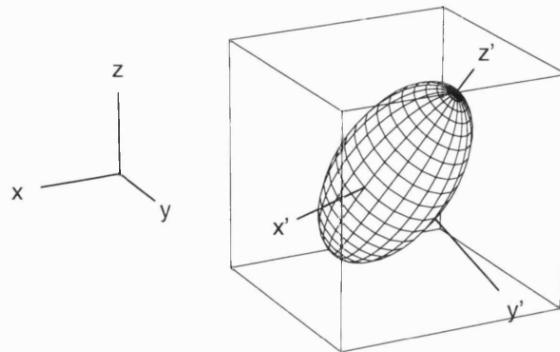
3.3.1 THE DIFFUSION TENSOR

The diffusion tensor model uses a 3×3 symmetric tensor to characterise the diffusion properties of the medium. The spin displacement probability density function in this framework takes the form (cf. equation 3.11):

$$P(\underline{r} \mid \Delta) = \frac{1}{\sqrt{(4\pi\Delta)^3 \det(\mathbf{D})}} \exp\left(-\frac{\underline{r}^T \mathbf{D}^{-1} \underline{r}}{4\Delta}\right) \quad (3.29)$$

where \underline{r} is the spin displacement vector from its original position, Δ is the effective diffusion time, \mathbf{D} is the diffusion tensor, and $\det(\mathbf{D})$ and \mathbf{D}^{-1} denote respectively the determinant and inverse of \mathbf{D} .

The diagonal elements of \mathbf{D} , namely D_{xx} , D_{yy} , and D_{zz} , reflect correlations between displacements in the same directions, whereas its off-diagonal elements, D_{xy} , D_{xz} , and D_{yz} , reflect correlations between molecular displacements in orthogonal directions [44]. For example, when observing diffusion in a lattice of semi-permeable parallel layers aligned obliquely with respect to the x and y axes, macroscopic displacements along



$$\mathbf{D} = \begin{pmatrix} 0.45 & -0.11 & -0.10 \\ -0.11 & 0.62 & 0.21 \\ -0.10 & 0.21 & 0.83 \end{pmatrix} \quad \mathbf{D}' = \begin{pmatrix} 0.4 & 0 & 0 \\ 0 & 0.5 & 0 \\ 0 & 0 & 1 \end{pmatrix}$$

Figure 3.8: a typical diffusion ellipsoid in the lab frame of reference, represented by the x , y and z axes. In this frame of reference, the diffusion tensor \mathbf{D} is represented by the matrix shown in the bottom left. The diagonalisation process aligns the axes of the coordinate system with the tensor's principal directions. In its *eigen* frame of reference, represented by the x' , y' and z' axes, the diffusion tensor \mathbf{D}' is represented by the matrix shown in the bottom right. Note that all off-diagonal elements are now zero. The elements along the diagonal corresponds to the tensor's *eigenvalues*. The corresponding *eigenvectors* are aligned along the x' , y' and z' axes.

x will appear to be correlated with displacements along y , resulting in a non-zero D_{xy} element. This formulation offers a number of advantages, outlined below.

THE EIGENSYSTEM

Although the actual elements of \mathbf{D} may vary as the laboratory frame is rotated, many of the properties described by the tensor will not change. Those properties are termed *rotationally invariant*, and are of particular interest since they will not depend on the orientation of the sample with respect to the laboratory frame. If the lattice of layers described above is rotated such that the layers lie perpendicular to the x -axis, displacements along x will cease to be correlated with displacements along y , and D_{xy} will be zero. In general, when the principal directions of the sample are aligned with the axis system, all off-diagonal elements disappear, leaving a diagonal tensor (see figure 3.8). In this frame of reference, the diagonal elements of the tensor are known as its *eigenvalues*, and the vectors pointing along the principal axes are known as its *eigenvectors*. Each eigenvalue is rotationally invariant, and is associated with its corresponding eigenvector, the elements of which are dependent on the orientation of the laboratory frame [44]. The eigenvalues are normally sorted in order of decreasing magnitude, and referred to as λ_1 , λ_2 , and λ_3 , while their corresponding eigenvectors are known as $\hat{\epsilon}_1$, $\hat{\epsilon}_2$, and $\hat{\epsilon}_3$. The eigenvector ($\hat{\epsilon}_1$)

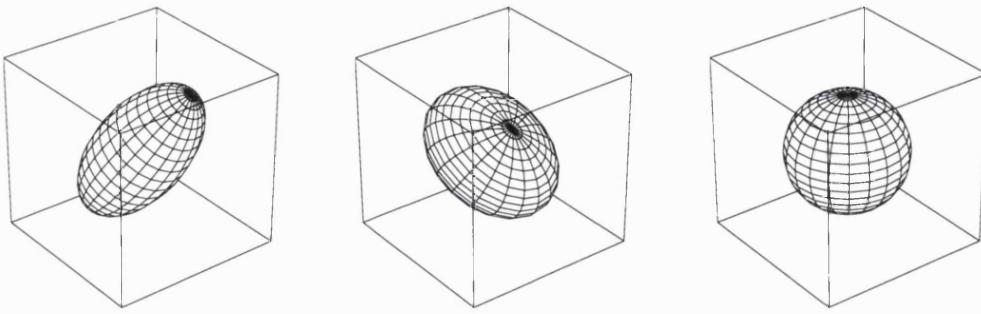


Figure 3.9: the properties of the diffusion tensor can be expressed in terms of the shape of its characteristic ellipsoid. Left: a prolate tensor ($\lambda_1 > \lambda_2 = \lambda_3$). Middle: an oblate tensor ($\lambda_1 = \lambda_2 > \lambda_3$). Right: an isotropic tensor ($\lambda_1 = \lambda_2 = \lambda_3$).

corresponding to the largest eigenvalue (λ_1) gives the direction of fastest diffusion, and is commonly known as the major eigenvector of the diffusion tensor.

The diffusion tensor is often represented visually by its characteristic ellipsoid [20], as shown in figure 3.8, corresponding to a surface of constant spin displacement probability. It can be seen from equation 3.29 that the surface of this ellipsoid is given by:

$$\frac{\underline{r}^T \mathbf{D}^{-1} \underline{r}}{2\Delta} = C \quad (3.30)$$

where C is a constant. In the diffusion tensor's eigensystem, all off-diagonal elements disappear, and this equation reduces to:

$$\frac{(r_1)^2}{2\Delta\lambda_1} + \frac{(r_2)^2}{2\Delta\lambda_2} + \frac{(r_3)^2}{2\Delta\lambda_3} = 1 \quad (3.31)$$

where C has been set to 1, and r_i represents the component of \underline{r} along $\hat{\epsilon}_i$. The dimensions of the diffusion ellipsoid along its axes therefore correspond to the root mean square spin displacements along the corresponding directions (cf. equation 3.5).

The properties of the diffusion tensor can be more clearly visualised using this representation. For example (figure 3.9), a tensor with $\lambda_1 > \lambda_2 \approx \lambda_3$ corresponds to a *prolate* ellipsoid, with a characteristic cigar-shaped appearance. A tensor with $\lambda_1 \approx \lambda_2 > \lambda_3$ corresponds to an *oblate* ellipsoid, with a characteristic pancake-shaped appearance. Finally, an isotropic tensor with $\lambda_1 = \lambda_2 = \lambda_3$ corresponds to a sphere.

The diffusion tensor can be derived from a combination of at least 6 diffusion-weighted measurements applied along non-collinear directions (in addition to measurements acquired to estimate the value of M_0 , the signal in the absence of any diffusion-weighting), by inverting equation 3.16:

$$M_i = M_0 e^{-\mathbf{b}_i \cdot \mathbf{D}} \quad (3.32)$$

The calculation is often done using linear regression, by rewriting the equation above as:

$$\mathbf{B} \cdot \underline{d} = \underline{m} \quad (3.33)$$

where:

$$\mathbf{B} = \begin{pmatrix} b_{1,xx} & b_{1,yy} & b_{1,zz} & 2b_{1,xy} & 2b_{1,xz} & 2b_{1,yz} & -1 \\ b_{2,xx} & b_{2,yy} & b_{2,zz} & 2b_{2,xy} & 2b_{2,xz} & 2b_{2,yz} & -1 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ b_{N,xx} & b_{N,yy} & b_{N,zz} & 2b_{N,xy} & 2b_{N,xz} & 2b_{N,yz} & -1 \end{pmatrix}$$

$$\underline{d} = \begin{pmatrix} D_{xx} \\ D_{yy} \\ D_{zz} \\ D_{xy} \\ D_{xz} \\ D_{yz} \\ \ln(M_0) \end{pmatrix} \quad \text{and} \quad \underline{m} = \begin{pmatrix} \ln(M_1) \\ \ln(M_2) \\ \ln(M_3) \\ \vdots \\ \ln(M_N) \end{pmatrix}$$

The matrix \mathbf{B} can then be inverted using standard techniques to find \underline{d} :

$$\underline{d} = \mathbf{B}^{-1} \cdot \underline{m} \quad (3.34)$$

The properties of the diffusion can also be visualised by plotting the apparent diffusion coefficient (estimated using the simple equation $M = M_0 e^{-bD}$) as a function of orientation, as shown in figure 3.10. This representation has the advantage that it is not dependent on the diffusion tensor model, which may not always be appropriate (see section 3.3.5).

THE TRACE OF THE DIFFUSION TENSOR

The trace of the diffusion tensor, $\text{tr}(\mathbf{D})$, is given by the sum of its diagonal elements, and is rotationally invariant [20]. In the eigensystem, the diagonal elements correspond to the eigenvalues, such that the average eigenvalue $\bar{\lambda}$ corresponds to $\text{tr}(\mathbf{D})/3$. Since the trace of the diffusion tensor is the mean diffusivity, averaging measurements made along a large number of uniformly distributed directions gives $\bar{\lambda}$. However, three ADC measurements made in orthogonal directions are sufficient to estimate the mean diffusivity. Alternatively, $\text{tr}(\mathbf{D})$ can also be measured directly using a trace-weighted sequence as described in section 3.2.2. In the normal human brain, the mean diffusivity is approximately $850 \times 10^{-6} \text{ mm}^2/\text{s}$ in grey and white matter, and close to that of free water in the CSF [45]. Trace maps such as that shown in figure 3.11 therefore appear relatively

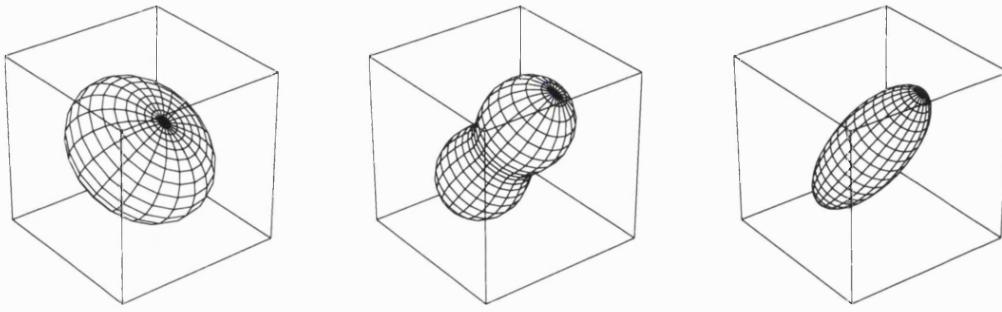


Figure 3.10: the relationship between the diffusion-weighted signal and the diffusion tensor. Left: the diffusion-weighted signal (measured using the same b -value) plotted as a function of orientation. The signal is lowest along the direction of fastest diffusion. Middle: the corresponding apparent diffusion coefficient (calculated using the simple equation $M = M_0 e^{-bD}$) plotted as a function of orientation. Right: the diffusion tensor estimated from the diffusion-weighted measurements shown on the left by linear regression (equation 3.33).

homogeneous over the entire brain, except in regions containing CSF.

DIFFUSION ANISOTROPY

The anisotropy of the diffusion process is a measure of its deviation from the isotropic diffusion case (where the ADC is independent of orientation). In biological tissues, this anisotropy arises from coherently oriented barriers to diffusion. The degree of anisotropy in bundles of neuronal axons has been shown to depend on a number of factors, in particular axonal density, fibre radius, and the degree of myelination [9, 16, 46].

A number of expressions have been proposed as objective markers of anisotropy [47, 48]. The first measures proposed were functions of the ADC as measured along two orthogonal directions [19]. Such an approach is however dependent on the orientation of the sample with respect to the laboratory frame, and is therefore not rotationally invariant. Anisotropy is now commonly measured using rotationally invariant measures, the most common of which are described below.

The volume ratio anisotropy index is defined as [49]:

$$A_{VR} = 1 - \frac{\prod_i \lambda_i}{\bar{\lambda}^3} = 1 - 27 \frac{\det(\mathbf{D})}{[\text{tr}(\mathbf{D})]^3} \quad (3.35)$$

where $\bar{\lambda}$ is the average eigenvalue, given as $\text{tr}(\mathbf{D})/3$, and $\prod_i \lambda_i$ is the product of the eigenvalues. This measure corresponds to the ratio of the volume of the diffusion tensor's ellipsoid to the volume of the sphere representing isotropic diffusion with $\text{ADC} = \bar{\lambda}$. In this definition the measure has been scaled such that $A_{VR} = 0$ corresponds to the isotropic

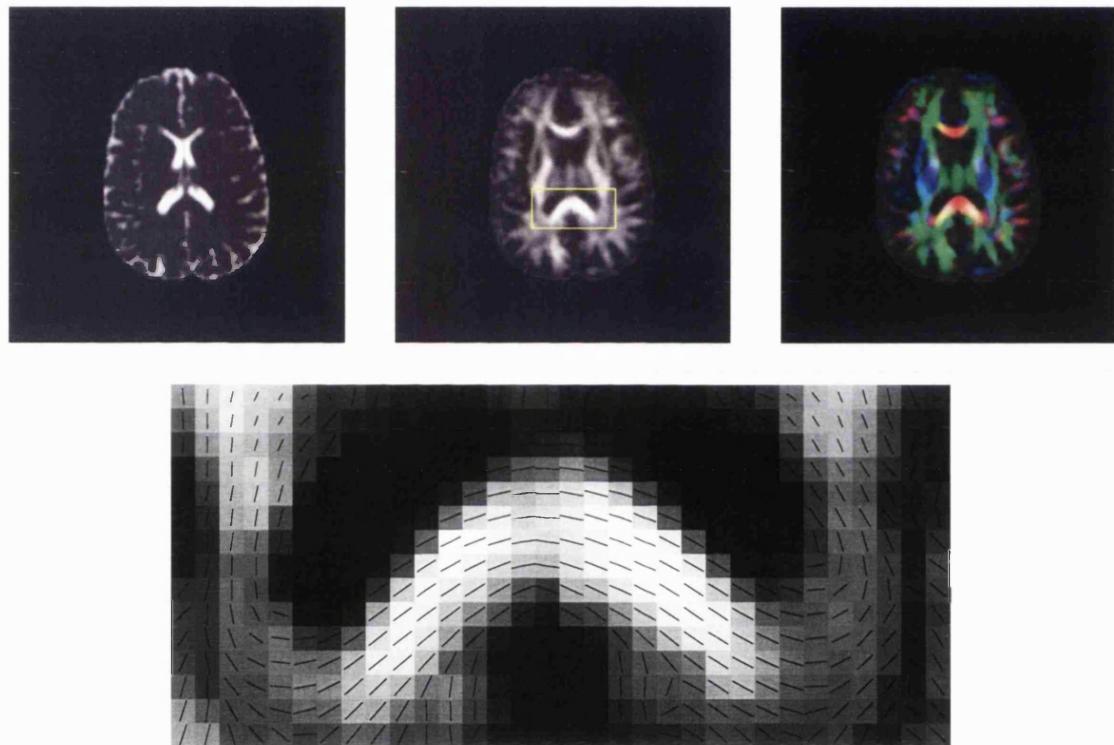


Figure 3.11: various maps available from the diffusion tensor. Top left: a map of the mean diffusivity. Top middle: a map of the fractional anisotropy (FA). Top right: a colour-coded map of the major eigenvector of the diffusion tensor weighted by the corresponding FA map. Bottom: an overlay of the major eigenvector onto a map of the fractional anisotropy for the region highlighted on the top middle image. The combination of these maps provides good contrast: grey matter has normal trace and low FA, white matter has normal trace and high FA, and CSF has high trace and low FA. See text for details.

anatomical region	FA	A_σ	A_{VR}
corpus callosum	0.65 ± 0.16	0.50 ± 0.02	0.86 ± 0.05
internal capsule	0.38 ± 0.06	0.39 ± 0.02	0.70 ± 0.08
optic radiation	—	0.31 ± 0.02	0.62 ± 0.12
frontal white matter	0.34 ± 0.12	0.26 ± 0.01	—
periventricular white matter	0.34 ± 0.08	—	0.27 ± 0.03
caudate nucleus	—	0.10 ± 0.01	0.08 ± 0.03

Table 3.1: diffusion anisotropy values measured in healthy adults in various regions of the brain. Fractional anisotropy (FA) values from Filippi *et al* [52], total anisotropy (A_σ) values from Shimony *et al* [53], and volume ratio (A_{VR}) values from Pierpaoli *et al* [54].

case, and $A_{VR} = 1$ corresponds to the highest possible anisotropy.

The relative anisotropy was introduced by Basser *et al* [50] and is defined as:

$$\text{RA} = \frac{\sqrt{\frac{1}{3} \sum (\lambda_i - \bar{\lambda})^2}}{\bar{\lambda}} = \frac{\sqrt{\mathbf{D}_a : \mathbf{D}_a}}{\bar{\lambda}\sqrt{3}} \quad (3.36)$$

where the colon operator denotes the tensor dot product as defined in equation 3.17, and $\mathbf{D}_a = \mathbf{D} - \bar{\lambda}\mathbf{I}$, such that \mathbf{D}_a contains only the anisotropic part of \mathbf{D} (\mathbf{I} is the 3×3 identity matrix). This measure corresponds to the standard deviation of the eigenvalues, normalised to the mean eigenvalue. This measure is similar to the *total anisotropy* index, A_σ , introduced by Conturo [51], which corresponds to $A_\sigma = \text{RA}/\sqrt{2}$.

The fractional anisotropy, also introduced by Basser *et al* [50], is defined as:

$$\text{FA} = \sqrt{\frac{3}{2}} \frac{\sqrt{\frac{1}{3} \sum (\lambda_i - \bar{\lambda})^2}}{\sqrt{\frac{1}{3} \sum \lambda_i^2}} = \sqrt{\frac{3}{2}} \sqrt{\frac{\mathbf{D}_a : \mathbf{D}_a}{\mathbf{D} : \mathbf{D}}} \quad (3.37)$$

This measure corresponds to the standard deviation of the eigenvalues, normalised to the root mean square eigenvalue. It is scaled by a factor of $3/2$ such that the maximum possible value of FA is 1.

These three rotationally invariant anisotropy indices all have the advantage that they can easily be calculated from the elements of the diffusion tensor; there is no need to diagonalise the tensor into its eigensystem, which is a computer-intensive process. In practice, the FA is the most commonly used anisotropy index, as it has been shown to be more robust to noise and to provide better contrast than other indices [48]. Typical values of the diffusion anisotropy for various tissues are shown in table 3.1, and a map of the fractional anisotropy acquired on a healthy volunteer is shown in figure 3.11.

THE MAJOR EIGENVECTOR OF THE DIFFUSION TENSOR

It is often assumed that the direction of fastest diffusion in brain white matter corresponds to the orientation of the underlying neuronal fibres (see section 3.3.3). This direction can be estimated from the diffusion tensor by evaluating its major eigenvector, and visualised using maps such as those shown in figure 3.11. These maps can be generated in two ways. First, the major eigenvector can be projected and overlaid onto a 2D anatomical slice, such as a map of diffusion anisotropy. Although such maps give an accurate representation of the direction of the major eigenvector when zooming in on a specific area of interest, the vectors are difficult to see when displaying full brain images. The second type of eigenvector map is generated by assigning a colour to each direction [55]. Colours on visual display units are usually defined by their red, green and blue components. A vector \hat{e} can therefore be represented by a colour, by setting the red component to $|\hat{e}_x|$, the green component to $|\hat{e}_y|$, and the blue component to $|\hat{e}_z|$. Therefore, regions of white matter where the fibres run primarily from left to right will appear red, regions where fibres run anterior to posterior will appear green, and regions where fibres run superior to inferior will appear blue. Such maps can be used to visualise white matter architecture over the whole brain, but do not provide an accurate representation of the orientation of the major eigenvector, as explained by Pierpaoli *et al* [55].

3.3.2 DIFFUSION ENCODING SCHEMES

In order to calculate the diffusion tensor, 7 variables need to be estimated: the tensor's 6 independent elements, and M_0 . This means that at least 7 measurements must be performed: one at low b -value and 6 at high b -value along non-collinear directions. A number of recent studies have explored the optimal set of diffusion encoding parameters that would allow the most robust determination of the diffusion tensor within a practical scan time. The parameters to be determined are the b -values to be used, and the directions of application of the diffusion encoding gradients.

Jones *et al* [28] have shown that for the optimal determination of various properties of the diffusion tensor, the following scheme should be employed: two b -values (b_1 and b_2) should be used, with $(b_2 - b_1)_{opt} = 1.09/\text{ADC}$; b_1 should be made as small as possible to maximise the SNR; the ratio of the number of high b -value scans acquired to low b -value scans acquired should be approximately 5.6:1 for trace measurements, or 11.3:1 for estimating the elements of the tensor; the set of orientations along which the diffusion-weighting is applied should be uniformly distributed over the surface of a hemisphere.

Papadakis *et al* [56] showed that although a particular sampling strategy may minimise the variance in the anisotropy of the tensor, the variance itself may be orientationally

dependent. However, this dependence becomes negligible when using a sampling scheme consisting of at least 20 uniformly distributed directions. It has also been shown that using more than approximately 20 directions minimises the entropy of fractional anisotropy maps [57].

A number of techniques for obtaining a set of uniformly distributed orientations along which to sample the ADC have been proposed. Jones *et al* [28] and Papadakis *et al* [56] used an electrostatic repulsion model, whereby each direction was represented by a pair of diametrically opposed charges at a fixed radius such that, at equilibrium, the charges are as far apart from each other as possible. Hasan *et al* [58] have proposed a number of possible measures of performance that could be minimised with respect to the gradient orientations to provide an optimal set. Tuch *et al* [59] obtained a set of directions from the tessellation of an icosahedron. Skare *et al* [60] proposed to use the condition number as a measure of a scheme's robustness to noise. Finally, Batchelor *et al* [61] showed that icosahedral schemes have properties that make them particularly suited to diffusion tensor measurements: schemes based on members of the icosahedral family have the same condition number as that using an infinite number of directions, and are rotationally invariant (the variance in the estimated diffusion tensor elements is independent of the orientation of the tensor with respect to the measurement axes).

Diffusion tensor acquisitions are therefore commonly done using a combination of low (i.e. close to zero) b -value measurements, and a larger number of high b -value (typically 1000 s/mm²) measurements with the diffusion encoding applied in a number of directions. The diffusion encoding should be applied in as many directions as the available scan time allows. Rotational invariance can be achieved either by using at least 20 uniformly distributed sampling directions [56, 57], and/or using an icosahedral scheme [61].

3.3.3 TRACTOGRAPHY

Brain white matter consists of fibre bundles or *tracts*, made up primarily of the axons of neurons linking different structures [62]. Diffusion in this type of structure is anisotropic, due to the greater hindrance to the movement of water molecules across the tightly packed axon bundles than along their axis [9]. The major white matter tracts are therefore clearly visible in anisotropy maps. The eigensystem decomposition of the diffusion tensor, described in section 3.3.1, also provides information about the direction of fastest diffusion, which can be exploited in tractography or fibre-tracking applications.

It is often assumed that the direction of fastest diffusion corresponds to the orientation of the underlying white matter tract fibres. This has been shown to be the case, for example, in the optic tract of the rat, using manganese-enhanced MRI [63]. It may then

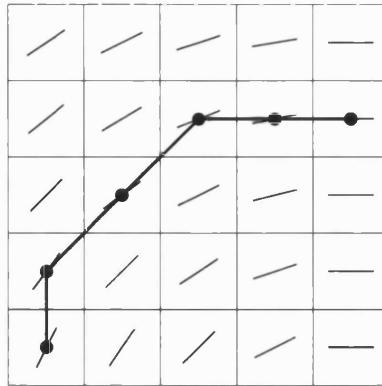


Figure 3.12: A schematic representation of voxel-linking tractography techniques. The figure shows the orientation of the major eigenvector within each voxel. At each step, voxels adjacent to the current voxel are examined to determine which one is most likely to be connected with the current voxel, according to a set of rules specific to the algorithm. A link is then established to the voxel selected, which becomes the current voxel. This process is repeated until an algorithm-specific termination criterion is satisfied.

be possible to follow this direction step by step through the data set, thus delineating the path of the tract of interest. In this way, neuronal connections between different anatomical regions of the brain may be inferred. To date, a number of different tractography algorithms have been proposed [64, 65, 66, 67, 68, 69, 70]. Most of these share a number of characteristics in common. For example, the track will often be terminated if it enters a region of low anisotropy, or if its curvature exceeds a predefined threshold. Many of these algorithms fall into two main categories, according to their approach to the problem: voxel-linking or streamlines. Each of these approaches has its own advantages and disadvantages, which will be discussed in more detail in chapter 5. A brief description of some of the most common techniques is given below.

VOXEL-LINKING TECHNIQUES

This approach is based on establishing links between adjacent voxels as shown in figure 3.12, according to a set of algorithm-specific rules. These links form a chain through the data set, thereby potentially enabling the identification of connections between different brain regions. The rules used to establish these links vary widely between different algorithms, making this approach very versatile. Another advantage of this approach is that the tracks generated are allowed to branch to different regions of the brain, such that the most likely connections from the seed point can be displayed. It is also possible to extend it so that each connection is given an indication of its likelihood.

Jones *et al* [64] allowed links to be established between adjacent voxels if a circular

arc could be drawn between the centre of the two voxels, such that the major eigenvectors of their corresponding diffusion tensors were tangential to that arc within a predefined tolerance. Their technique allowed the delineation of groups of voxels likely to belong to the same white matter structure.

The technique introduced by Poupon *et al* [70] establishes a ‘forward’ and a ‘backward’ link for each voxel, according to the orientation of the major eigenvector. Links are established in each direction to the adjacent voxel that minimises the curvature of the track. This technique intrinsically allows the track to branch, since a particular voxel may establish a connection to another that is already connected to two other voxels. The algorithm as described is also combined with a method to regularise the major eigenvector field, using a ‘spaghetti plate’ model to minimise the curvature in the field.

Parker *et al* [67] have recently proposed a technique based on the evolution of a front from an initial user-defined seed point, using the principles of level set theory and the fast marching algorithm. The front evolves at a rate governed by the orientation of the major eigenvector with respect to the surface of the front. A path is then calculated from every point in the brain back to the seed point by following the gradient of the map of the time of arrival of the front, and a connectivity index is assigned to each path. Only the highest connectivity paths are eventually displayed. This allows multiple connections to be established to the seed point, thus intrinsically allowing the tracks generated to branch or diverge. The technique also provides a pseudo-probabilistic measure for the likelihood of each track given the data.

Koch *et al* [69] recently proposed a Monte Carlo approach to establish pseudo-probabilistic estimates of connectivity. Starting from a user-defined seed point, a random ‘jump’ is made to one of three adjacent voxels in the forward direction with respect to the direction of fastest diffusion in the current voxel. A fixed number of such ‘jumps’ are performed for each iteration of the algorithm. After a large number of iterations, the ‘probability’ of connectivity for a particular voxel is given by the number of times the algorithm ‘jumped’ to that voxel, normalised to the number of iterations performed.

Voxel-linking techniques are however limited in two respects. First, the number of orientations that the track can follow is limited to those joining adjacent voxels (a total of 26 directions). This results in a much lower angular resolution (45°) than can be achieved with DTI (the major eigenvector can typically be determined to within 10° [63, 71]). Second, the path of each track is constrained to pass through the centre of the different voxels along it. This introduces errors in the estimation of the true path that may cause the track to venture into unrelated structures and ultimately lead to a very different region of the brain.

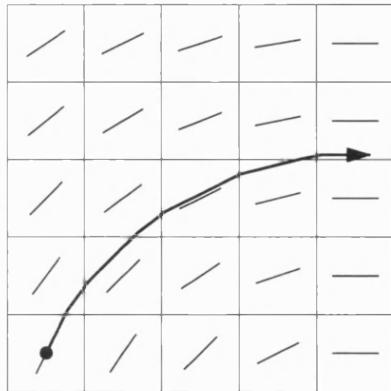


Figure 3.13: a schematic representation of a streamlines algorithm. The figure shows the orientation of the major eigenvector within each voxel. Starting from a user-defined seed point, a small step is repeatedly taken along the direction of the major eigenvector, until an algorithm-specific termination criterion is satisfied. In contrast to voxel-linking techniques, the tensor field is treated as continuous, such that the path of the track is not restricted to the centre of each voxel.

THE STREAMLINES APPROACH

The diffusion tensor field can be expressed as a continuous function of position. The simplest such description is obtained by assuming that the tensor at any point in space is that of the nearest voxel. More sophisticated methods are available that involve interpolation, resulting in a smoother description of the tensor field. In such a framework, tracks do not need to jump from voxel to voxel, but can follow any continuous path in three dimensional space. Fibre-tracking techniques based on the streamlines technique generate tracks by stepping along the presumed direction of the underlying white matter tract (normally the major eigenvector of the diffusion tensor) by a small amount for each iteration, until some terminating criterion is reached, as illustrated in figure 3.13. The size of each step is normally referred to as the *step size*, and should be smaller than the dimensions of a voxel.

Mori *et al* [65] proposed the well-known *fibre assignment by continuous tracking* (FACT) algorithm, which uses the simple non-interpolated continuous tensor field approximation described in the previous paragraph. Conturo *et al* [66] proposed a similar algorithm, using an interpolated description of the tensor field. Each element of the tensor at the point of interest is given by the tri-linear interpolation of the 8 nearest neighbours' corresponding tensor elements. Basser *et al* [68] approximate the tensor field using a set of B-spline functions that provides estimates of the 2nd and higher derivatives of the major eigenvector. This allows the use of the 4th order Runge-Kutta method, and hence a more precise determination of the track path. In addition, the process of fitting B-splines

to the tensor data intrinsically imposes constraints on the form of the tensor field and increases its smoothness.

Although streamlines techniques do not have the angular or spatial constraints of the voxel-linking approach, they do have a number of disadvantages. Since only one direction is considered at each step, the tracks generated are not allowed to branch. There is only one track per seed point, and no indication is given of other possible connections, or of the likelihood of the track generated. The major eigenvector is sensitive to noise and partial volume effects (see section 3.3.4), and may be corrupted. A track generated using this corrupted orientation may establish a connection that does not exist in reality, without giving any indication of its likelihood or of possible alternatives.

3.3.4 ARTEFACTS

In addition to artefacts specific to DWI outlined in section 3.2.3, the low SNR and low spatial resolution of DWI causes additional problems particular to tractography. These effects are discussed in more detail in chapter 5.

NOISE CONTAMINATION

The presence of noise in the diffusion-weighted images will propagate to the tensor elements [28, 72, 73, 74]. As mentioned in section 3.3.2, optimising the acquisition scheme is essential to minimise the sensitivity of the tensor to the variance in the diffusion data. However, some noise will inevitably be present in the DTI data, affecting the various maps that can be derived.

In particular, the anisotropy of the tensor will tend to be overestimated in the presence of noise, especially at low values of the anisotropy [75]. Consider the case of a perfectly isotropic tensor, which would have a zero anisotropy index: if noise is present in the images, the estimated tensor will deviate from its ideal isotropy, raising its anisotropy index. Thus, at low SNR, anisotropy maps will be biased in addition to the expected increase in the variance.

The presence of noise also introduces an error in the estimated direction of the major eigenvector, depending on the anisotropy of the tensor, the SNR and the diffusion encoding scheme [72, 75]. This is particularly relevant for tractography applications, since a corrupted orientation may cause the track to venture into an adjacent, yet unrelated structure, establishing connections where none exists in reality.

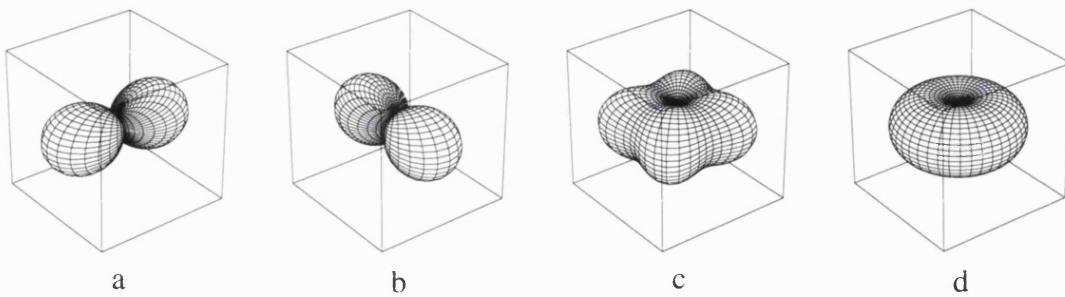


Figure 3.14: the effect of partial volume averaging. **(a):** the profile of the apparent diffusion coefficient measured as a function of orientation for a single fibre population (equivalent to a single prolate diffusion tensor, see figure 3.10). **(b):** the diffusion profile for a fibre population at right angles to, but otherwise identical to that in (a). **(c):** the apparent diffusion coefficient profile that would be measured in a voxel consisting of equal fractions of the fibre populations in (a) and (b), in the slow exchange limit. Note that the peaks of the diffusion profile do not coincide with the orientations of either fibre populations present. **(d):** the diffusion tensor that would be estimated from the diffusion profile in (c).

PARTIAL VOLUME EFFECTS

Due to the typically low spatial resolution of DWI, multiple distinct fibre populations may contribute to the signal from a single voxel. Unfortunately, the diffusion tensor model is unable to describe more than one fibre population at any one point [59, 76, 77, 78]. In this case, the diffusion tensors of the various fibre populations present within the voxel will be ‘averaged’ in some way to give a single tensor. This introduces significant errors in estimates of the anisotropy, and has profound implications for fibre-tracking applications [79, 80].

Consider for example a voxel consisting of a mixture of two distinct fibre populations, with one oriented at an angle with respect to the other. This situation could happen for example where two fibre tracts intersect, or where the tracts simply run very close to each other. Although the diffusion tensor within each fibre population may be highly anisotropic and *prolate* (i.e. $\lambda_1 > \lambda_2 \simeq \lambda_3$), the diffusion tensor estimated within that voxel may appear quite different. In the particular case that the fibre populations have identical diffusion properties, contribute equally to the signal, and intersect each other at right angles, the tensor estimated would appear *oblate* (i.e. $\lambda_1 \simeq \lambda_2 > \lambda_3$), and the anisotropy would be significantly reduced (see figure 3.14). In general, where two or more fibre populations are present, the diffusion tensor will tend to become more oblate or even isotropic, depending on the arrangement of the fibres, leading to a reduction of the anisotropy [76]. In order to minimise this problem, it has recently been proposed

to use the variance in the ADC measured using a high angular resolution scheme (see section 3.3.5) as an alternative measure of diffusion anisotropy that is not derived from the tensor model [81].

The impact of such a situation on tractography applications is considerable. One consequence of the reduction in anisotropy is that tracks may be terminated prematurely. Of far greater concern however is the fact that the major eigenvector of the diffusion tensor will not in general be consistent with any of the fibre populations present [59]. Tracking through regions affected by partial volume effects is therefore very unreliable, as illustrated by Pierpaoli *et al* [82], and is a problem that is exacerbated by not knowing which areas are affected.

Although the impact of partial volume effects can be reduced by increasing the resolution of the images, this would inevitably reduce the SNR, introducing other problems previously described. To address these issues, alternatives to the tensor model (described in the next section) have been proposed that may be better suited to the problem. However, even using more sophisticated models, it is not clear whether the distinction between crossing and ‘kissing’² fibres can be made using diffusion-weighted MRI, since the diffusion characteristics are expected to be similar.

3.3.5 ALTERNATIVE MODELS FOR DIFFUSION

The diffusion tensor model is valid for free Gaussian diffusion. However, it is widely acknowledged that the diffusion anisotropy observed in the brain is due to the presence of barriers to diffusion, which would invalidate the model. Despite this, it remains the most commonly used model of diffusion in the brain as measured by NMR, due to its simplicity, and the fact that it seems adequate in regions free from partial volume effects [39]. However, deviations from the tensor model can readily be observed in regions containing significant partial volume effects, such as crossing fibres [59, 78, 81].

As mentioned in section 3.1.3, in the slow exchange regime, the NMR signal from the various populations will add together. In this case, a single tensor will not be an adequate fit to the set of ADCs measured over a large number of directions. To describe such a system, two approaches can be employed, namely *q*-space and high angular resolution imaging, each of which requires particular data acquisition conditions to be fulfilled.

q-SPACE

The *q*-space formalism states that, assuming the narrow-pulse approximation (see section 3.1.1), there is a Fourier relationship between the signal measured as a function of

²Fibres are said to ‘kiss’ if they pass close to each other without crossing, forming a \cap shape.

the diffusion-encoding gradient and the spin displacement probability density function (PDF) [83]. Thus, using a diffusion time Δ and a gradient pulse duration δ , the signal $M(\underline{x}, \underline{q})$ measured at position \underline{x} in the sample as a function of $\underline{q} = \gamma\delta\underline{G}$ is given by:

$$M(\underline{x}, \underline{q}) = \rho(\underline{x}) \int P(\underline{r} \mid \underline{x}, \Delta) e^{i \underline{q} \cdot \underline{r}} d\underline{r} \quad (3.38)$$

where $\rho(\underline{x})$ is the spin density (ignoring the effects of relaxation), and $P(\underline{r} \mid \underline{x}, \Delta)$ is the spin displacement PDF, giving the probability of a spin initially at position \underline{x} having diffused by a displacement \underline{r} after a diffusion time Δ (see equation 3.10). Although the spin displacement PDF provides a complete description of the diffusing system, q -space imaging is very difficult to perform in practice, for a number of reasons. First, a very large number of measurements need to be performed to sample 3-dimensional q -space adequately. For example, to obtain a $10 \times 10 \times 10$ spin displacement map, 500 measurements need to be performed (only half of the number of samples are needed due to the symmetry of the spin displacement PDF) [84]. Second, to avoid truncation artefacts, q -space needs to be sampled up to large q -values, whilst satisfying the narrow-pulse approximation [85]. This requires extremely short and intense gradient pulses, which are not currently feasible on clinical systems.

Wedge *et al* [84] have recently proposed a technique based on q -space imaging, dubbed *diffusion spectroscopic imaging* (DSI), that ignores the narrow-pulse approximation constraint. In this case, the spin displacement measured is the difference between the *average* position of the spin during the first pulse and its *average* position during the second pulse [86]. If the diffusion is restricted, the average spin positions during each pulse will tend to move towards the centre of the restricting compartment, and hence closer together, as the pulse duration increases. The spin displacement will therefore be under-estimated using long gradient pulse durations. Nevertheless, this technique is capable of showing structures in areas of crossing fibres that are not visible using diffusion tensor imaging. However, it suffers from an impractical imaging time (up to 4 hours), making it unsuitable for clinical use.

HIGH ANGULAR RESOLUTION IMAGING

In tractography, it is the orientation of the underlying tract that is of interest. It may therefore be sufficient to estimate the *angular* distribution of the spin displacement PDF. Tuch *et al* [87] proposed a technique that samples the diffusion-weighted signal along a large number of diffusion gradient orientations, using a constant b -value. This technique also shows structures that cannot be seen using conventional diffusion tensor imaging, and has the advantage of a shorter imaging time than its q -space counterpart.

Alexander *et al* [78] and Frank *et al* [77] have suggested that the signal acquired using such a technique could be fitted using the set of even spherical harmonic functions. Statistical tests can then be performed to highlight areas of the sample where the tensor model does not fit the data (the tensor model is equivalent to truncating the spherical harmonic series to 2nd order).

So far, high angular resolution data have been acquired using schemes that have been optimised for the acquisition of diffusion tensor data. Further work is needed to optimise the acquisition scheme for this particular technique. In addition, it remains unclear how the ADC profile acquired using high angular resolution imaging relates to the underlying fibre architecture. A number of models have been proposed [88, 59], and further work is needed to extract the fibre orientation from such data.

3.3.6 UNRESOLVED ISSUES IN TRACTOGRAPHY

To date, a number of techniques have been proposed to perform fibre-tracking using DTI data. However, these algorithms are difficult to validate *in-vivo* due to the lack of a gold standard. It has been shown to be feasible to use manganese ions (Mn^{2+}) as a neuronal tract tracer [63, 89]. However, it has so far only been possible to investigate tracts that can be accessed easily (such as the optic and olfactory tracts). Moreover, due to the toxicity of manganese, it is not possible to use this technique on the human brain *in-vivo*. Another common method to provide evidence for the validity of an algorithm is to compare the tracks generated with known anatomy. However, this does not provide any indication of the robustness of the algorithm to, for example, noise or partial volume effects, or to the properties of the underlying tract, such as cross-sectional radius, curvature, anisotropy or the presence of pathology.

Existing techniques therefore need to be further assessed and validated. Furthermore, existing algorithms suffer from various limitations inherent in their design, and there is a need for the development of new, more robust techniques. Some of these issues are the subject of chapters 5 and 6.

REFERENCES

- [1] CRANK, J. *The mathematics of diffusion*. Oxford University Press, Oxford, UK, 1975.
- [2] LE BIHAN, D., AND BASSER, P. J. *Molecular Diffusion and nuclear magnetic resonance*. In Le Bihan [90], 1995, chapter 1.
- [3] EINSTEIN, A. *Investigation on the theory of the Brownian movement*. Dover, New York, 1926.
- [4] PAULSON, O. B., HERTZ, M. M., BOLWIG, T. G., AND LASSEN, N. A. Filtration and diffusion of water across the blood-brain-barrier in man. *Microvascular Research* 13 (1977), 113–124.

- [5] NICHOLSON, C., AND PHILLIPS, J. M. Ion diffusion modified by tortuosity and volume fraction in the extracellular microenvironment of the rat cerebellum. *Journal of Physiology* 321 (1981), 225–257.
- [6] STEJSKAL, E. O., AND TANNER, J. E. Spin diffusion measurement: spin echoes in the presence of a time-dependent field gradient. *Journal of Chemical Physics* 42 (1965), 288–292.
- [7] CALLAGHAN, P. T. *The measurement of motion using spin echoes*. In [91], 1991, chapter 6.
- [8] NEEMAN, M., FREYER, J. P., AND SILLERUD, L. O. *b factor: effects of imaging gradients on diffusion measurements by MR imaging*. In Le Bihan [90], 1995, chapter 5.
- [9] BEAULIEU, C. The basis of anisotropic water diffusion in the nervous system – a technical review. *NMR in Biomedicine* 15 (2002), 435–455.
- [10] NORRIS, D. G. The effects of microscopic tissue parameters on the diffusion weighted magnetic resonance imaging experiment. *NMR in Biomedicine* 14 (2001), 77–93.
- [11] THOMAS, D. L., LYTHGOE, M. F., PELL, G. S., CALAMANTE, F., AND ORDIDGE, R. J. The measurement of diffusion and perfusion in biological systems using magnetic resonance imaging. *Physics in Medicine and Biology* 45 (2000), R97–R138.
- [12] PFEUFFER, J., FLÖGEL, U., DREHER, W., AND LEIBFRITZ, D. Restricted diffusion and exchange of intracellular water: theoretical modelling and diffusion time dependence of ^1H NMR measurements on perfused glial cells. *NMR in Biomedicine* 11 (1998), 19–31.
- [13] VON MEERWALL, E. D. Interpreting pulsed gradient spin echo diffusion experiments in polydisperse specimens. *Journal of Magnetic Resonance* 50 (1982), 409–416.
- [14] BUCKLEY, D. L., BUI, J. D., PHILLIPS, M. I., ZELLES, T., INGLIS, B. A., PLANT, H. D., AND BLACKBAND, S. J. The effect of ouabain on water diffusion in the rat hippocampal slice measured by high resolution NMR imaging. *Magnetic Resonance in Medicine* 41 (1999), 137–142.
- [15] LATOUR, L. L., SVOBODA, K., MITRA, P. P., AND SOTAK, C. H. Time-dependent diffusion of water in a biological model system. *Proceedings of the National Academy of Sciences* 91 (1994), 1229–1233.
- [16] LE BIHAN, D. Molecular diffusion, tissue microdynamics and microstructure. *NMR in Biomedicine* 8 (1995), 375–386.
- [17] TOPGAARD, D., AND SÖDERMAN, O. Experimental determination of pore shape and size using q -space NMR microscopy in the long diffusion-time limit. *Magnetic Resonance Imaging* 21 (2003), 69–76.
- [18] CLEVELAND, G. G., CHANG, D. C., HAZLEWOOD, C. F., AND RORSCHACH, H. E. Nuclear magnetic resonance measurement of skeletal muscle: anisotropy of the diffusion coefficient of the intracellular water. *Biophysics Journal* 16 (1976), 1043–1053.
- [19] MOSELEY, M. E., COHEN, Y., KUCHARZYK, J., MINTOROVITCH, J., ASQARI, H. S., WENDLAND, M. F., TSURUDA, J., AND NORMAN, D. Diffusion-weighted MR imaging of anisotropic water diffusion in cat central nervous system. *Radiology* 176 (1990), 439–445.
- [20] BASSER, P. J., MATTIELLO, J., AND LE BIHAN, D. MR diffusion tensor spectroscopy and imaging. *Biophysics Journal* 66 (1994), 259–267.
- [21] BASSER, P. J., MATTIELLO, J., AND LE BIHAN, D. Estimation of the effective self-diffusion *tensor* from the NMR spin echo. *Journal of Magnetic Resonance B* 103 (1994), 247–254.
- [22] BLOCH, F. Nuclear induction. *Physical Review* 70 (1946), 460–474.
- [23] TORREY, H. C. Bloch equations with diffusion terms. *Physical Review* 104 (1956), 563–565.
- [24] MORI, S., AND VAN ZIJL, P. C. M. Diffusion weighting by the trace of the diffusion tensor within a single scan. *Magnetic Resonance in Medicine* 33 (1995), 41–52.
- [25] WONG, E. C., COX, R. W., AND SONG, A. W. Optimized isotropic diffusion weighting. *Magnetic Resonance in Medicine* 34 (1995).

- [26] MATTIELLO, J., BASSER, P. J., AND LE BIHAN, D. Analytical expressions for the b matrix in NMR diffusion imaging and spectroscopy. *Journal of Magnetic Resonance A* 108 (1994), 131–141.
- [27] MATTIELLO, J., BASSER, P. J., AND LE BIHAN, D. The b matrix in diffusion tensor echo-planar imaging. *Magnetic Resonance in Medicine* 37 (1997), 292–300.
- [28] JONES, D. K., HORSFIELD, M. A., AND SIMMONS, A. Optimal strategies for measuring diffusion in anisotropic systems by magnetic resonance imaging. *Magnetic Resonance in Medicine* 42 (1999), 515–525.
- [29] VAN GELDEREN, P., DE VLEESCHOUWER, M. H. M., DESPRES, D., PEKAR, J., VAN ZIJL, P. C. M., AND MOONEN, C. T. W. Water diffusion and acute stroke. *Magnetic Resonance in Medicine* 31 (1994), 154–163.
- [30] TANNER, J. E., AND STEJSKAL, E. O. Restricted self-diffusion of protons in colloidal systems by the pulsed-gradient, spin-echo method. *Journal of Chemical Physics* 49 (1968), 1768–1777.
- [31] CALLAGHAN, P. T. *Introductory nuclear magnetic resonance*. In [91], 1991, chapter 2.
- [32] REESE, T. G., HEID, O., WEISSKOFF, R. M., AND WEDEEN, V. J. Reduction of eddy-current-induced distortion in diffusion MRI using a twice-refocused spin echo. *Magnetic Resonance in Medicine* 49 (2003), 177–182.
- [33] ANDERSON, A. W., AND GORE, J. C. Analysis and correction of motion artifacts in diffusion weighted imaging. *Magnetic Resonance in Medicine* 32 (1994), 379–387.
- [34] ORDIDGE, R. J., HELPERN, J. A., QING, Z. X., KNIGHT, R. A., AND NAGESH, V. Correction of motional artifacts in diffusion-weighted MR images using navigator echoes. *Magnetic Resonance Imaging* 12 (1994), 455–460.
- [35] BUTTS, K., PAULY, J., DE CRESPIGNY, A., AND MOSELEY, M. Isotropic diffusion-weighted and spiral-navigated interleaved EPI for routine imaging of acute stroke. *Magnetic Resonance in Medicine* 38 (1997), 741–749.
- [36] NGUYEN, Q., CLEMENCE, M., AND ORDIDGE, R. J. The use of intelligent re-acquisition to reduce scan time in MRI degraded by motion. In *Proceedings of the ISMRM* (1998), vol. 6, p. 134.
- [37] SKARE, S., AND ANDERSSON, J. L. R. On the effects of gating in diffusion imaging of the brain using single shot EPI. *Magnetic Resonance Imaging* 19 (2001), 1125–1128.
- [38] JEZZARD, P., BARNETT, A. S., AND PIERPAOLI, C. Characterization of and correction for eddy current artifacts in echo planar diffusion imaging. *Magnetic Resonance in Medicine* 39 (1998), 801–812.
- [39] BASSER, P. J., AND JONES, D. K. Diffusion-tensor MRI: theory, experimental design and data analysis – a technical review. *NMR in Biomedicine* 15 (2002), 456–467.
- [40] BOESCH, C., GRUETTER, R., AND MARTIN, E. Temporal and spatial analysis of fields generated by eddy-currents in superconducting magnets: optimization of corrections and quantitative characterization of magnet/gradient systems. *Magnetic Resonance in Medicine* 20 (1991), 268–284.
- [41] CALAMANTE, F., PORTER, D. A., GADIAN, D. G., AND CONNELLY, A. Correction for eddy current induced B_0 shifts in diffusion-weighted echo-planar imaging. *Magnetic Resonance in Medicine* 41 (1999), 95–102.
- [42] HASELGROVE, J. C., AND MOORE, J. R. Correction for distortion of echo-planar images used to calculate the apparent diffusion coefficient. *Magnetic Resonance in Medicine* 36 (1996), 960–964.
- [43] ANDERSSON, J. L. R., AND SKARE, S. A model-based method for retrospective correction of geometric distortions in diffusion-weighted EPI. *NeuroImage* 16 (2002), 177–199.
- [44] LE BIHAN, D., MANGIN, J.-F., POUPEON, C., CLARCK, C. A., PAPPATA, S., MOLKO, N., AND CHABRIAT, H. Diffusion tensor imaging: concepts and application. *Journal of Magnetic Resonance Imaging* 13 (2001), 534–546.

[45] DE CRESPIGNY, A. J., MARKS, M. P., ENZMANN, D. R., AND MOSELEY, M. E. Navigated diffusion imaging of normal and ischemic human brain. *Magnetic Resonance in Medicine* 33 (1995), 720–728.

[46] BEAULIEU, C., AND ALLEN, P. S. Determinants of anisotropic water diffusion in nerves. *Magnetic Resonance in Medicine* 31 (1994), 394–400.

[47] ARMITAGE, P. A., AND BASTIN, M. E. Selecting an appropriate anisotropy index for displaying diffusion tensor imaging data with improved contrast and sensitivity. *Magnetic Resonance in Medicine* 44 (2000), 117–121.

[48] PAPADAKIS, N. G., XING, D., HOUSTON, G. C., SMITH, J. M., SMITH, M. I., JAMES, M. F., PARSONS, A. A., HUANG, C. L.-H., AND HALL, LAURANCE D. CARPENTER, T. A. A study of rotationally invariant and symmetric indices of diffusion anisotropy. *Magnetic Resonance Imaging* 17 (1999), 881–892.

[49] PIERPAOLI, C., AND BASSER, P. J. Towards a quantitative assessment of diffusion anisotropy. *Magnetic Resonance in Medicine* 36 (1996), 893–906.

[50] BASSER, P. J., AND PIERPAOLI, C. Microstructural and physiological features of tissues elucidated by quantitative-diffusion-tensor MRI. *Journal of Magnetic Resonance B* 111 (1996), 209–219.

[51] CONTURO, T. E., MCKINSTRY, R. C., AKBUDAK, E., AND ROBINSON, B. H. Encoding of anisotropic diffusion with tetrahedral gradients: a general mathematical diffusion formalism and experimental results. *Magnetic Resonance in Medicine* 35 (1996), 399–412.

[52] FILIPPI, M., CERCIGNANI, M., INGLESE, M., HORSFIELD, M. A., AND COMI, G. Diffusion tensor magnetic resonance imaging in multiple sclerosis. *Neurology* 56 (2001), 304–311.

[53] SHIMONY, J. S., MCKINSTRY, R. C., AKBUDAK, E., ARONOVITZ, J. A., SNYDER, A. Z., LORI, N. F., CULL, T. S., AND CONTURO, T. E. Quantitative diffusion-tensor anisotropy brain MR imaging: normative human data and anatomical analysis. *Radiology* 212 (1999), 770–784.

[54] PIERPAOLI, C., JEZZARD, P., BASSER, P. J., BARNETT, A., AND DI CHIRO, G. Diffusion tensor MR imaging of the human brain. *Radiology* 201 (1996), 637–648.

[55] PAJEVIC, S., AND PIERPAOLI, C. Color schemes to represent the orientation of anisotropic tissues from diffusion tensor data: application to white matter tract mapping in the human brain. *Magnetic Resonance in Medicine* 42 (1999), 526–540.

[56] PAPADAKIS, N. G., MURRILLS, C. D., HALL, L. D., HUANG, C. L.-H., AND CARPENTER, A. A. Minimal gradient encoding for robust estimation of diffusion anisotropy. *Magnetic Resonance Imaging* 18 (2000), 671–679.

[57] ARFANAKIS, K., CORDES, D., HAUGHTON, V. M., CAREW, J. D., AND MEYERAND, M. E. Independent component analysis applied to diffusion tensor MRI. *Magnetic Resonance in Medicine* 47 (2002), 354–363.

[58] HASAN, K. M., PARKER, D. L., AND ALEXANDER, A. L. Comparison of gradient encoding schemes for diffusion-tensor MRI. *Journal of Magnetic Resonance Imaging* 13 (2001), 769–780.

[59] TUCH, D. S., REESE, T. G., WIEGELL, M. R., MAKRIS, N., BELLIVEAU, J. W., AND WEDEEN, V. J. High angular resolution diffusion imaging reveals intravoxel white matter fiber heterogeneity. *Magnetic Resonance in Medicine* 48 (2002), 577–582.

[60] SKARE, S., HEDEHUS, M., MOSELEY, M. E., AND LI, T.-Q. Condition number as a measure of noise performance of diffusion tensor data acquisition schemes with MRI. *Journal of Magnetic Resonance* 147 (2000), 340–352.

[61] BATCHELOR, P. G., ATKINSON, D., HILL, D. L. G., CALAMANTE, F., AND CONNELLY, A. Anisotropic noise propagation in diffusion tensor MRI sampling schemes. *Magnetic Resonance in Medicine* (2003), *in press*.

[62] MORI, S., AND VAN ZIJL, P. C. M. Fiber tracking: principles and strategies – a technical review. *NMR in Biomedicine* 15 (2002), 468–480.

[63] LIN, C. P., Y., T. W., CHENG, H. C., AND CHEN, J. H. Validation of diffusion tensor magnetic resonance axonal fiber imaging with registered manganese-enhanced optic tracts. *NeuroImage* 14 (2001), 1035–1047.

[64] JONES, D. K., SIMMONS, A., WILLIAMS, S. C. R., AND HORSFIELD, M. A. Non-invasive assessment of axonal fiber connectivity in the human brain via diffusion tensor MRI. *Magnetic Resonance in Medicine* 42 (1999), 37–41.

[65] MORI, S., CRAIN, B. J., CHACKO, V. P., AND VAN ZIJL, P. C. M. Three-dimensional tracking of axonal projections in the brain by magnetic resonance imaging. *Annals of Neurology* 45 (1999), 265–269.

[66] CONTURO, T. E., LORI, N. F., CULL, T. S., AKBUDAK, E., SNYDER, A. Z., SHIMONY, J. S., MCKINSTRY, R. C., BURTON, H., AND RAICHLE, M. E. Tracking neuronal fiber pathways in the living human brain. *Proceedings of the National Academy of Sciences* 96 (1999), 10422–10427.

[67] PARKER, G. J. M., WHEELER-KINGSHOTT, C. A. M., AND BARKER, G. J. Estimating distributed anatomical connectivity using fast marching methods and diffusion tensor imaging. *IEEE Transactions on Medical Imaging* 21 (2002), 505–512.

[68] BASSER, P. J., PAJEVIC, S., PIERPAOLI, C., DUDA, J., AND ALDROUBI, A. In vivo fiber tractography using DT-MRI data. *Magnetic Resonance in Medicine* 44 (2000), 625–632.

[69] KOCH, M. A., NORRIS, D. G., AND HUND-GEORGIADIS, M. An investigation of functional and anatomical connectivity using magnetic resonance imaging. *NeuroImage* 16 (2002), 241–250.

[70] POUPON, C., CLARK, C. A., FROUIN, V., RÉGIS, J., BLOCH, I., LE BIHAN, D., AND MANGIN, J.-F. Regularization of diffusion-based direction maps for the tracking of brain white matter fascicles. *NeuroImage* 12 (2000), 184–195.

[71] JONES, D. K. Determining and visualizing uncertainty in estimates of fiber orientation from diffusion tensor MRI. *Magnetic Resonance in Medicine* 49 (2003), 7–12.

[72] ANDERSON, A. W. Theoretical analysis of the effects of noise on diffusion tensor imaging. *Magnetic Resonance in Medicine* 46 (2001), 1174–1188.

[73] BASTIN, M. E., ARMITAGE, P. A., AND MARSHALL, I. A theoretical study of the effect of experimental noise on the measurement of anisotropy in diffusion imaging. *Magnetic Resonance Imaging* 16 (1998), 773–385.

[74] SKARE, S., LI, T.-Q., NORDELL, B., AND INGVAR, M. Noise considerations in the determination of diffusion tensor anisotropy. *Magnetic Resonance Imaging* 18 (2000), 659–669.

[75] BASSER, P. J., AND PAJEVIC, S. Statistical artifacts in diffusion tensor MRI (DT-MRI) caused by background noise. *Magnetic Resonance in Medicine* 44 (2000), 41–50.

[76] ALEXANDER, A. L., HASAN, K. M., LAZAR, M., TSURUDA, J. S., AND PARKER, D. L. Analysis of partial volume effects in diffusion-tensor MRI. *Magnetic Resonance in Medicine* 45 (2001), 770–780.

[77] FRANK, L. R. Characterisation of anisotropy in high angular resolution diffusion-weighted MRI. *Magnetic Resonance in Medicine* 47 (2002), 1083–1099.

[78] ALEXANDER, D. C., BARKER, G. J., AND ARRIDGE, S. R. Detection and modelling of non-Gaussian apparent diffusion coefficient profiles in human brain data. *Magnetic Resonance in Medicine* 48 (2002), 331–340.

[79] TOURNIER, J.-D., CALAMANTE, F., KING, M. D., GADIAN, D. G., AND CONNELLY, A. Limitations and requirements of diffusion tensor fiber tracking: an assessment using simulations. *Magnetic Resonance in Medicine* 47 (2002), 701–708.

[80] LORI, N. F., AKBUDAK, E., SHIMONY, J. S., CULL, T. S., SNYDER, A. Z., GUILLORY, R. K., AND CONTURO, T. E. Diffusion tensor fiber tracking of human brain connectivity: acquisition methods, reliability analysis and biological results. *NMR in Biomedicine* 15 (2002), 493–515.

- [81] FRANK, L. R. Anisotropy in high angular resolution diffusion weighted MRI. *Magnetic Resonance in Medicine* 45 (2001), 935–939.
- [82] PIERPAOLI, C., BARNETT, A., PAJEVIC, S., CHEN, R., PENIX, L., VIRTÀ, A., AND BASSER, P. Water diffusion changes in Wallerian degeneration and their dependence on white matter architecture. *NeuroImage* 13 (2001), 1174–1185.
- [83] CORY, D. G., AND GARROWAY, A. N. Measurement of translational displacement probabilities by NMR: an indicator of compartmentation. *Magnetic Resonance in Medicine* 14 (1990), 435–444.
- [84] WEDEEN, V. J., REESE, T. G., TUCH, D. S., WEIGEL, M. R., DOU, R. M., WEISKOFF, R. M., AND CHESSLER, D. Mapping fibre orientation spectra in cerebral white matter with Fourier-transform diffusion MRI. In *Proceedings of the ISMRM* (2000), vol. 8, p. 82.
- [85] BASSER, P. J. Relationships between diffusion tensor and q -space MRI. *Magnetic Resonance in Medicine* 47 (2002), 392–397.
- [86] MITRA, P. P., AND HALPERIN, B. I. Effects of finite gradient-pulse widths in pulsed-gradient diffusion measurements. *Journal of Magnetic Resonance A* 113 (1995), 94–101.
- [87] TUCH, D. S., WEISSKOFF, R. M., BELLIVEAU, J. W., AND WEDEEN, V. J. High angular resolution diffusion imaging of the human brain. In *Proceedings of the ISMRM* (1999), vol. 7, p. 321.
- [88] JANSONS, K. J., AND ALEXANDER, D. C. Spin echo attenuation to diffusion displacement density: a general inversion for measurements on a sphere. In *Proceedings of the ISMRM* (2002), vol. 10, p. 1171.
- [89] PAUTLER, R. G., SILVA, A. C., AND KORETSKY, A. P. *In vivo* neuronal tract tracing using manganese-enhanced magnetic resonance imaging. *Magnetic Resonance in Medicine* 40 (1998), 740–748.
- [90] LE BIHAN, D., Ed. *Diffusion and perfusion magnetic resonance imaging*. Raven Press, Ltd., New York, 1995.
- [91] CALLAGHAN, P. T. *Principles of nuclear magnetic resonance microscopy*. Oxford University Press, Oxford, UK, 1991.

CHAPTER 4

IMAGING SEQUENCE IMPLEMENTATION

CONTENTS

4.1	Tetra-ortho sequence	79
4.2	20 direction sequence	81
4.2.1	Optimisation of diffusion gradient timing	82
4.2.2	Remaining eddy-current artefact reduction	84
4.2.3	Offline reconstruction	86
References		93

The images presented in this work have been acquired using one of two sequences, both of which are described in this chapter. The sequence initially used to acquire diffusion tensor data was based heavily on an existing sequence, written in-house for the acquisition of the trace of the diffusion, which was extended to acquire images according to the tetrahedral-orthogonal diffusion encoding scheme [1]. This sequence, described in the first section, produced images of sufficient quality for the calculation of anisotropy maps and for simple tractography applications.

This sequence, however, suffered from a number of drawbacks. The scanner's pulse programming environment has a number of limitations, which in our case affected the efficiency of data collection. In particular, it was not possible to write the sequence in a single pulse program, and a pulse program cannot be started until all calculations for the previous one have been completed. For reasons outlined below, a constant phase-encoding (CPE) EPI read-out was used, so the calculations required for image reconstruction took a relatively long time to complete (see section 2.2.2). Since the sequence had to be implemented using several pulse programs, this introduced a significant time lapse between starting successive programs (approximately two minutes for 25 128×128 slices). Also, it has recently been suggested that applying the diffusion encoding in a large number of

directions uniformly distributed over the surface of a hemisphere may reduce the bias in the calculation of the tensor and its anisotropy [2]. For these reasons, a new sequence was developed that would acquire data in a much more time efficient manner, using a 20 direction diffusion encoding scheme. It is described in the second section.

Different diffusion-encoding gradients will in general induce different eddy-currents, and hence different receiver delays. With a blipped EPI read-out, the receiver delay would need to be adjusted separately for every gradient orientation. However, the receiver gain needs to be kept constant over all the pulse programs if the data are to be consistently scaled. On the Siemens Vision system used in this work, adjusting the receiver delay and gain separately is impractical. In addition, the adjustments introduce extra ‘dead time’ into the acquisition. However, these adjustments only need to be performed once when using a CPE-EPI read-out. For this reason, a CPE-EPI read-out was used for both sequences.

4.1 TETRA-ORTHO SEQUENCE

The DTI sequence initially used was a simple modification of that routinely used to collect average ADC maps in the clinical examinations at Great Ormond Street Hospital. This consisted of a spin-echo EPI sequence, with two diffusion gradient pulses placed either side of the refocusing pulse (see figure 4.1). Due to the limitations described above, the sequence was split into two pulse programs, to be run sequentially. A correction scan consisting of two lines through the centre of k -space was acquired immediately after the 90° excitation pulse to correct for Nyquist ghosting artefacts (see section 2.3.2). In order to alleviate the effects of eddy-currents induced by the diffusion-encoding gradients, two opposite polarity dummy gradients were placed one immediately before the 90° excitation pulse, the other after the EPI read-out, such that they would not affect the image contrast. The remaining eddy-current-induced artefact (a slice-dependent image shift) was corrected by acquiring an additional data set without phase-encoding [3].

The original sequence for the acquisition of ADC maps consisted of acquiring images with the diffusion encoding applied along three orthogonal directions (x , y , z) and one $b=0$ image. As described in section 3.3.2, a full characterisation of the diffusion tensor requires the acquisition of at least 6 images with the diffusion-weighting applied in non-collinear directions, in addition to the $b=0$ image. The tetrahedral-orthogonal diffusion encoding scheme was used in this sequence, as it has been shown to perform well for the acquisition of diffusion tensor data [1]. The diffusion-encoding directions used for each pulse program are shown in table 4.1. The imaging parameters used in the acquisition were the following: TE = 110 ms, matrix = 128×128 , FOV = 256 mm, slice

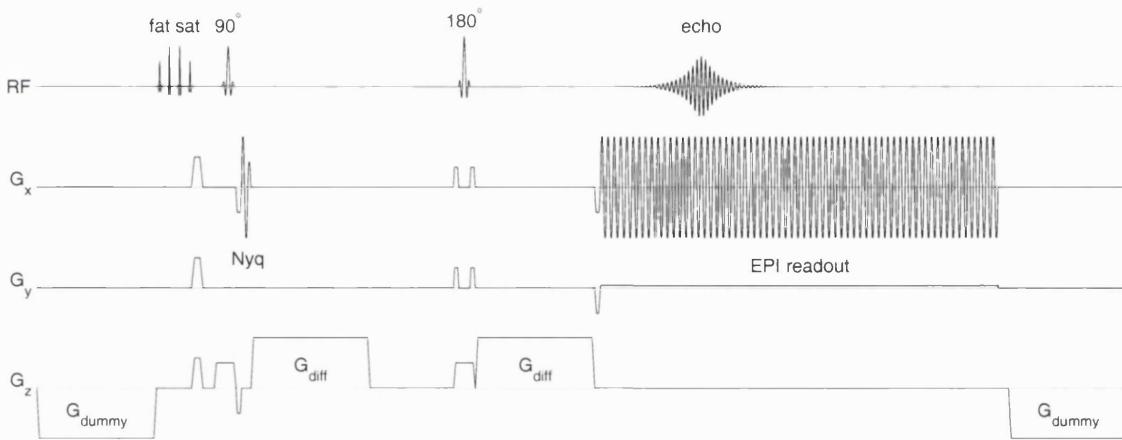


Figure 4.1: The tetrahedral-orthogonal pulse sequence implemented in Great Ormond Street Hospital. This consists of a simple PGSE sequence (see section 3.2.2), with two additional negative polarity ‘dummy’ gradients at either end of the sequence to alleviate the effects of eddy-currents. A set of binomial fat saturation pulses are applied prior to the 90° RF excitation pulse to selectively excite and spoil the lipid signal. A set of 2 reference scans (Nyq) are acquired immediately after the 90° excitation pulse to correct Nyquist ghost artefacts. The EPI read-out consists of a constant phase-encoding gradient sequence (see section 2.2.2).

thickness = 2 mm (resulting in a $2 \times 2 \times 2$ mm voxel size), with the slices acquired in interleaved order. The diffusion encoding gradients were applied with $\delta = 27.1$ ms and $\Delta = 52.3$ ms, giving a b -value of 1066 s/mm 2 . The measurements were repeated 6 times to give a final SNR of approximately 32 in the non-diffusion-weighted ($b=0$) image. This resulted in a total acquisition time of approximately 45 minutes for 25 contiguous slices covering a 5 cm slab, 28 minutes of which were taken up by image reconstruction and storage.

As can be seen from figure 4.2, the images produced were of a sufficient quality for the production of suitable fractional anisotropy maps and fibre-tracking (see chapter 5). However, as implemented, the tetra-ortho sequence suffered from a number of limitations that prompted the development of an improved sequence. In particular, the scan time was

ortho	tetra
0, 0, 0	$\frac{1}{\sqrt{3}} (1, 1, 1)$
1, 0, 0	$\frac{1}{\sqrt{3}} (-1, 1, 1)$
0, 1, 0	$\frac{1}{\sqrt{3}} (1, -1, 1)$
0, 0, 1	$\frac{1}{\sqrt{3}} (-1, -1, 1)$

Table 4.1: the diffusion-encoding gradient direction vectors used for the tetrahedral-orthogonal sequence.

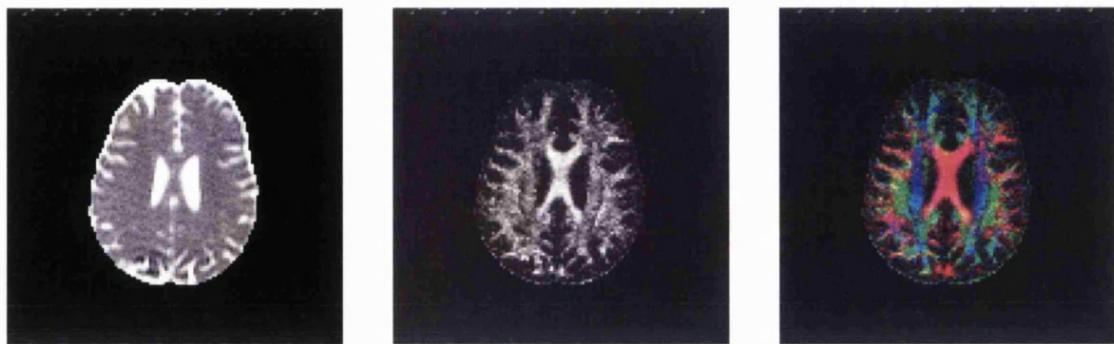


Figure 4.2: maps calculated from DTI data acquired on a healthy volunteer using the tetrahedral-orthogonal sequence. Left: mean diffusivity map. Middle: fractional anisotropy map. Right: colour-coded major eigenvector map.

too long to be practical, and the number of diffusion-encoded directions was limited to 7.

4.2 20 DIRECTION SEQUENCE

The tetra-ortho sequence relied on the acquisition of an additional set of data to correct for eddy-currents distortions. However, it has been shown that these effects can be significantly reduced by the use of a twice-refocused spin-echo EPI sequence [4], which may eliminate the need for such a correction. Assuming that the remaining eddy-current artefacts were small enough to be tolerated, the data acquisition part of such a sequence could be run in half the time required for the tetra-ortho sequence, while producing images of similar quality. This approach was therefore used in the improved sequence, and the timing of the gradients was optimised as described in section 4.2.1. The sequence as implemented is shown in figure 4.3.

As for the first sequence, the sequence had to be split into a number of different pulse programs, and a CPE-EPI read-out was used. In order to minimise the amount of ‘dead time’ due to image reconstruction between successive acquisitions, the sequence was modified to output the raw data rather than the fully reconstructed images, leaving the reconstruction to be performed offline at a later stage. Note however that since it is a relatively quick step, the Nyquist ghost correction (described in section 4.2.3) was still performed on the scanner. In addition, the sequence program was modified such that data processing and storage could be performed while the acquisition was still running, rather than having to wait until all the data had been acquired. With these modifications, the dead time due to the processing of 60 interleaved 128×128 slices, acquired in 15 seconds, was reduced from 60 to 10 seconds. The implementation of the offline reconstruction is described in more detail in section 4.2.3.

The minimum echo time of the sequence is limited by the gradient strength and the

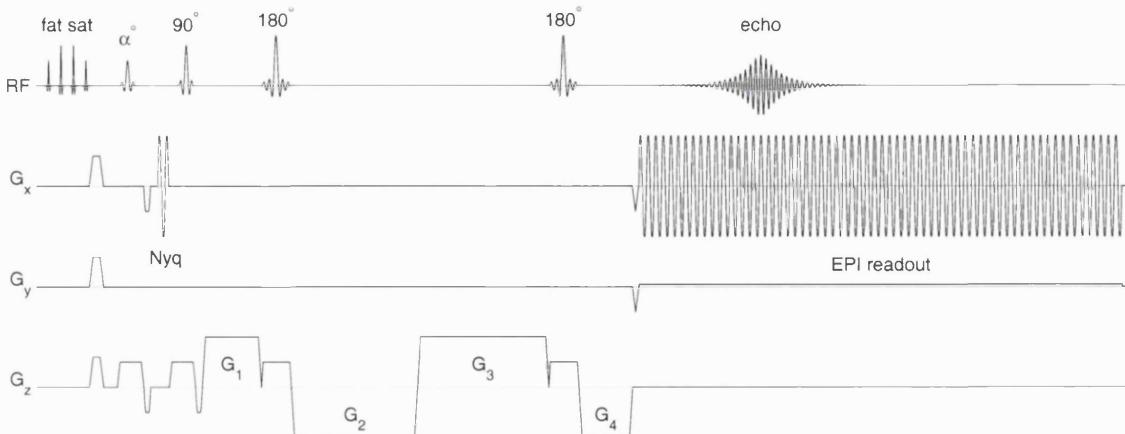


Figure 4.3: the twice-refocused diffusion-weighted spin-echo sequence implemented in Great Ormond Street Hospital. This type of sequence is less susceptible to eddy-current artefacts (see section 3.2.2). In this implementation, the gradient timings have been optimised to null eddy-currents with a time constant of the order of the echo time, while keeping the echo time as short as possible. A set of binomial fat saturation pulses are applied first to selectively excite and spoil the lipid signal. A set of 3 reference scans are acquired after an α RF pulse to correct Nyquist ghost artefacts ($\alpha = 10^\circ$). The EPI read-out consists of a constant phase-encoding gradient sequence (see section 2.2.2).

required b -value (see next section). However, the minimum TE may be reduced by acquiring the Nyquist ghost correction scans before the 90° excitation pulse, such that the diffusion gradients can be applied as soon as possible. In addition, using this approach, it is possible to acquire an extra reference scan, allowing a more robust Nyquist ghost correction (as described in section 2.3.2) without any impact on the TE. Three reference scans were therefore acquired after a 10° RF excitation pulse, immediately prior to the main 90° pulse.

Using a diffusion-encoding scheme consisting of at least 20 directions has been shown to alleviate the orientational bias present in anisotropy measurements (see section 3.3.2) [2, 5]. For this reason, the sampling scheme used for this sequence consisted of 20 directions, obtained using the electrostatic repulsion model proposed by Jones *et al* [6], in addition to 3 $b=0$ images.

With these optimisations, a data set consisting of 60 contiguous 128×128 slices, covering the whole brain, with 23 images per slice ($2 \times 2 \times 2$ mm voxel size) can be acquired in 13 minutes, to produce images such as those shown in figure 4.4.

4.2.1 OPTIMISATION OF DIFFUSION GRADIENT TIMING

The use of a twice-refocused sequence imposes a number of constraints on the gradient timings. Referring to the sequence diagram shown in figure 4.5, these can be expressed

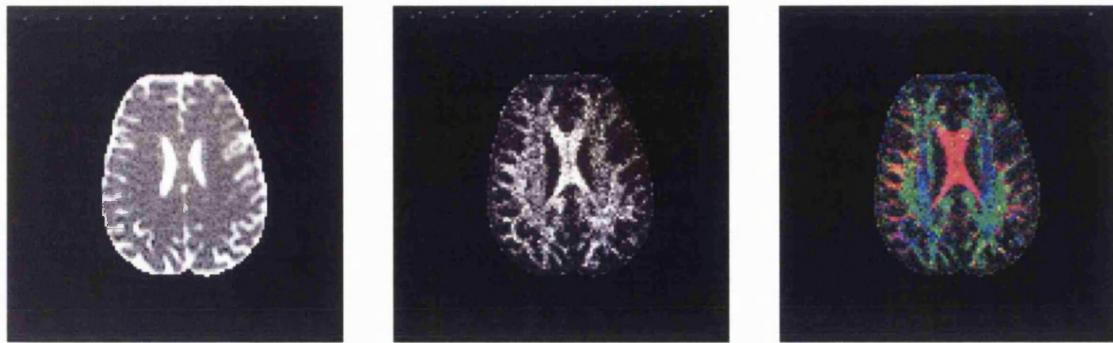


Figure 4.4: maps calculated from DTI data acquired on a healthy volunteer using the 20 direction sequence. Left: mean diffusivity map. Middle: fractional anisotropy map. Right: colour-coded major eigenvector map. Note that with this sequence 23 images are acquired per slice location, compared to 48 using the tetrahedral-orthogonal sequence.

as:

$$\begin{aligned}
 \Delta_1 + \Delta_2 + \Delta_3 &= TE && \text{total time} \\
 \Delta_1 + \Delta_3 &= \Delta_2 && \text{echo arrival} \\
 \delta_1 + \delta_2 &= \delta_3 + \delta_4 && \text{rephasing}
 \end{aligned} \tag{4.1}$$

To minimise the echo time, the delay between δ_2 and δ_3 is set to zero. This leaves one degree of freedom in the timings for these gradients, which can be exploited to minimise eddy-current artefacts. As outlined by Heid [7], the treatment of eddy-currents in diffusion imaging can be greatly simplified if the following assumptions are made: eddy-currents are produced only at gradient ramps, and decay exponentially to zero with a decay constant λ ; the effects of imaging gradients can be ignored; and all ramps used for the diffusion gradients are identical (apart from polarity) and have infinitesimal duration. The eddy-current amplitude $E(t)$ for a particular λ can then be expressed as a sum of exponential terms:

$$E(t) = \sum_i \sigma_i e^{-(t-t_i)/\lambda} \quad \text{where } \sigma_i = \pm 1 \tag{4.2}$$

where i spans over all relevant gradient ramps that occurred before t , and σ_i and t_i represent the polarity and time of onset of ramp i . This is shown schematically in figure 4.5. It has been suggested that eddy-currents with intermediate decay rates (i.e. $\lambda TE \approx 1$) are the primary source of artefacts in DWI [7], so the gradient timings should be chosen such that eddy-currents are minimised around the echo time, when data are being acquired. The eddy-current behaviour of the gradient timings can be assessed using plots of eddy-current amplitude at TE versus λTE , and the optimal combination can be found by altering them by trial and error subject to the constraints in equation 4.1. However,

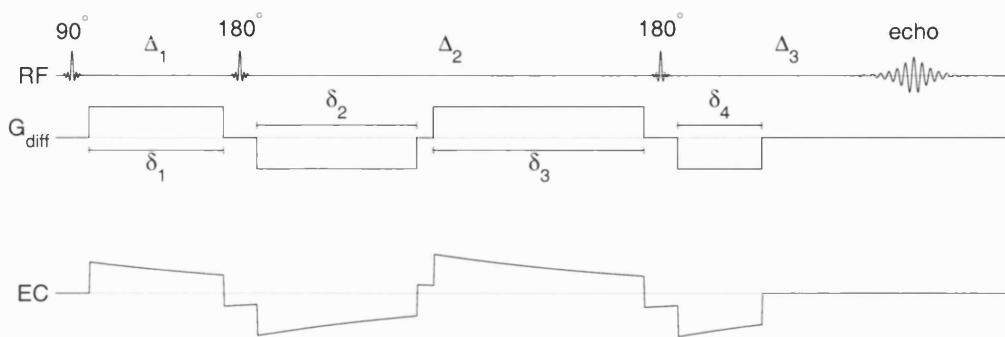


Figure 4.5: a schematic representation of the twice-refocused diffusion-weighted spin-echo sequence, illustrating the eddy-current nulling made possible by the use of this sequence. Eddy-currents (EC) are induced by the fast gradient switching that occurs during gradient ramps, and decay exponentially to zero. By choosing appropriate gradient timings, eddy-currents with a particular time constant can be effectively nulled around the echo time, when the signal is being collected.

the b -value also depends on the gradient timings, and in practice the echo time must be allowed to change such that the desired b -value can be obtained. The final gradient pulse times used were: $\delta_1 = 10.06$ ms, $\delta_2 = 22.97$ ms, $\delta_3 = 23.92$ ms, and $\delta_4 = 9.11$ ms, giving an echo time of 110 ms, a b -value of 1000 s/mm 2 , and the eddy-current behaviour shown in figure 4.6.

4.2.2 REMAINING EDDY-CURRENT ARTEFACT REDUCTION

Despite the use of a twice-refocused sequence to minimise eddy-current, there was a remaining artefact in the EPI images. This artefact could be seen on a sagittal reformat of a set of axial interleaved images acquired on a phantom, shown in figure 4.7, consistent with the presence of a long term eddy-current. Note that the images shown were acquired using the non-diffusion-encoded sequence. This artefact is therefore not due to the diffusion gradients, and is also present with identical characteristics using the diffusion-weighted sequence programs, irrespective of the direction of application of the diffusion encoding. The first few slices acquired (i.e. slices 1, 3, 5...) are clearly shifted with respect to the rest of the acquisition, and this shift tends to decrease for later slices. This is most likely due to a long-term eddy-current inducing a B_0 shift, the amplitude of which increases until a steady state is reached, such that slices acquired beyond that point are shifted by the same amount.

This artefact was removed by introducing a number of dummy loops in order to reach a steady state before data are actually acquired. Although this approach does not remove the actual shift in the images, the slices are now aligned with respect to each other. It was

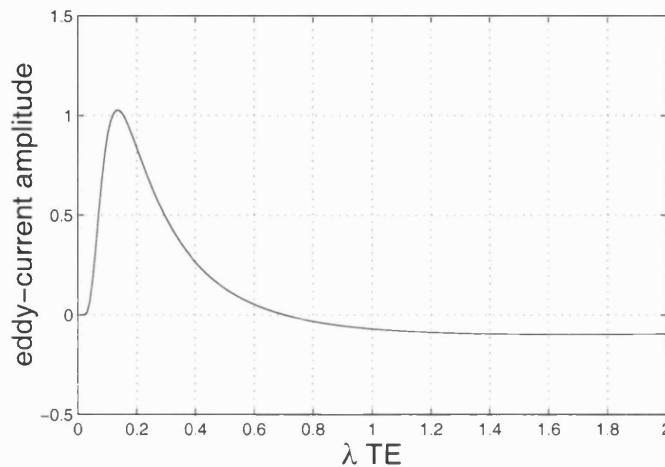


Figure 4.6: remaining eddy-current amplitude (arbitrary units) as a function of the time constant of the eddy-currents, λ (expressed in units of $1/\text{TE}$). As can be seen, residual eddy-currents with time constants around the echo time are minimised.

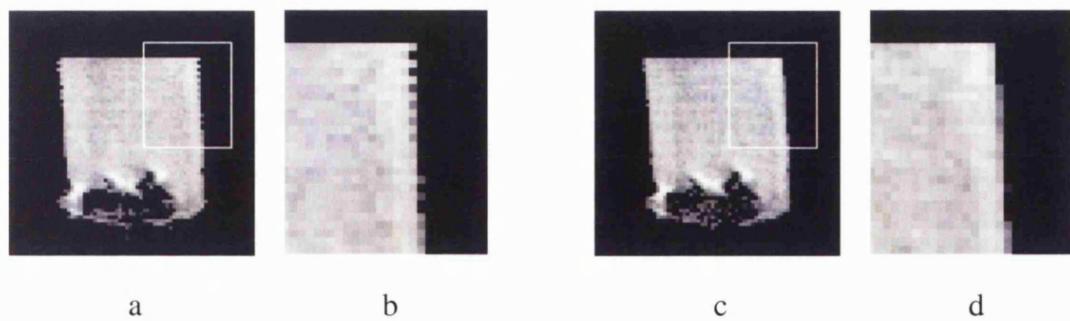


Figure 4.7: minimising the effects of long term eddy-currents using a dummy loop. (a): a sagittal slice through a set of contiguous 2D axial slices acquired on an agar phantom without a dummy loop. (b): the region highlighted in (a), showing an initial mismatch between even and odd slices. (c): the same slice acquired using a dummy loop to reach a steady state prior to data collection. (d): the region highlighted in (c), showing that the mismatch present in (b) has been reduced to a negligible level by the dummy loop.

found that using 10 such dummy loops was sufficient to effectively remove the artefact, as can be seen in figure 4.7.

The dummy loop was also useful to locate the source of the eddy-current. Gradient pulses can be selectively removed from the dummy loop to see which ones are responsible for the artefact. It was found that after removing the slice selection gradient and its refocusing lobe from the dummy loop, the artefact was present with the same intensity as in the absence of the dummy loop. Taking out other gradient pulses in the dummy loop had negligible effect on the appearance of the images. The slice selection gradient and its refocusing lobe are therefore the most likely cause of this artefact.

4.2.3 OFFLINE RECONSTRUCTION

As previously mentioned, the reconstruction was implemented offline in order to speed up the DTI scan time. Reconstructing data acquired using CPE-EPI involves the steps outlined in figure 4.8. Each of these is described in detail in the following sections. The implementation was written from scratch in C++ on a GNU/Linux workstation, and is based closely on the Siemens implementation as found on the Vision system used in this work.

NYQUIST GHOST CORRECTION

The initial step in the reconstruction of the images is the Nyquist ghost correction. Note that since this step can be performed relatively quickly on the scanner, it was not implemented offline. The algorithm as performed on the Siemens Vision is described here in detail for completeness.

As described in section 2.3.2, Nyquist ghosting arises out of inconsistencies, due to a relative shift in the time domain, and/or a global phase difference between alternate lines of k -space. Although there are a number of methods to correct these errors, this section will describe only the technique that is implemented on our Siemens Vision scanner.

A Nyquist ghost correction scan is acquired by sampling the central line of k -space 3 times, using the same frequency-encoding gradients as are used for the image data acquisition, but without any phase encoding. The timing and phase errors present in the correction scan are identical to those present in the corresponding EPI acquisition (ignoring the effects of the phase-encoding gradient). However, since all three lines correspond to the same data, inconsistencies between them can be estimated and used to correct those present in the image data. In this implementation, the third line is averaged with the first, to prevent the misinterpretation of off-resonance frequencies (see section 2.3.2 for details).

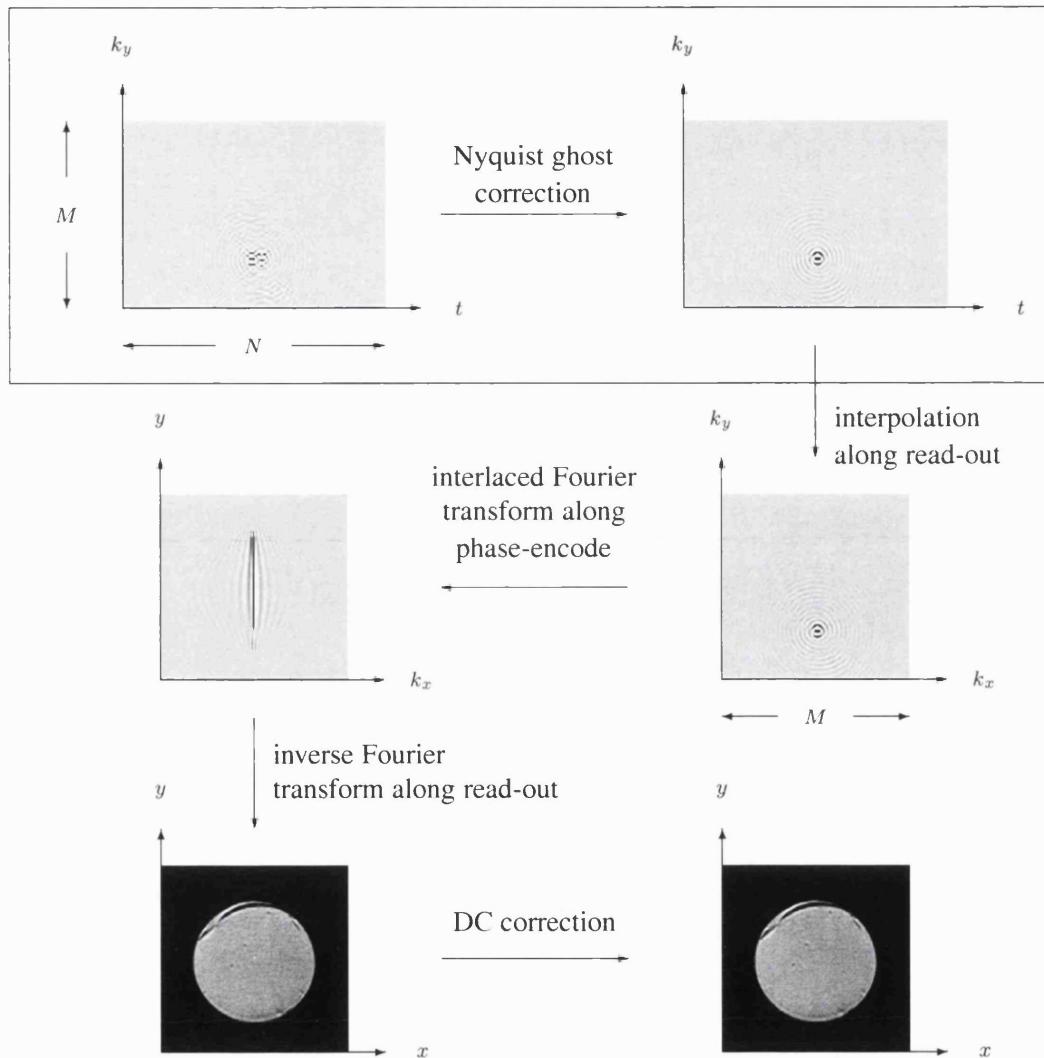


Figure 4.8: overview of the image reconstruction process. The first step after line reversal is the Nyquist ghost correction, described on page 86. Each line along the read-out direction is then resampled onto a regular grid using sinc interpolation (see page 88). The interlaced Fourier transform is then used to perform the required 1D Fourier transform along the phase-encoding direction without the need for regridding (see page 90). Finally, after a simple 1D Fourier transform along the read-out direction, the DC offset artefact is corrected as described on page 92. Note that the steps contained within the frame are performed on the scanner, while all others take place offline.

The Fourier shift Theorem states [8]:

If $f(t)$ has the Fourier transform $F(\nu)$, then $f(t - a)$ has the Fourier transform $e^{-2\pi i a \nu} F(\nu)$.

This means that a shift in the time domain translates into a linear phase term in the frequency domain. Therefore, after a one-dimensional Fourier transform along each line of k -space, the inconsistencies mentioned above should be fully characterised by a combination of constant and linear phase terms. However, under realistic imaging conditions, the data will also contain some non-linear phase errors. These additional terms can also be estimated by using all the phase information present in the corrections scans, as explained below.

Each line of the correction scan, labelled as $\underline{c}^{t,\pm}$, is weighted by a Hanning filter, then Fourier transformed to yield the corresponding $\underline{c}^{\nu,\pm}$ (\pm refers to the direction of acquisition of that particular line). The purpose of the filtering is to reduce the contribution from the highest order terms, such that very rapid phase variations are minimised. The Hanning filter used is defined as:

$$H(n) = \begin{cases} (1 + \cos(\pi(n - L_1)/S_H))/2 & 0 \leq n < L_1 \\ 1 & L_1 \leq n \leq L_2 \\ (1 + \cos(\pi(n - L_2)/S_H))/2 & L_2 < n < N \end{cases} \quad (4.3)$$

where $L_1 = S_H/2$ and $L_2 = N - S_H/2$ (as implemented, $N = 180$ and $S_H = 128$). The phase information contained in the vectors $\underline{c}^{\nu,\pm}$ can then be applied to the image data. Each line of image data along the read direction, denoted by the vector $\underline{d}^{t,\pm}$, is Fourier transformed to yield $\underline{d}^{\nu,\pm}$. Each element of $\underline{d}^{\nu,\pm}$ then has its phase altered by the negative phase of the corresponding element of $\underline{c}^{\nu,\pm}$, in the following manner:

$$d_i^{\nu,corr,\pm} = \frac{d_i^{\nu,\pm} [\underline{c}_i^{\nu,\pm}]^*}{|\underline{c}_i^{\nu,\pm}|} \quad (4.4)$$

where z^* denotes the complex conjugate of z , and $|z|$ denotes the magnitude of z . Finally, an inverse Fourier transform is applied to each $\underline{d}^{\nu,corr,\pm}$ to yield the corrected data.

INTERPOLATION ALONG READ-OUT

With a constant phase-encoded EPI read-out sequence, the data points are sampled at regular intervals in time, but not in k -space. The sampling rate along the read-out direction can be corrected using sinc interpolation, such that a sample of N data points, acquired at regular time intervals δt , can be mapped to a set of M samples, arranged at regular

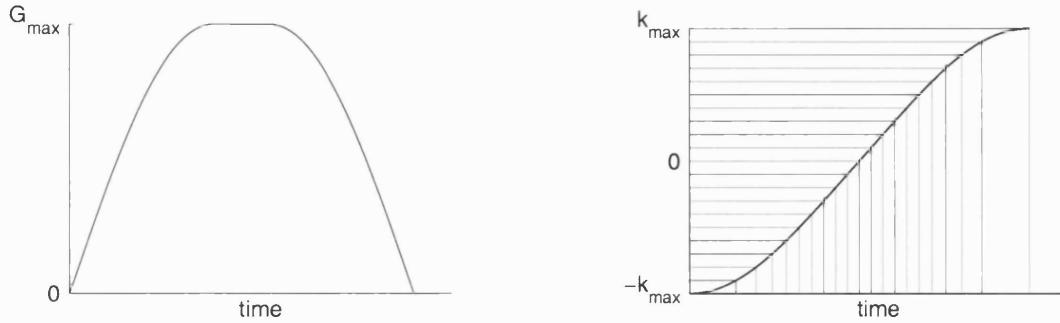


Figure 4.9: linear sampling in time means that consecutive samples are not necessarily equidistant in k -space. Left: the EPI read-out gradient waveform dictates the k -space velocity and hence the position in k -space as a function of time. Right: this can be used to calculate the positions in time of a set of points equidistant in k -space. The signal intensity at these time points can then be calculated from the time-sampled data by sinc interpolation.

k -space intervals δk along the read-out direction (figure 4.9). The interval between any two data samples that was constant in the time domain will now depend on the magnitude of the gradient at that time (see equation 2.6). For accurate resampling, the maximum interval in k -space between these samples must be smaller than the desired constant k -space interval. For this reason, the number of data points actually acquired along one line of k -space, denoted by N , will be greater than the number of data points that will eventually be regularly sampled in k -space, denoted by M (in the current implementation, $N = 180$ and $M = 128$).

Given a sample of N data points, sampled at regular time intervals δt , the intensity at an arbitrary time point τ as given by sinc interpolation is:

$$I(\tau) = \frac{\sum_{n=1}^N S(n) \operatorname{sinc}(\tau/\delta t - n)}{\sum_{n=1}^N \operatorname{sinc}(\tau/\delta t - n)} \quad \text{for } \delta t \leq \tau \leq N\delta t \quad (4.5)$$

In order to reduce the amount of computation, the summation is usually only done over the immediate neighbourhood of the point of interest. In this case, truncation artefacts need to be avoided by use of a suitable filter. Here, the sinc function is weighted by a Hamming filter, the equation of which is given below:

$$H(x) = \begin{cases} 0.54 + 0.46 \cos(\pi x/E) & |x| \leq E \\ 0 & \text{otherwise} \end{cases} \quad (4.6)$$

where E is the extent of the filter, corresponding to the number of data points either side of the point of interest to be included in the summation (in this implementation, $E = 20$).

The intensity at time τ can be more easily expressed as a matrix operation:

$$I(\tau) = [\underline{w}(\tau)]^T \cdot \underline{s}^t \quad (4.7)$$

where \underline{s}^t is the vector of length N containing the regularly time-sampled intensities, and $\underline{w}(\tau)$ is a vector, also of length N , containing the weighting factors given by:

$$\begin{aligned} w_i(\tau) &= \frac{a_i(\tau)}{\sum_{j=1}^N a_j(\tau)} \quad \text{for } 1 \leq i \leq N \\ \text{where } a_i(\tau) &= \text{sinc}(\tau/\delta t - i) H(\tau/\delta t - i) \end{aligned} \quad (4.8)$$

To map the set of N regularly time-sampled data points to M data points evenly distributed in k -space, the time of acquisition of each desired k -space data point needs to be known (figure 4.9). These times can be obtained from the waveform used to drive the read-out gradient, using equation 2.6, and stored in the *read-out raster*. This is a vector \underline{u} of length M , where each element u_i gives the time coordinate for point i at position k_i in k -space. The intensity at k_i is thus given by combining equation 4.7 with the read-out raster \underline{u} :

$$I(k_i) = \underline{w}(u_i)^T \cdot \underline{s}^t \quad (4.9)$$

Equation 4.9 gives the intensity for the point at position k_i . The vector \underline{s}^k , containing the intensities for the desired M regularly sampled data points along one line of k -space, is thus formed by setting each element \underline{s}_i^k to $I(k_i)$, such that:

$$\underline{s}^k = \mathbf{W} \cdot \underline{s}^t \quad \text{where } W_{ij} = w_j(u_i) \quad (4.10)$$

Since this operation is repeated for each frequency-encoded line of k -space, it can be generalised to:

$$\mathbf{S}^k = \mathbf{W} \cdot \mathbf{S}^t \quad (4.11)$$

where \mathbf{S}^t is the matrix formed by setting each column to the corresponding vector \underline{s}^t , and each column of \mathbf{S}^k corresponds to a line of k -space regularly sampled along the read-out direction.

THE INTERLACED FOURIER TRANSFORM

The final stage of image formation is a 2D inverse Fourier transform. However, data along the phase-encoding direction are not regularly sampled, and therefore require additional processing before the Fourier transform can be done. Although in principle these data could be resampled using interpolation as described above, it is much more efficient to use

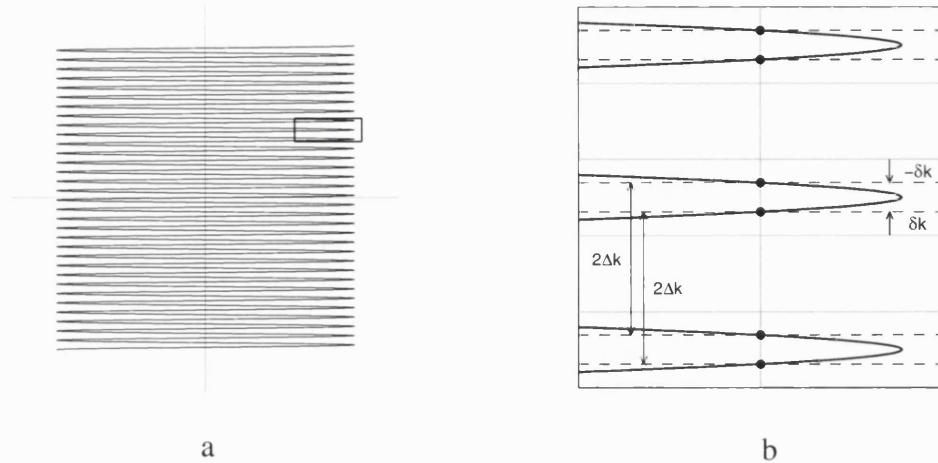


Figure 4.10: (a): a constant phase-encoded EPI trajectory; (b): the section of k -space highlighted in (a), enlarged to show a set of data points lying along the same column of k -space. Looking at the set of odd data points, it can be seen that they are regularly sampled with twice the ideal k -space sampling interval, and that they are displaced from their ideal position in k -space by δk . Similarly, the set of even data points are displaced by $-\delta k$. This property can be exploited to perform the *interlaced* Fourier transform, as described in the text.

the *interlaced* Fourier Transform, which allows a 1D Fourier transform to be performed directly on each phase-encoded line of data, subject to the conditions detailed below.

An enlargement of a constant phase-encoding gradient EPI k -space trajectory can be seen in figure 4.10, showing a single line of data along the phase-encoding direction that we are interested in reconstructing, overlaid onto the k -space trajectory. It can be seen that the set of odd data points along that line is regularly sampled in k -space, with each point being shifted along the phase-encode direction by δk relative to its ideal position for a constant sampling rate. Similarly, the even data points are shifted by $-\delta k$. The interlaced Fourier transform makes use of this property to perform a one dimensional Fourier transform, correcting for the shifts in k -space using the Fourier shift theorem (see page 88). Note that it is not possible to apply this technique to resample data in the read direction, since these cannot be broken up into regularly sampled components.

The inverse Fourier transform of \underline{p}^k , the vector containing the M data points along the particular phase-encoded line of k -space of interest, will contain significant ghosting depending on the corresponding value of δk . However, the vector $\underline{p}^{k,+}$, formed from \underline{p}^k by setting all even elements to zero, is regularly sampled. Its inverse Fourier transform $\underline{p}^{x,+}$ does not contain any ghosting, but is severely aliased, since the field of view has effectively been halved. Its phase will also be corrupted, since $\underline{p}^{k,+}$ is shifted in k -space by δk . The latter can be corrected using the Fourier shift theorem, modified for the inverse

Fourier transform:

If $F(k)$ has the inverse Fourier transform $f(x)$, then $F(k - \delta k)$ has the inverse Fourier transform $e^{+2\pi i x \delta k} f(x)$.

Therefore, the phase of $\underline{p}^{x,+}$ can be corrected by applying an adequate linear phase term:

$$\underline{p}^{x,corr,+} = \underline{\phi}^+ \diamond \underline{p}^{x,+} \quad \text{where } \phi_q^+ = e^{-2\pi i q \delta k} \quad (4.12)$$

The symbol \diamond is introduced here to denote element-by-element multiplication¹. The set of even data points can be corrected in the same way:

$$\underline{p}^{x,corr,-} = \underline{\phi}^- \diamond \underline{p}^{x,-} \quad \text{where } \phi_q^- = e^{+2\pi i q \delta k} \quad (4.13)$$

The full, corrected, inverse Fourier transformed line of k -space can then be obtained by recombining the two halves of the data:

$$\underline{p}^{x,corr} = \underline{p}^{x,corr,+} + \underline{p}^{x,corr,-} \quad (4.14)$$

This process can be applied to each line of k -space provided the corresponding shift δk is known. These can be obtained from the gradient waveforms using equation 2.6, and stored in the *phase-encode raster*. This is a vector \underline{v} of length M containing the k -space shifts corresponding to each phase-encoded line. Equations 4.12, 4.13 and 4.14 can then be combined and simplified to:

$$\mathbf{P}^{x,corr} = \Phi^+ \diamond \mathbf{P}^{x,+} + \Phi^- \diamond \mathbf{P}^{x,-} \quad \text{where } \Phi_{qr}^\pm = e^{\mp 2\pi i q v_r} \quad (4.15)$$

where $\mathbf{P}^{x,\pm}$ is the matrix obtained by setting each column to its corresponding $\underline{p}^{x,\pm}$ vector. After this process, the data set will have undergone a complete Fourier transform along the phase-encode direction, and the final image can be obtained by performing a standard Fourier transform along the read-out direction.

DC-OFFSET CORRECTION

The signal acquired during a MRI examination will typically be contaminated by some unwanted DC signal, which will give rise to an incorrect value for the central pixel (the zero frequency spot). In traditional 2D FT imaging, artefacts due to the presence of this offset can be eliminated using phase-cycling [9]. However, this option is not available

¹In other words, if $\underline{a} = \underline{b} \diamond \underline{c}$, then $a_i = b_i c_i$. This also applies to matrices: if $\mathbf{A} = \mathbf{B} \diamond \mathbf{C}$, then $A_{ij} = B_{ij} C_{ij}$.

in EPI, since a single excitation is used to acquire the whole image. In this case, the magnitude of the DC offset must be estimated to remove the artefact. This may be done in a number of ways; the method described here is based closely on the Siemens Vision implementation.

The inverse Fourier transform is performed along one of the image axes, for example, the y -axis. The x -axis now contains information in the image domain, while the y -axis contains data in the k -space domain. The DC offset artefact, which in image space was limited to the central pixel, is now present as a constant offset along the entire central line in the y direction. By averaging the two lines either side of the central line, a good estimate of the correct intensity for the central pixel can be obtained, as follows.

We define the vector \underline{b} to be the geometric average of \underline{a}^- and \underline{a}^+ , the two lines either side of the central DC-contaminated line \underline{a} , as follows:

$$b_i = \pm \sqrt{\underline{a}_i^- \times \underline{a}_i^+} \quad (4.16)$$

Since there are two possible values for each b_i , care must be taken to select the appropriate one. On an Argand diagram, b_i will be positioned such that it bisects the angle formed by \underline{a}_i^- and \underline{a}_i^+ . There are two such angles, one acute, the other obtuse. Assuming that the phase does not vary too rapidly across the image, the appropriate value of b_i should be that which bisects the *acute* angle.

The Fourier transform of \underline{b} gives an estimate of the central line free of any DC-offset, estimated from the two lines either side of it. The central pixel can therefore be replaced by the corresponding point on this corrected line. The same process can be repeated along the other axis, and the two values thus obtained can be averaged to provide a more robust estimate for the central pixel intensity.

REFERENCES

- [1] HASAN, K. M., PARKER, D. L., AND ALEXANDER, A. L. Comparison of gradient encoding schemes for diffusion-tensor MRI. *Journal of Magnetic Resonance Imaging* 13 (2001), 769–780.
- [2] PAPADAKIS, N. G., MURRILLS, C. D., HALL, L. D., HUANG, C. L.-H., AND CARPENTER, A. A. Minimal gradient encoding for robust estimation of diffusion anisotropy. *Magnetic Resonance Imaging* 18 (2000), 671–679.
- [3] CALAMANTE, F., PORTER, D. A., GADIAN, D. G., AND CONNELLY, A. Correction for eddy current induced B_0 shifts in diffusion-weighted echo-planar imaging. *Magnetic Resonance in Medicine* 41 (1999), 95–102.
- [4] REESE, T. G., HEID, O., WEISSKOFF, R. M., AND WEDEEN, V. J. Reduction of eddy-current-induced distortion in diffusion MRI using a twice-refocused spin echo. *Magnetic Resonance in Medicine* 49 (2003), 177–182.

- [5] ARFANAKIS, K., CORDES, D., HAUGHTON, V. M., CAREW, J. D., AND MEYERAND, M. E. Independent component analysis applied to diffusion tensor MRI. *Magnetic Resonance in Medicine* 47 (2002), 354–363.
- [6] JONES, D. K., HORSFIELD, M. A., AND SIMMONS, A. Optimal strategies for measuring diffusion in anisotropic systems by magnetic resonance imaging. *Magnetic Resonance in Medicine* 42 (1999), 515–525.
- [7] HEID, O. Eddy current-nulled diffusion weighting. In *Proceedings of the ISMRM* (2000), vol. 8, p. 799.
- [8] BRACEWELL, R. N. *The basic theorems*, third ed. McGraw-Hill International Editions, Singapore, 2000, chapter 6.
- [9] CALLAGHAN, P. T. *Introductory nuclear magnetic resonance*. Oxford University Press, Oxford, UK, 1991, chapter 2.

CHAPTER 5

LIMITATIONS AND REQUIREMENTS OF DIFFUSION TENSOR FIBRE-TRACKING: ASSESSMENT USING SIMULATIONS

CONTENTS

5.1	Introduction	95
5.2	Methods	97
5.2.1	Simulated data	97
5.2.2	Tracking algorithm	100
5.2.3	Simulations	100
5.2.4	Real Data	102
5.2.5	Measures of reliability of tracking	103
5.3	Results and discussion	104
5.3.1	Model A	104
5.3.2	Model B	109
5.3.3	Real Data	112
5.4	Conclusion	113
	References	114

5.1 INTRODUCTION

As mentioned in chapter 3, diffusion tensor MRI (DT-MRI) is a promising tool for the study of the microscopic properties of tissue [1]. In particular, it can be used to infer the directionality of the white matter fibre architecture within a voxel. With this information, studies of human white matter connectivity *in vivo* are now potentially feasible, although a number of issues remain to be addressed.

The lack of a relevant *in vivo* gold standard makes the validation of DT-MRI tractography techniques particularly difficult. It has been shown that the direction of fastest diffusion as measured by DT-MRI does correspond to the orientation of the fibres, both in excised rabbit myocardium [2] and *in vivo* (using manganese-enhanced MRI) in the optic tract of the rat [3]. However, such techniques are only applicable in limited circumstances, and cannot in general be used to assess or validate DT-MRI tractography techniques. For this reason, fibre tracking algorithms are often assessed by comparing tracking results to an appropriate anatomical atlas [4, 5, 6, 7]. However, the validation is then restricted to major white matter pathways that are well known, and biological variation cannot be taken into account. There is therefore a need for alternative methods to assess the characteristics of particular DT-MRI fibre tracking algorithms, and simulations on synthetic data in which the fibre geometry is known exactly can be used for this purpose.

In particular, the limitations and requirements of DT-MRI fibre tracking have not been fully assessed. Fibre tracking relies on the use of good quality diffusion-weighted images, together with a suitable algorithm to generate the tracks. However, the acquisition of a set of images of sufficiently high quality is time-consuming, and in practice, quality will be compromised to allow shorter scan times. For example, a number of studies have reported results using high quality data acquired from healthy cooperative volunteers in 30-40 minutes [5, 6, 8]. However, Jones *et al* [9] argue that scans longer than 15 minutes are unlikely to be well tolerated by patients, and proposed a DTI acquisition sequence that could be run within 15 minutes, designed to optimise the compromise between the various image acquisition parameters. Different tracking algorithms may be susceptible to noise, partial volume effects, motion artefacts, etc., to varying degrees. It is therefore essential to establish which structures can be tracked reliably and under what conditions this can be achieved. Simulations can again provide a means of assessing these issues.

Lori *et al* [10] have recently presented results using an algorithm proposed by Conturo *et al* [5] showing the importance of signal to noise ratio (SNR) for reliable tracking. Basser *et al* [6] have assessed the behaviour of their own algorithm in a range of specific situations (for example crossing and ‘kissing’ fibres). Nevertheless, many aspects have not been addressed, such as the dependence of these algorithms on anisotropy, partial volume effects, curvature, and step size (see section 5.2.2).

This chapter describes how the properties of a tracking algorithm can be studied in a controlled and systematic way using simulations. In particular, the effects of SNR, anisotropy, fibre curvature and cross-sectional radius, interpolation and other algorithm-specific parameters are investigated using both synthetic and real data.

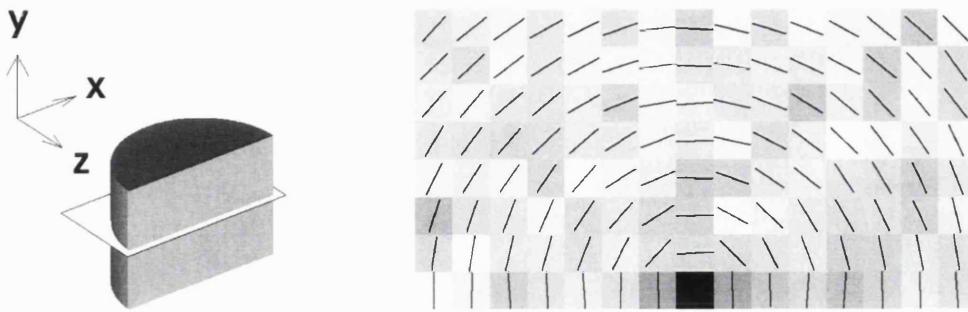


Figure 5.1: Model A – a half cylinder of constant anisotropy aligned along a semi-circular path around a central axis, as shown on the left. Right: a slice through the simulated data set, showing a fractional anisotropy map (SNR = 32) with an overlay of the major eigenvector projected onto the slice.

5.2 METHODS

For the purposes of this study, the popular streamlines algorithms is assessed, as introduced by Mori *et al* [4] (using nearest-neighbour interpolation) and by Conturo *et al* [5] (using tri-linear interpolation). This algorithm is based on following the major eigenvector of the diffusion tensor stepwise, starting from a user defined seed point. It was applied to two types of simulated diffusion tensor fields (models A & B below), as well as to real data acquired from a healthy volunteer. The simulated data allow a controlled study where the true outcome is known in advance, whereas the real data can be used to evaluate the relevance of the simulation results to real *in vivo* fibre structures. All distances in this study are expressed in units of voxels, such that the actual dimensions of the simulated fibre can be derived from the image resolution.

5.2.1 SIMULATED DATA

Three-dimensional simulated data were generated using Matlab (The Mathworks Inc., Natick, MA) for two models, A and B, with isotropic voxel sizes and assuming an orthogonal-tetrahedral diffusion encoding [5, 11] with $b\bar{\lambda} = 1$, where b is the diffusion-weighted b -factor, and $\bar{\lambda}$ is the mean eigenvalue of the diffusion tensor, corresponding to the average apparent diffusion coefficient (see section 3.1).

Model A consisted of a half cylinder of constant anisotropy aligned along a circular path around a central axis (figure 5.1). The fractional anisotropy (FA) [12] of the material could be varied. Various radii of curvature could be studied by seeding from points at different distances from the central axis.

Model B consisted of a fibre of circular cross-section and constant anisotropy aligned along a semicircular path embedded within a homogeneous background (figure 5.2). This background was either isotropic, or anisotropic with the direction of fastest diffusion ori-

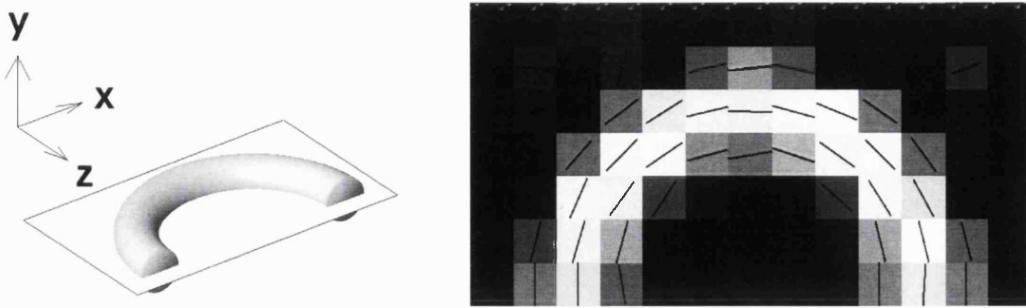


Figure 5.2: Model B – a fibre of constant anisotropy aligned along a semi-circular path, embedded within a homogeneous background, as shown on the left. Right: a slice through the data set, showing a fractional anisotropy map (SNR = 32) with an overlay of the major eigenvector projected onto the slice. In this example, the background is isotropic (excluding the effects of noise).

ented perpendicular to the plane of the fibre. The FA within the fibre, the FA in the background outside the fibre, the radius of the path of the fibre, and the fibre cross-sectional radius could be varied independently.

To ensure that the dimensions of the data sets did not place restrictions on the tracking process, the data sets generated were considerably larger than required. Semi-circular models were chosen to allow a simple simulation of the more highly curved sections of a fibre tract, where errors in fibre tracking are most likely to occur. The main difference between the two models is that model A consists of a homogeneous medium, while model B consists of a fibre embedded in a background medium. A comparison of the two models should enable the examination of the effects of partial volume.

COMPUTATION OF THE SIMULATED DATA

For both models, the origin of the coordinate system was aligned with the axis of rotation of the fibres, giving the following expressions for the eigenvectors $\{\hat{\epsilon}_i\}$ within the fibre:

$$\hat{\epsilon}_1 = \frac{1}{x^2 + y^2} \begin{pmatrix} y \\ -x \\ 0 \end{pmatrix} \quad \hat{\epsilon}_2 = \frac{1}{x^2 + y^2} \begin{pmatrix} x \\ y \\ 0 \end{pmatrix} \quad \hat{\epsilon}_3 = \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix} \quad (5.1)$$

The diffusion tensor \mathbf{D} could be evaluated from these using the following expression:

$$\mathbf{D} = \mathbf{E} \Lambda \mathbf{E}^T \quad (5.2)$$

where:

$$\mathbf{E} = \begin{pmatrix} \hat{\epsilon}_1 & \hat{\epsilon}_2 & \hat{\epsilon}_3 \end{pmatrix} \quad \text{and} \quad \Lambda = \begin{pmatrix} \lambda_1 & 0 & 0 \\ 0 & \lambda_2 & 0 \\ 0 & 0 & \lambda_3 \end{pmatrix}$$

where $\{\lambda_i\}$ correspond to the eigenvalues of the diffusion tensor. For both models, the diffusion tensor was assumed to be axially symmetric within the fibre (i.e. $\lambda_1 > \lambda_2 = \lambda_3$). The FA within the fibre was set by adjusting the eigenvalues $\{\lambda_i\}$ according to the following equations:

$$\begin{aligned} \lambda_1 &= \bar{\lambda}(1 + 2\alpha) \\ \lambda_2 = \lambda_3 &= \bar{\lambda}(1 - \alpha) \\ \text{where } \alpha &= \frac{\text{FA}}{\sqrt{3 - 2(\text{FA})^2}} \end{aligned}$$

These expressions were chosen to ensure that $\bar{\lambda}$ remains constant, and were derived from the definition of the FA (equation 3.37).

For model B, the diffusion tensor in the background outside the fibre was set using expressions similar to that above:

$$\mathbf{D} = \bar{\lambda} \begin{pmatrix} 1 - \beta & 0 & 0 \\ 0 & 1 - \beta & 0 \\ 0 & 0 & 1 + 2\beta \end{pmatrix} \quad (5.3)$$

where β controls the amount of anisotropy, and is given by the same equation as α above. In this case, the largest eigenvalue corresponds to the z -axis. The value of $\bar{\lambda}$ is uniform in the *in vivo* human brain [13], and was therefore kept constant. Note however that $\bar{\lambda}$ is significantly higher in cerebro-spinal fluid (CSF), and that the effects of partial volume with CSF have therefore not been assessed.

A point was deemed inside the fibre if it satisfied the following condition:

$$\left(\sqrt{x^2 + y^2} - R \right)^2 + z^2 \leq r^2 \quad (5.4)$$

where R is the radius of curvature of the fibre, and r is its cross-sectional radius.

The values of the tensor elements derived from these expressions were numerically integrated over each voxel to produce the data sets. Using this approach, partial volume effects were automatically incorporated into the models. An approximation to Rician noise [14, 15] was added in the base diffusion-weighted images, and the SNR was defined as S_0/σ , where S_0 is the $b = 0$ image intensity and σ is the standard deviation of the noise.

5.2.2 TRACKING ALGORITHM

The tracking algorithm used is a simple streamlines technique, based on that proposed by Mori *et al* [4] and Conturo *et al* [5]. It differs from both these techniques in the termination criteria used: tracks are terminated if they venture outside the fibre, or once they have reached the end plane (see later). In a real tracking experiment, the geometry of the fibres is not known, and different termination criteria would be used (for example, thresholds on curvature or anisotropy). The algorithm operates as follows: starting from a user-defined seed point \underline{r}_0 , the track is generated by calculating the next point along the path from the current point iteratively, in the following manner:

1. The base image intensities $\{S_n(\underline{r}_i)\}$ are calculated at the current point \underline{r}_i , depending on the choice of interpolation method.
2. The corresponding diffusion tensor $\mathbf{D}(\underline{r}_i)$ is calculated from the base image intensities assuming the relationship $S_n = S_0 e^{-b_n \cdot \mathbf{D}}$ [16].
3. The normalised major eigenvector $\hat{\epsilon}_1(\underline{r}_i)$ of the diffusion tensor is computed.
4. The next point \underline{r}_{i+1} along the path is calculated by stepping a distance s (the step size) along the direction of the major eigenvector: $\underline{r}_{i+1} = \underline{r}_i + s \hat{\epsilon}_1(\underline{r}_i)$.

Two types of interpolation methods were used in this study: nearest-neighbour and tri-linear. When using tri-linear interpolation, the base image intensities are calculated at each step along the track by linearly interpolating to the point of interest from the 8 nearest discrete data points in 3D space (the data points forming the corners of the one pixel-wide cube containing the point of interest). When using nearest neighbour interpolation, steps 1 to 4 can all be computed prior to performing the actual tracking, making the whole process much less computationally intensive. Tri-linear interpolation should allow a smoother representation of the tensor field, thus enabling more accurate tracking in regions of high curvature.

5.2.3 SIMULATIONS

The tracking algorithm was tested by computing 100 tracks for each combination of model parameters. For each of those 100 tracks, a new data set was generated according to the current model parameters, but with different random noise. A large range of conditions was studied:

- SNR = 8, 16, 32, 64, 128 and 100,000
- FA = 0.2, 0.4, 0.6 and 0.8

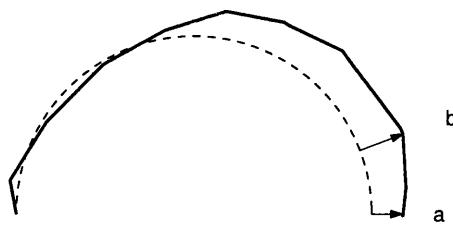


Figure 5.3: the dotted line indicates an ideal track; the solid line indicates an example trajectory obtained from simulations. The outcome measures shown are: (a) end-point x -offset and (b) maximum deviation. The end-point y -offset is measured perpendicular to the page. Not shown here are the RMS deviation and the success rate. See text for details.

- Background FA = 0 and 0.2 (Model B)
- Step size = 0.02, 0.1, 0.2, 0.4, 0.8 and 1.6 pixels
- Radius of curvature = 0.5, 1, 2 and 4 pixels
- Fibre cross-sectional radius = 0.25, 0.5, 1 and 2 pixels (Model B)
- nearest-neighbour or tri-linear interpolation

To assess the quality of a particular track, several output parameters were computed:

- The end-point x - and y -offsets (see figure 5.3): the position of the computed end-point relative to the true end-point in pixel units.
- The maximum departure in pixels of the computed track from the ideal track (see figure 5.3): this gives an estimate of the size of the structures that can be tracked. This parameter is always positive.
- The root mean square (RMS) departure in pixels of the computed track from the ideal track: this provides a more global measure of the quality of the track.
- The success rate: for model A, the track was deemed to have failed if it ventured outside the data set; for model B, if the track ventured outside the fibre by more than one pixel, the simulation was terminated and the track was classed as a failure. The threshold of one pixel was chosen since the signal from a voxel that is partial volume averaged with the fibre will contain a contribution from the fibre. A track may thus be as much as one voxel away from the fibre, yet still be influenced by signal originating from it.

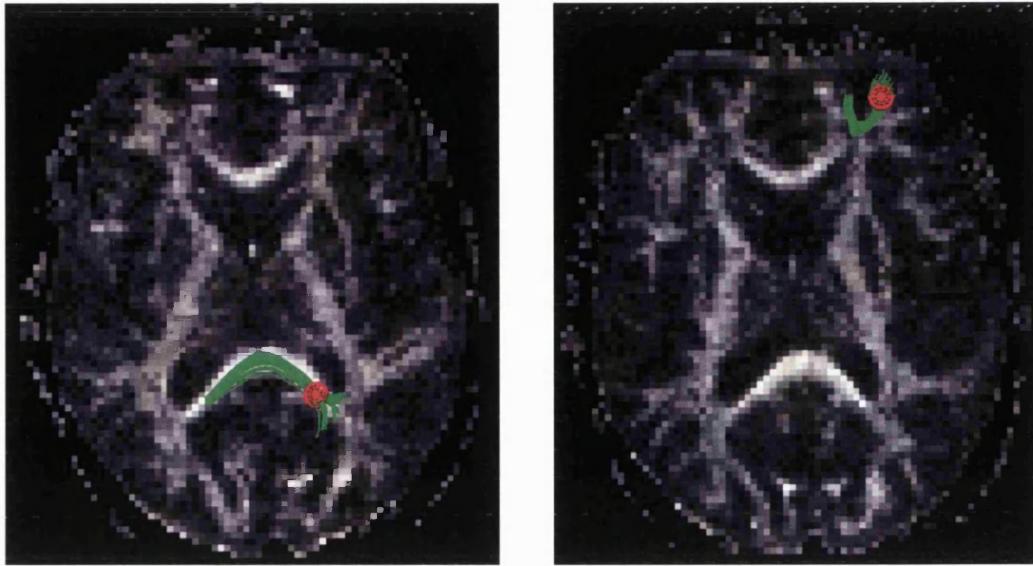


Figure 5.4: FA maps for two axial slices in a healthy volunteer, showing the fibre bundles used to test the tracking algorithm. Tracks are shown in green, and the end ROI (see section 5.2.4) is shown in red. Fibres in the corpus callosum on the left, in sub-cortical white matter on the right.

5.2.4 REAL DATA

Measurements were acquired from a healthy volunteer on a Siemens Vision 1.5T scanner using a SE-EPI sequence with a pair of diffusion gradients on either side of the refocusing RF pulse, described in section 4.1 (TE = 110ms; 128×128; 256 mm FOV; 2 mm slice thickness; 2×2×2 mm voxels). An orthogonal-tetrahedral diffusion encoding [5, 11] was used with the following parameters: $\delta/\Delta = 27.1/52.3$ ms, b -values = 0 and 1066 s/mm². Reference scans were acquired for the online correction of B_0 eddy-current effects [17]. Twenty-five contiguous axial slices were acquired covering a 5 cm slab centred on the corpus callosum. Six measurements were acquired, and the SNR in the base images was varied by adding different numbers of images, with SNR ranging from 13 for a single acquisition to 32 for all 6 averages. When using less than the full 6 measurements, an approach similar to the bootstrap method can be used to generate several different data sets with the same SNR [18]. Thus, 6 different data sets were generated for each SNR level (except at the highest level, since there is only one combination for all 6 measurements).

Two fibre bundles were selected with properties similar to those studied in the simulations: one in the corpus callosum (mean FA = 0.8, high cross-sectional radius and low curvature), the other in sub-cortical white matter (mean FA = 0.3, low cross-sectional radius and high curvature). These are shown in figure 5.4. Tracking was initiated from a seed region of interest (ROI), and a track was considered successful if it entered the end ROI. The seed ROI consisted of a cubic region equivalent in size to one voxel, within

which seed points were arranged in a three-dimensional $5 \times 5 \times 5$ grid. Note that different seed points within the same voxel will in general result in different tracks [10]. The end ROI consisted of a spherical volume of radius two pixels. These ROIs were selected empirically by performing the tracking given the optimal conditions available (SNR = 32, step size = 0.02 pixels, using tri-linear interpolation). Those points within the seed ROI from which the tracking failed given the optimal conditions were not used in the subsequent experiments. This led to a total of 125 seed points being used for the seed ROI in the corpus callosum, and 121 for that in the sub-cortical white matter. Such an approach is necessary due to the unknown characteristics of the actual fibre.

The experiments were performed with nearest-neighbour and tri-linear interpolation for a range of step sizes (0.02 to 1.6 pixels) and SNR, the latter obtained from using different numbers of averages. The outcome measures mentioned above for the simulated data cannot be calculated for the *in vivo* data, since the true fibre tract is not known. Therefore, the measure of success used was the proportion of the seed points that entered the end ROI.

5.2.5 MEASURES OF RELIABILITY OF TRACKING

For model A, the reliability measure RM_A = average maximum departure + 2SD was defined. Approximately 98% (assuming a one-sided Gaussian distribution, since the departure is always positive) of all tracks computed have a maximum departure smaller than this value. Therefore if it is required that the computed track stay within a distance d of the ideal track, the contour $RM_A = d$ indicates the range of parameters that allows reliable tracking, i.e. tracking with a 98% success rate. A lower RM_A indicates improved tracking.

For model B, the reliability measure was defined as RM_B = fibre radius for which the success rate is 98%. Once again, the contour $RM_B = d$ would indicate the combination of parameters with which a fibre of radius d can be tracked reliably. This measure was chosen as it is comparable to the measure used for model A.

As previously mentioned, neither of the above measures can be applied to real data since the underlying fibre structure is not known. Therefore, the success rate for the real data was defined as the proportion of seeded points with successful tracks (i.e. entering the end ROI). A limitation of this measure is that it does not allow a direct comparison with the contour plots from the simulated data.



Figure 5.5: example tracks ($n = 20$) computed using model A to illustrate the effects of noise and interpolation. Top row: using tri-linear interpolation; Bottom row: using nearest-neighbour interpolation. Left: SNR = 8; middle: SNR = 32; right: SNR = 128. The ideal track is shown by a dotted line (visible in the left hand images). All plots were produced with a fractional anisotropy of 0.4 and a radius of curvature of 4 pixels, using a step size of 0.02 pixels.

5.3 RESULTS AND DISCUSSION

5.3.1 MODEL A

To illustrate the variability in the tracks obtained using model A as a function of SNR, each sub-figure in figure 5.5 shows 20 tracks generated from the same seed point at different SNR levels, using either nearest-neighbour or tri-linear interpolation. As can be seen in the figure, the variability in the tracks is reduced by increasing the SNR. It can also be seen that the precision of the tracks is further improved by using tri-linear interpolation.

Some interesting simple relationships were observed with the simulations using model A. As might be expected, increasing the step size increases the average end-point x -offset. This relationship was found to be approximately linear and independent of SNR, FA, or the type of interpolation (figure 5.6). This means that, in principle, it is possible to improve the quality of the tracking simply by reducing the step size. The average end-point y -offset was not significantly different from zero for all conditions. This was expected, since there is no preferred direction for deviations along this axis. The standard deviation (SD) of the end-point x - and y -offsets, however, had a strong dependence on the SNR, FA and type of interpolation. In particular, using tri-linear (rather than nearest-neighbour) interpolation produced little change in the average end-point offset, but did reduce its standard deviation (figure 5.6). In other words, tri-linear interpolation makes the tracking more reproducible. As expected, increasing the SNR and the FA decreases the SD, also making the tracking more reproducible.

The relationships observed with the maximum departure and RMS departure were

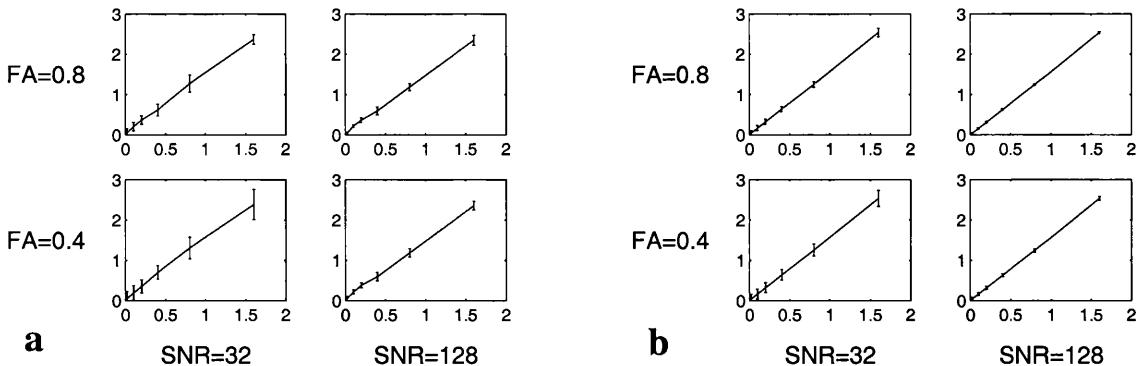


Figure 5.6: plots of the end-point x -offset (y -axis) as a function of step size (x -axis), both in pixel units, (a) using nearest-neighbour interpolation, (b) using tri-linear interpolation. For each set, the top row corresponds to $FA = 0.8$, the bottom row to $FA = 0.4$; the left column to $SNR = 32$, the right column to $SNR = 128$. The error bars represent two SDs either side of the mean value. All plots correspond to the case of a radius of curvature of two pixels.

similar to that seen with the end-point offset. However, when using nearest-neighbour interpolation, as the step size decreases, the average maximum departure does not tend to zero (see figure 5.7(a)). This is due to a limitation of nearest-neighbour interpolation: the track must remain straight within each voxel, and therefore cannot remain close to a curved ideal track. Using tri-linear interpolation, the track can follow the ideal path much more closely. Therefore, tracking using nearest-neighbour interpolation introduces systematic departures from the ideal track, especially when the curvature of the structure is high. This was also true for the RMS departure (figure 5.8). It should be noted that although sub-voxel accuracy can be obtained in this model with an appropriate step size (see figures 5.6 and 5.7), small departures from the true path can have a significant impact when tracking on real data, in particular if the track ventures into an adjacent structure with an orientation unrelated to the original structure. Such problems are not seen with this model since it consists of a single structure.

The various outcome measures demonstrated little or no dependence on curvature (except at very low step sizes; see above). However, it is difficult to interpret this result, since a simulated path with a lower radius of curvature is intrinsically shorter. While a lower radius of curvature would be expected to make tracking more difficult, a shorter path should lead to a lower percentage of tracking errors. It is possible, therefore, that the two effects to some extent cancel each other out. Although this may be a limitation of this particular model, in a real situation, a curved fibre-path with a lower radius of curvature would in general be shorter, and the semi-circular model may represent a reasonably realistic approximation to in-vivo fibre structures.

In summary, it can be seen from figures 5.6 and 5.7 that using interpolation and a low

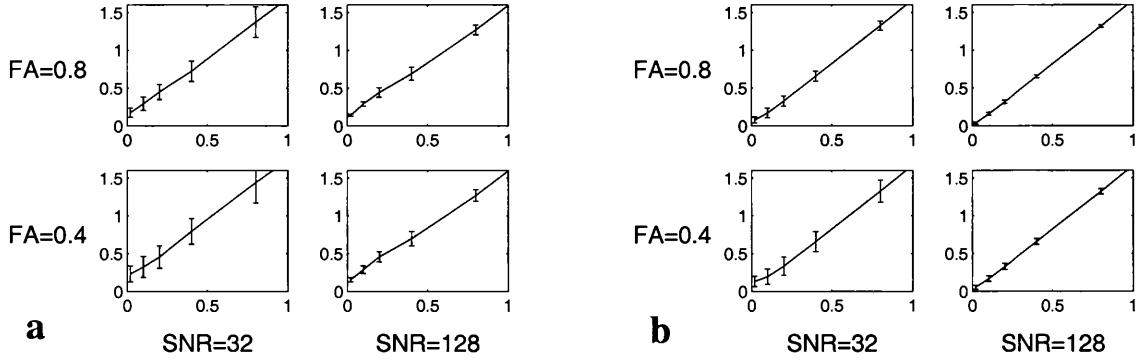


Figure 5.7: Plots of the maximum departure (y -axis) as a function of step size (x -axis), both in pixel units, (a) using nearest-neighbour interpolation, (b) using tri-linear interpolation. For each set, the top row corresponds to $FA = 0.8$, the bottom row to $FA = 0.4$, the left column to $SNR = 32$, and the right column to $SNR = 128$. The error bars represent two standard deviations either side of the mean value. All plots correspond to the case of a radius of curvature of two pixels. Note that the scale is different from that in figure 5.6.

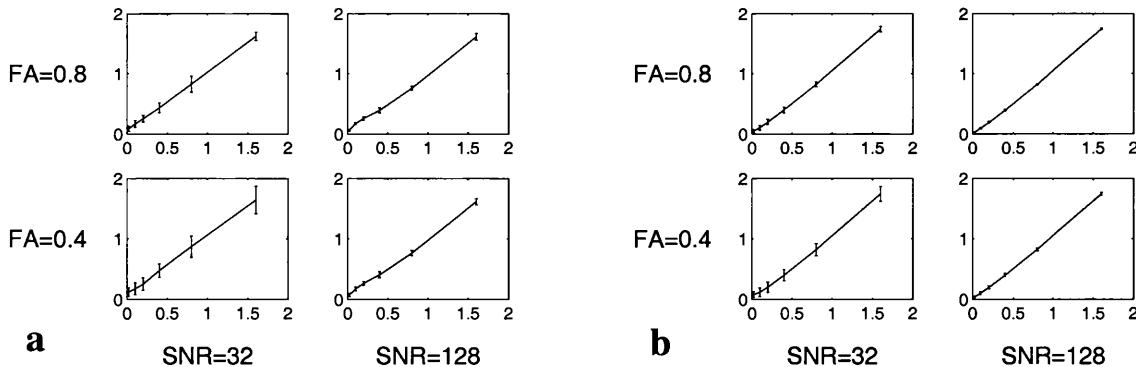


Figure 5.8: Plots of the root mean square departure (y -axis) as a function of step size (x -axis), both in pixel units, (a) using nearest-neighbour interpolation, (b) using tri-linear interpolation. For each set, the top row corresponds to $FA = 0.8$, the bottom row to $FA = 0.4$, the left column to $SNR = 32$, and the right column to $SNR = 128$. The error bars represent two standard deviations either side of the mean value. All plots correspond to the case of a radius of curvature of two pixels.

step size can produce very accurate tracking (small end-point offset, maximum departure and RMS departure). However, the precision (as represented by the error bars), and thus the reliability of the tracking, will also depend on the SNR, FA and radius of curvature.

Figure 5.9 shows contour plots for model A showing the dependence of RM_A (see section 5.2.5) on step size and SNR. Such plots are used extensively in this chapter, and will therefore be explained in detail here. The parameters SNR and step size were chosen because they are to some extent under user control. Analysis of these plots can be used to determine the image quality required to track a particular white matter structure with known characteristics, or conversely whether tracking this structure is feasible for a given image quality. They are used by considering the positions of the individual contours: tracking is more easily achieved if the position of a contour with a given RM_A lies closer to the top left corner of the plot (i.e. when a low SNR and a high step size are sufficient to track reliably). The effect of particular parameters (e.g. FA, interpolation, etc.) on tracking reliability can be studied by examining their impact on the position of the contours.

As expected from the results of figures 5.6 and 5.7, the reliability of tracking improves as the SNR increases and as the step size decreases: the contours corresponding to increased reliability in figure 5.9 are closer to the bottom right of the plot. The effect of the interpolation method can be studied by comparing plots generated using one method to those generated using the other method: contours obtained using tri-linear interpolation are closer to the top left corner than the equivalent contours produced using nearest-neighbour interpolation. This indicates that using tri-linear interpolation improves the quality of the tracking. Conversely, as the FA decreases, the contours move towards the bottom right corner, indicating reduced tracking reliability.

As an example of the use of these contour plots, consider tracking with the aim of not deviating from the ideal track by more than 0.5 pixel for a structure with $FA = 0.8$ and a data set with an $SNR = 32$ (figures 5.9(a&b)). According to the simulation results, this can be achieved reliably (i.e. with a 98% success rate) using a step size of 0.2 pixels and tri-linear interpolation (figure 5.9(b)). Alternatively, it can be done using nearest-neighbour interpolation by decreasing the step size to 0.1 pixels (figure 5.9(a)).

By comparing figure 5.9(a) (nearest-neighbour interpolation, $FA = 0.8$) and figure 5.9(d) (tri-linear interpolation, $FA = 0.4$), it can be seen that using tri-linear interpolation has a more important effect on the reliability of tracking than a doubling of the anisotropy (corresponding contours in figure 5.9(d) are closer to the top left than in figure 5.9(a)). It is interesting to note that although the anisotropy is a property of the particular tract of interest, over which there is no control, the choice of interpolation method is under user control.

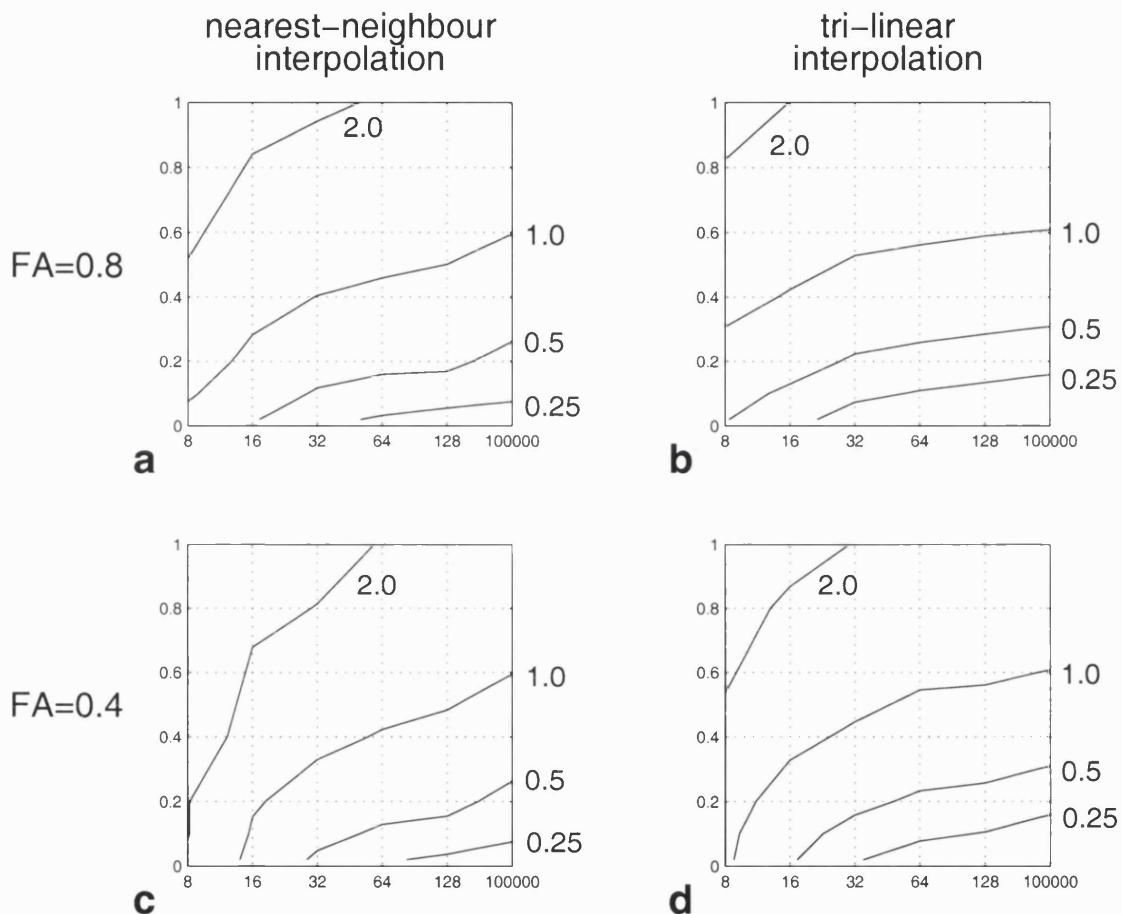


Figure 5.9: Contours of the measure of reliability, RM_A (see text), as a function of SNR (x -axis) and step size (y -axis), in units of pixels. (a) & (c): using nearest-neighbour interpolation, (b) & (d): using tri-linear interpolation. (a) & (b): $FA = 0.8$, (c) & (d): $FA = 0.4$. Each contour is labelled with its corresponding RM_A value (in pixels). Note that the x -axis uses a non-linear scale. All plots correspond to the case of a radius of curvature of two pixels. It can be seen that a higher SNR, a higher FA, a lower step size and the use of tri-linear interpolation all improve the quality of the tracking.

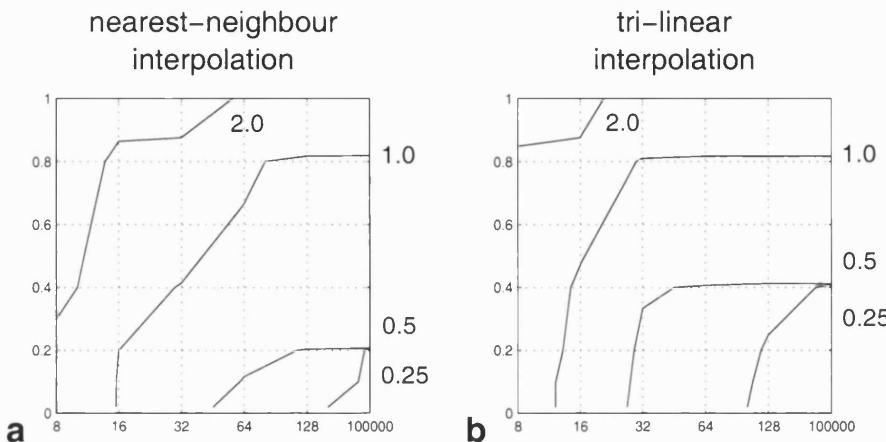


Figure 5.10: Contours of the measure of reliability, RM_B (see text), as a function of SNR (x-axis) and step size (y-axis), in units of pixels. (a): nearest-neighbour interpolation, (b): tri-linear interpolation. Both plots correspond to the case of a radius of curvature of two pixels, FA within the fibre = 0.8, and an isotropic background. It can be seen that a higher SNR, a lower step size and the use of interpolation all improve the quality of the tracking.

The results from figure 5.9 indicate that in this model, it is possible to track without deviating further than a quarter of a pixel. Although this suggests that it may be possible to track fibres with cross-sectional radii as small as a quarter of a pixel reliably, it should be noted that model A is free from partial volume effects with other structures; with different models or in real data, tracking such small structures may not be possible (see model B).

In summary, figure 5.9 indicates that the tracking reliability increases with higher SNR and lower step size for small fibre structures. For lower values of the FA, an even higher SNR is needed to track as successfully as at higher FA. The use of tri-linear interpolation allows successful tracking with a larger step size and/or a lower SNR.

5.3.2 MODEL B

Due to the nature of this model, the average end-point offset and the maximum and RMS departure values will contain a significant amount of bias. For example, it is not possible for the maximum departure to be greater than 2 voxels when tracking a fibre with a cross-sectional radius of 1 voxel, since the track would have failed. The results for these outcome measures are therefore not presented.

The contour plots for model B (figure 5.10) show the dependence of RM_B on step size and SNR. The general trend is the same as that for model A, in that the reliability of the tracking increases as the SNR increases and as the step size decreases. However, the effects of partial volume are immediately apparent: a much higher SNR is required to track small fibres using this model. The importance of interpolation can be readily

appreciated from this figure. For example, the results suggest that it is not possible to track a fibre with a cross-sectional radius of 0.5 pixels at an SNR of 32 reliably without using tri-linear interpolation (figure 5.10), irrespective of the step size used (cf. figures 5.9(a&b)). When interpolation is used, this can be achieved with a step size of 0.3 pixels or less, indicating that lowering the step size cannot always replace the use of interpolation.

Note that for model B, it is possible to track reliably using a higher step size than the results from model A suggest (figures 5.9(b) & 5.10(b)), which seems to indicate that partial volume effects are in some ways beneficial. This can only be the case when the background is isotropic (corresponding for example to grey matter): in these circumstances, the signal from the background contains no orientational preference, and any directional information is due to the fibre. It may thus be possible to continue tracking the fibre even though the track is technically outside the fibre. However, these effects would be expected to have increasingly detrimental effects as the anisotropy in the background increases, corresponding to increasing partial volume averaging with other white matter structures. Figure 5.11 shows the effect of the background FA on the tracking for model B. This figure shows contour plots using two different values of FA within the fibre and two different values of the background FA outside the fibre, all using tri-linear interpolation. When the FA within the fibre is reduced, a higher SNR is needed to track with the same reliability. When the background FA is increased, tracking small fibres reliably requires both a higher SNR and a lower step size, due to more important partial volume effects with the background. Note that this effect can be seen even with relatively small levels of anisotropy in the background. For example, however small the step size, the minimum SNR required to track a fibre with $FA = 0.8$ and $radius = 0.5$ pixels in an isotropic medium is approximately 30 (figure 5.11(b)) (note that this is not possible for any step size if the SNR is lower than this value). If the FA within the fibre is reduced to 0.4, the SNR needs to be increased to about 64 to track the same fibre reliably (figure 5.11(a)), corresponding to a four-fold increase in the scan time required. If on the other hand the background FA is increased to 0.2, the SNR needs to be increased to about 45 (figure 5.11(d)). When the two effects are combined (i.e. fibre $FA = 0.4$, background $FA = 0.2$), it becomes extremely difficult to track such a fibre reliably (figure 5.11(c)). Note that it is possible to track a fibre with cross-sectional radius = 0.25 (figures 5.11(a,c&d)) only when the conditions for tracking are optimal (cf. model A, figures 5.9(b&d)). Therefore, although the trends are similar for both models A & B, it should be noted from model B that partial volume effects can be important, in particular for small structures in an anisotropic background, even for relatively small levels of background anisotropy.

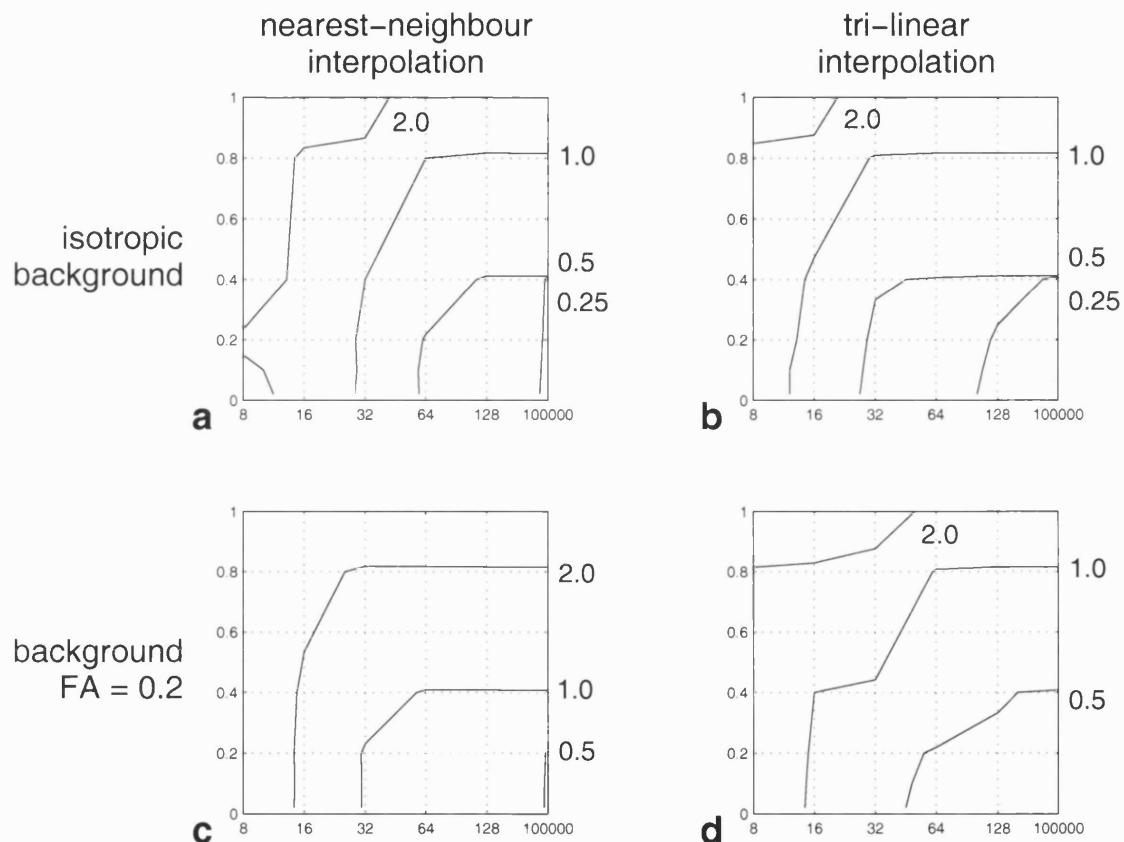


Figure 5.11: Contours of the measure of reliability, RM_B (see text), as a function of SNR (x -axis) and step size in units of pixels (y -axis). (a) & (b): isotropic background ($FA = 0$), (c) & (d): anisotropic background ($FA = 0.2$) aligned perpendicular to the plane of the fibre. (a) & (c): fibre $FA = 0.4$, (b) & (d): fibre $FA = 0.8$. All plots were produced using tri-linear interpolation and a radius of curvature of two pixels.

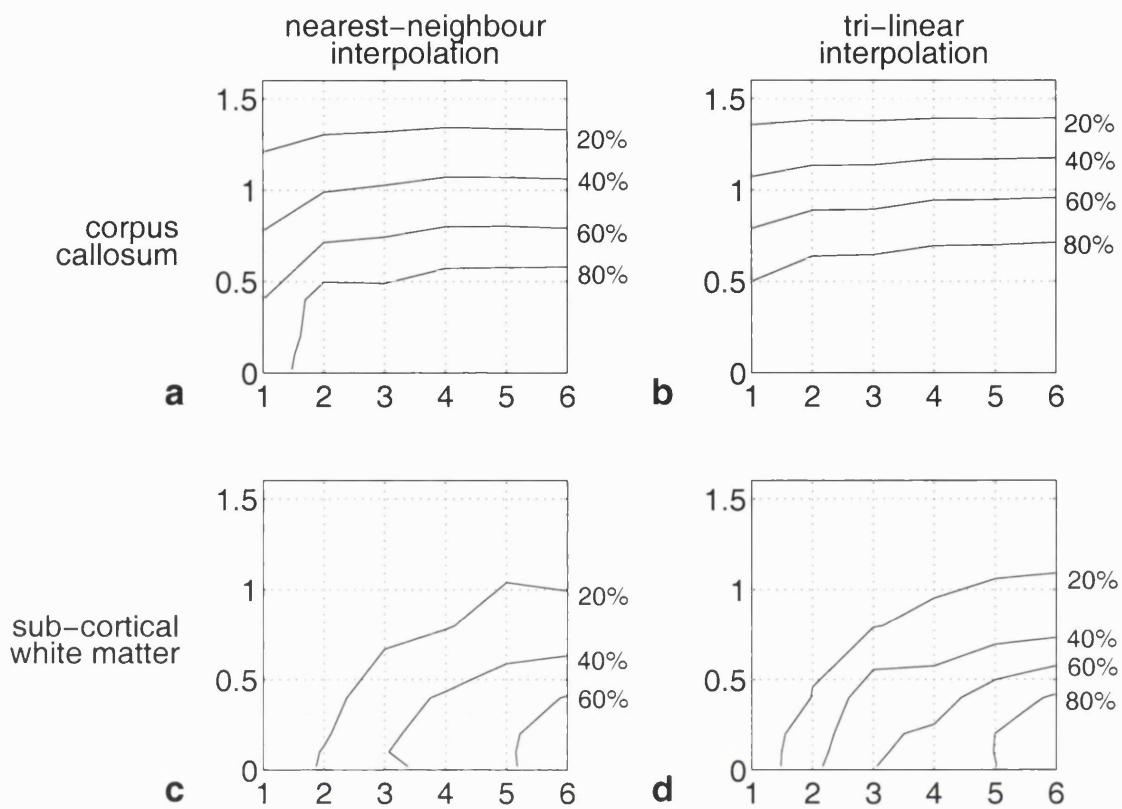


Figure 5.12: Contours of the success rate for tracking on real data as a function of the number of acquisitions (x -axis), and step size in pixel units (y -axis). (a) & (b): corpus callosum, (c) & (d): sub-cortical white matter. (a) & (c): nearest-neighbour interpolation, (b) & (d): tri-linear interpolation. The SNR ranges from 13 for a single acquisition to 32 for all 6 acquisitions. Each contour is labelled with its corresponding success rate.

5.3.3 REAL DATA

Results obtained from the real data shown in figure 5.4 are plotted on figures 5.12(a&b) for tracking in the corpus callosum and figures 5.12(c&d) for tracking in the sub-cortical white matter. These show contours of success rate as a function of the number of averages (x -axis) and the step size of the tracking algorithm (y -axis). Varying the number of acquisitions allows the study of SNR ranging from 13 to 32 (from one to six averages). It should be noted that since a maximum SNR of only 32 was used, the actual shapes of the contours are not well defined, particularly in the areas of high step size or low SNR where the reproducibility of tracking is decreased (equivalent to an increase in the SD in figures 5.6 & 5.7).

Tracking reliably in the corpus callosum was possible using a wide range of SNR and step size (figures 5.12(a&b)). In particular, 80% success could be achieved even with one average (SNR = 13) using tri-linear interpolation, as was expected from the simulations

using model A. Interpolation did not have a very significant effect when tracking in the corpus callosum. On the other hand, in the sub-cortical white matter the interpolation had a substantial effect (the 80% success rate contour is not visible in figure 5.12(c)). As suggested by the simulations using model B, relatively reliable tracking in the sub-cortical white matter could only be achieved with a very restricted combination of parameters: using interpolation with the highest SNR available and a low step size (figure 5.12(d)).

The results reflect the trends seen in the simulations. In particular, the success rate is increased at higher SNR, lower step size and by using tri-linear interpolation. By comparing the results for the two fibres, it can be seen that the success rate is also increased at higher FA, higher fibre cross-sectional radius and higher radius of curvature. In addition, assuming that the fibre in the corpus callosum is less contaminated by partial volume effects, it can be seen that the success rate is much improved when such effects are not significant.

5.4 CONCLUSION

In the absence of a gold standard for DT-MRI tractography, simulations where the exact fibre geometry is known provide a framework for investigating different aspects of fibre tracking in a controlled way. The results presented here indicate that the use of tri-linear interpolation produces more reliable results than nearest-neighbour interpolation, and the former should therefore be used whenever possible. In addition, the information provided by the simulations may help to determine whether a particular acquisition sequence will produce images of sufficient quality to track structures with particular characteristics (FA, radius of curvature, etc). Simulations may also be helpful in deciding on the structures that should be considered for tracking for a given quality of image.

The simulations presented in this study provide an indication of the capabilities of the algorithm selected, in terms of its dependence on SNR, partial volume effects, etc. However, they can also be extended to other algorithms, and could be used in their fine-tuning and assessment. For the particular algorithm implemented in the present work, potential developments include evaluation of other interpolation methods (spline or sinc), tensor field regularisation techniques [8, 19] and a context-sensitive step size (as suggested by Basser *et al* [6]) based on results from contour plots such as figure 5.11. More complex models may be used to investigate the behaviour of these algorithms when faced with more challenging situations (for example crossing or ‘kissing’ fibres, or branching) [6].

Due to uncertainties concerning the true *in vivo* fibre structure and other practical limitations (e.g. limited SNR), it is difficult to use simulations to assess directly the reliability of tracking performed on real data. Moreover, the models used are simplistic with respect

to the conditions found *in vivo*, in that they address very specific issues. However, trends seen in the simulations were also seen in real data, which suggests that simulations can be related to *in vivo* fibre tracking and do produce meaningful results. Therefore, simulations such as those presented in this work can make an important contribution to the assessment of diffusion tensor tractography.

REFERENCES

- [1] LE BIHAN, D., MANGIN, J.-F., POUPEON, C., CLARCK, C. A., PAPPATA, S., MOLKO, N., AND CHABRIAT, H. Diffusion tensor imaging: concepts and application. *Journal of Magnetic Resonance Imaging* 13 (2001), 534–546.
- [2] HOLMES, A. A., SCOLLAN, D. F., AND WINSLOW, R. L. Direct histological validation of diffusion tensor MRI in formaldehyde-fixed myocardium. *Magnetic Resonance in Medicine* 44 (2000), 157–161.
- [3] LIN, C. P., Y., T. W., CHENG, H. C., AND CHEN, J. H. Validation of diffusion tensor magnetic resonance axonal fiber imaging with registered manganese-enhanced optic tracts. *NeuroImage* 14 (2001), 1035–1047.
- [4] MORI, S., CRAIN, B. J., CHACKO, V. P., AND VAN ZIJL, P. C. M. Three-dimensional tracking of axonal projections in the brain by magnetic resonance imaging. *Annals of Neurology* 45 (1999), 265–269.
- [5] CONTURO, T. E., LORI, N. F., CULL, T. S., AKBUDAK, E., SNYDER, A. Z., SHIMONY, J. S., MCKINSTRY, R. C., BURTON, H., AND RAICHLE, M. E. Tracking neuronal fiber pathways in the living human brain. *Proceedings of the National Academy of Sciences* 96 (1999), 10422–10427.
- [6] BASSER, P. J., PAJEVIC, S., PIERPAOLI, C., DUDA, J., AND ALDROUBI, A. In vivo fiber tractography using DT-MRI data. *Magnetic Resonance in Medicine* 44 (2000), 625–632.
- [7] PARKER, G. J. M., WHEELER-KINGSHOTT, C. A. M., AND BARKER, G. J. Estimating distributed anatomical connectivity using fast marching methods and diffusion tensor imaging. *IEEE Transactions on Medical Imaging* 21 (2002), 505–512.
- [8] POUPEON, C., CLARK, C. A., FROUIN, V., RÉGIS, J., BLOCH, I., LE BIHAN, D., AND MANGIN, J.-F. Regularization of diffusion-based direction maps for the tracking of brain white matter fascicles. *NeuroImage* 12 (2000), 184–195.
- [9] JONES, D. K., WILLIAMS, S. C. R., GASSTON, D., HORSFIELD, M. A., SIMMONS, A., AND HOWARD, R. Isotropic resolution diffusion tensor imaging with whole brain acquisition in a clinically acceptable time. *Human Brain Mapping* 15 (2002), 216–230.
- [10] LORI, N. F., AKBUDAK, E., SHIMONY, J. S., CULL, T. S., SNYDER, A. Z., GUILLOREY, R. K., AND CONTURO, T. E. Diffusion tensor fiber tracking of human brain connectivity: acquisition methods, reliability analysis and biological results. *NMR in Biomedicine* 15 (2002), 493–515.
- [11] HASAN, K. M., PARKER, D. L., AND ALEXANDER, A. L. Comparison of gradient encoding schemes for diffusion-tensor MRI. *Journal of Magnetic Resonance Imaging* 13 (2001), 769–780.
- [12] BASSER, P. J., AND PIERPAOLI, C. Microstructural and physiological features of tissues elucidated by quantitative-diffusion-tensor MRI. *Journal of Magnetic Resonance B* 111 (1996), 209–219.
- [13] DE CRESPIGNY, A. J., MARKS, M. P., ENZMANN, D. R., AND MOSELEY, M. E. Navigated diffusion imaging of normal and ischemic human brain. *Magnetic Resonance in Medicine* 33 (1995), 720–728.
- [14] GUDBJARTSSON, H., AND PATZ, S. The Rician distribution of noisy MRI data. *Magnetic Resonance in Medicine* 34 (1995), 910–914.

- [15] HENKELMAN, R. M. Measurement of signal intensities in the presence of noise in MR images. *Medical Physics* 12 (1985), 232–233.
- [16] BASSER, P. J. Inferring microstructural features and the physiological state of tissues from diffusion-weighted images. *NMR in Biomedicine* 8 (1995), 333–344.
- [17] CALAMANTE, F., PORTER, D. A., GADIAN, D. G., AND CONNELLY, A. Correction for eddy current induced B_0 shifts in diffusion-weighted echo-planar imaging. *Magnetic Resonance in Medicine* 41 (1999), 95–102.
- [18] JONES, D. K. Determining and visualizing uncertainty in estimates of fiber orientation from diffusion tensor MRI. *Magnetic Resonance in Medicine* 49 (2003), 7–12.
- [19] TENCH, C. R., MORGAN, P. S., BLUMHARDT, L. D., AND CONSTANTINESCU, C. Improved white matter fiber tracking using stochastic labelling. *Magnetic Resonance in Medicine* 48 (2002), 677–683.

CHAPTER 6

FRONT EVOLUTION TRACTOGRAPHY

CONTENTS

6.1	Introduction	117
6.2	The FRET algorithm	118
6.2.1	Overview	118
6.2.2	The fibre orientation density function	119
6.2.3	Sampling orientation space	121
6.2.4	Front evolution	122
6.2.5	Surface reconstruction	124
6.2.6	Down-sampling the front	126
6.2.7	Propagation of the index of connectivity	127
6.2.8	Seed point selection	128
6.3	Methods	128
6.3.1	Data acquisition	128
6.3.2	Assessment of the FRET algorithm	128
6.3.3	Streamlines tracking algorithm	129
6.4	Results and discussion	129
6.4.1	Dependence on the signal to noise ratio	130
6.4.2	Dependence on initial seed region	135
6.4.3	Dependence on algorithm parameters	135
6.4.4	Dependence on partial volume effects	137
6.4.5	Comparison with streamlines	138
6.4.6	The index of connectivity	138
6.5	Conclusion	140
	References	140

6.1 INTRODUCTION

As mentioned in section 3.3, diffusion tensor imaging (DTI) can provide important information related to tissue micro-architecture in the brain *in-vivo*. The anisotropy of the diffusion process in white matter tracts is of particular interest for tractography, since it may be used to infer the direction of the underlying fibres, and thus follow the supposed path of white matter tracts in the brain *in-vivo*. Most fibre-tracking algorithms to date equate the direction of fastest diffusion, as given by the orientation of the major eigenvector of the diffusion tensor, directly with the direction of the fibres. However, as described in sections 3.2.3 and 3.3.4, DTI is particularly sensitive to the effects of both partial volume averaging and noise, leading to errors in the major eigenvector direction and artificially increased anisotropy values. The impact on the tracks generated can be considerable: if an incorrect major eigenvector direction causes the track to venture into an adjacent, yet unrelated structure, an apparent connection may be established where none exists in reality.

Ideally, fibre tracking would be performed on very high quality data (high signal-to-noise ratio (SNR), high spatial resolution, a large number of diffusion encoding directions, etc.). However, in practice, image quality will frequently be compromised in favour of feasible imaging times. In particular, the voxel size is often increased in order to boost the SNR, leading to considerable partial volume effects. When more than one fibre population is present within a single voxel, the diffusion tensor will generally fail to describe any of these adequately (see section 3.3.4).

Another problem is that a single large fibre bundle may split up into a number of smaller tracts, each projecting to a different area of the brain, such that there are a number of possible solutions emanating from a single seed point. Many of the existing techniques [1, 2, 3, 4] generate a single track per seed point, potentially leading to a considerable loss of information.

It is clear that there is a degree of uncertainty in the measurement of white matter tract orientation using DT-MRI [5], and that the level of uncertainty depends on a number of factors. This uncertainty in the true path of the tract will tend to increase as a function of the distance tracked. A number of pseudo-probabilistic techniques have recently been proposed to address this issue [6, 7, 8, 9, 10]. These assume a model to estimate the uncertainty in the orientation of the fibres, and often work by investigating a large number of possible paths emanating from a single seed point. Each connection thus generated is assigned an index of connectivity with the seed point, derived from the particular model assumed.

In such a pseudo-probabilistic framework, the tractography problem may be expressed

in terms of a front emanating from the seed point, the evolution of which is governed by the diffusion data at this point [6]. This chapter describes the implementation of a novel front evolution technique for fibre-tracking using diffusion-weighted MRI, to which the acronym FRET (FRont Evolution Tractography) has been assigned.

6.2 THE FRET ALGORITHM

6.2.1 OVERVIEW

A general outline of the various steps involved in the FRET algorithm is given below, and illustrated in figures 6.1 to 6.7. Each step is described in more detail in the following sections. The FRET algorithm works by iteratively evolving a front from a user-defined location, using the information contained in the diffusion data to infer the appropriate direction of evolution. For each iteration, the surface of the front is characterised by a set of points that adequately samples it. Each of these points will be used to generate a new set of points describing the surface of the next iteration of the front. In this context, a point that is used to generate new points is referred to as a *parent point*, and a point being generated in this way is referred to as a *child point*. The set of child points generated from the same parent is referred to as a *child front*. Thus, each point on the current surface of the front will be used as a parent point to generate its own local child front. The resulting set of child points will eventually constitute the surface of the next iteration of the front.

Information about the strength of the connectivity between the region where the front was initiated (the seed region) and the current point is carried by an associated *index of connectivity*, which is assigned to each child point by considering both its parent's index of connectivity, and the local diffusion data. The contribution of the local diffusion data is derived from the so-called fibre orientation density function (see section 6.2.2). This approach allows the investigation of a large number of potential paths emanating from a particular region, and to some extent to account for the effects of noise and partial volume effects.

A number of other issues need to be addressed for this algorithm to produce adequate results. In particular, only a finite number of directions can be used in practice to generate the set of points making up each child front, and the optimal sampling strategy will be a compromise between the efficiency of the algorithm and the adequacy of the sampling. In addition, not all child points generated during the course of one iteration will lie on the surface of the next iteration of the front. Finally, it is not computationally feasible to use all the child points generated during one iteration as parent points in the next iteration, and a down-sampling strategy must be implemented.

A ‘fast marching’ tractography algorithm [6] has recently been presented that uses a similar concept of front evolution. However, fast marching differs from the FRET algorithm in several important respects, the two most significant of which are as follows. First, fast marching is a voxel-by-voxel technique, which limits the possible directions for tracking to those connecting the centres of two neighbouring voxels (a total of 26 directions), whereas FRET treats both the orientation space and the tensor field as continuous. Second, fast marching uses information from the time of arrival of the front to generate the tracks, whereas the surface of the front as obtained using FRET *is* the final track.

6.2.2 THE FIBRE ORIENTATION DENSITY FUNCTION

The index of connectivity assigned to a point on the front is derived from the local diffusion information, assuming a model for the uncertainty in the direction of the underlying fibres. This model can be expressed in terms of the fibre orientation density function (ODF). This function gives an indication of the likelihood that an underlying fibre tract is aligned along any particular direction, given the local diffusion data. This function dictates the evolution of the front, using the assumptions made in the model. For example, most streamlines-based techniques (e.g. Mori *et al* [1], Conturo *et al* [2], Basser *et al* [3]) follow the direction of the major eigenvector, and hence implicitly use a delta function pointing along that direction as their ODF, in which case both noise and partial volume effects are effectively ignored.

As mentioned in section 3.3.5, the exact relationship between the diffusion-weighted data and the underlying fibre orientation in the presence of noise and partial volume effects is unknown. Moreover, the most adequate model depends on the type of acquisition performed. The q -space approach [11] (see section 3.3.5) involves impractical scan times, and is thus not considered as a suitable candidate for this application. Although high angular resolution imaging may reveal structures that DTI fails to describe, no adequate model yet exists relating the diffusion-weighted signal intensities to the orientation of the underlying fibre population. It has been suggested that a number of distinct diffusion tensors could be fitted to the ADC profile (corresponding to a multi-compartment model in the slow exchange regime) [12]. However, this was found to be unstable when more than two tensors were included in the fitting. It is worth pointing out that the FRET algorithm is not restricted to any particular ODF, and if a more appropriate model relating diffusion to fibre orientation is developed, the particular ODF being used can relatively easily be replaced.

Although the diffusion tensor model has been shown to be inadequate in certain regions of the brain (section 3.3.5), it is well characterised, and its mathematical properties

are well documented. For this reason, the ODF used in this study is an empirically determined function of the local diffusion tensor \mathbf{D} . For a highly anisotropic prolate diffusion tensor, it is assumed that a single fibre population is present, which is aligned along the major eigenvector of the tensor. The ODF should therefore have a narrow distribution, having maximum intensity along the direction of the major eigenvector. In the case of an oblate diffusion tensor, it is assumed that several fibre populations may be present, indicative of either fibre crossing or fibre branching. The ODF should then have a pancake-shaped distribution, with those directions in the plane of the tensor's major axes having similar probabilities. For low anisotropy tensors, it is assumed that the fibres are not arranged coherently, and all directions should have zero or low probability.

The form of the equation for the ODF used in this study was chosen empirically to provide a good approximation to the expected behaviour as described above. At each point, the raw diffusion-weighted intensities are computed by trilinear interpolation from the raw data, and then used to calculate \mathbf{D} by least-squares linear fitting. If $\hat{\mathbf{u}}$ is a unit vector along the direction of interest, then the value returned by the ODF used in this study is given by:

$$P = \alpha \exp[\beta(|\underline{v}| - \lambda_1)] \quad (6.1)$$

where:

$$\begin{aligned} \alpha &= \frac{1}{1 + \exp[a_1(a_2 - \text{FA})]} \\ \beta &= \frac{a_3}{1 + \exp[a_4(a_5 - \text{FA})]} \\ \underline{v} &= \mathbf{D} \cdot \hat{\mathbf{u}} \end{aligned}$$

$|\underline{v}|$ is the magnitude of the vector \underline{v} , λ_1 is the major eigenvector of the diffusion tensor, and FA is the value of the fractional anisotropy [13] at the point of interest. Since the maximum value of $|\underline{v}|$ is λ_1 , the expression $\exp[\beta(|\underline{v}| - \lambda_1)]$ has maximum value 1 when the direction of interest $\hat{\mathbf{u}}$ corresponds to the major eigenvector, and tends to zero as the magnitude of \underline{v} decreases. For isotropic tensors, $|\underline{v}| = \lambda_1$, and the index of connectivity is constant for all directions. For anisotropic tensors, $|\underline{v}|$ is smaller than λ_1 for any direction other than the major eigenvector, and the index of connectivity for those directions will be smaller, depending on the value of β . To ensure that a small (near zero) index of connectivity is given for isotropic tensors, the factor α is a function of anisotropy. To ensure that highly anisotropic tensors yield narrow, well-defined ODFs, β is also a function of anisotropy.

The parameters $\{a_i\}$ control aspects of the behaviour of this ODF, for example its width and amplitude given a particular diffusion tensor. These parameters were set em-

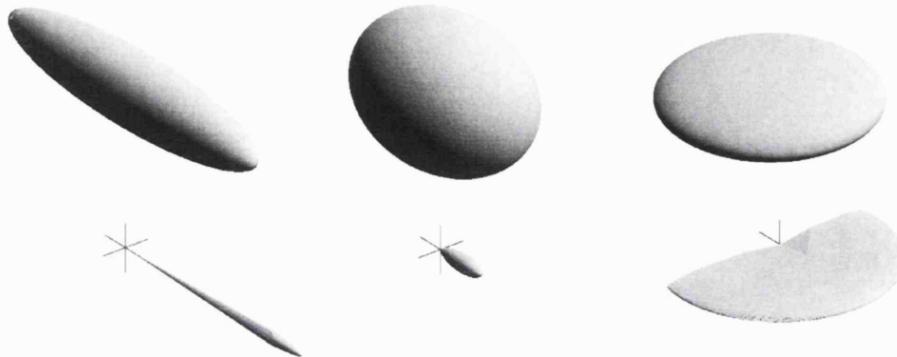


Figure 6.1: example diffusion tensors (top) and their corresponding fibre orientation density functions (bottom) generated as described in section 6.2.2. Left: a highly anisotropic tensor representing a fibre population for which the orientation is relatively well defined. Middle: a low anisotropy tensor gives rise to a much broader distribution with a reduced index of connectivity. Right: an oblate tensor, representing the case of crossing fibres, where only the plane likely to contain the fibre orientation can be determined.

pirically, such that the ODF generated for a particular tensor displayed the desired behaviour. Three such tensors were used in this process: a low anisotropy prolate tensor, a high anisotropy prolate tensor, and an oblate tensor. In this way, the parameters $\{a_i\}$ were set to $\{10, 0.3, 200, 6, 0.5\}$. Figure 6.1 shows the three tensors used in this process, along with the corresponding ODFs generated using the final set of parameters; as can be seen, the chosen empirical ODF describes the expected behaviour adequately.

6.2.3 SAMPLING ORIENTATION SPACE

Each child front consists of a set of points, each obtained by stepping away from the parent point by a small distance. In order for the algorithm to produce optimal results, as many directions as possible should be sampled. However, sampling too many directions would imply that each child front is made up of a large number of points, increasing the amount of computation required and impacting on the efficiency of the algorithm. Since it is the most likely directions that are of interest, it is more efficient to take more samples along those orientations than along others.

Since the ODF used in this study is based on the diffusion tensor, the most likely orientation for the fibre will coincide with the major eigenvector. A suitable sampling scheme can then be generated as illustrated in figure 6.2. A set of unit vectors that sample orientation space evenly is generated (by tessellation of a regular polyhedron, yielding 66 vectors). Each vector is then multiplied by the tensor: if the tensor is prolate, the vectors will tend to align with its major axis. If the tensor is oblate, the vectors will tend to move into the plane of the two major axes. The directions of this set of vectors are then fed into

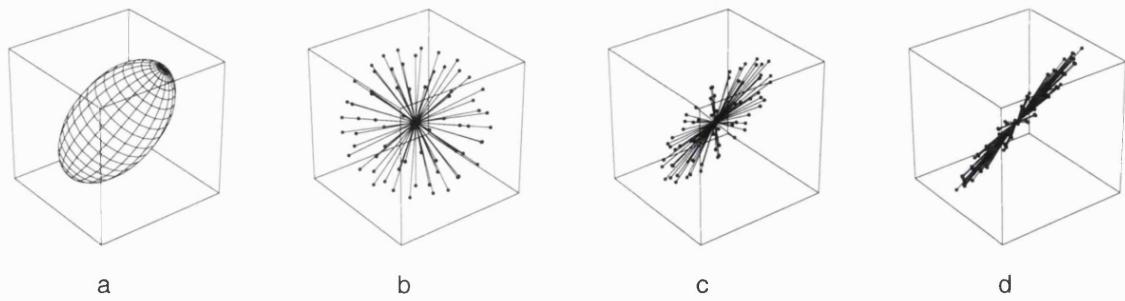


Figure 6.2: given a typical prolate diffusion tensor (a), the most likely orientations can be sampled more densely through the following process. Starting from a set of vectors that sample orientation space uniformly (b), each vector is multiplied by the tensor raised to the power n , such that it tends to align with the tensor's major axis. The resulting set of vectors with $n = 2$ (c) and $n = 4$ (d).

the ODF to generate the child front.

The sampling density can be increased further by using a modified version of the diffusion tensor with an increased anisotropy for the multiplication process. This can be done easily by raising the tensor to a power n , determined empirically (in this study, $n = 4$). Thus, the set of vectors $\{\underline{u}_i\}$ used to sample the ODF is given by:

$$\underline{u}_i = \mathbf{D}^n \cdot \underline{r}_i \quad (6.2)$$

where $\{\underline{r}_i\}$ is the set of evenly distributed vectors.

In addition, a curvature constraint was introduced to prevent the tracks generated from turning back on themselves. A threshold was set on the angle between each vector in the set $\{\underline{u}_i\}$ generated above and the current direction of tracking (see below). If this angle was exceeded, the corresponding point in the child front was discarded. The threshold depends on the minimum radius of curvature specified by the operator and the step size used in the algorithm, which is also specified by the operator (in general, it was found that a minimum radius of curvature of 2 mm produced adequate results).

6.2.4 FRONT EVOLUTION

As previously mentioned, the front is defined by a set of points lying on its outermost surface. During each iteration, a child front is generated from each of the parent points that currently define the front (see figure 6.3), and all child fronts are subsequently merged to form the front for the next iteration (see section 6.2.5 below). Each point on the child front has an associated direction of sampling (the line joining it to its parent point), a direction of evolution (the normal to the front at this point), and an index of connectivity

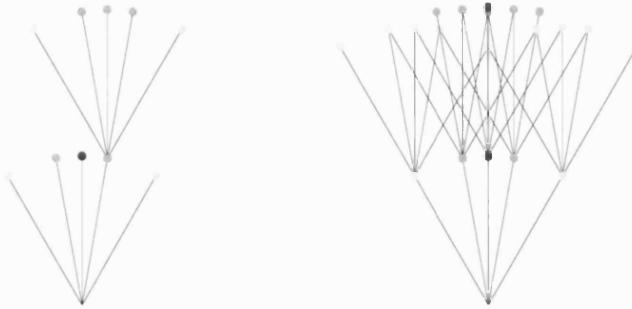


Figure 6.3: overview of the front evolution process, for the case of a single seed point. Left: the seed point yields a child front, consisting of a set of points, each of which has been assigned its index of connectivity (represented by the grey levels in this example). Each of these points in turn generates its own child front. Right: the set of points thus generated form a crude representation of the outer surface of the front.

with the seed point. There are a number of ways that these child fronts could be generated, and these have profound implications for the properties of the algorithm.

It is possible to make the speed of evolution a function of the ODF, such that the time of arrival of the front gives an indication of the level of connectivity to that region (see figure 6.4). This is similar to the approach taken by Parker *et al* [6]. In this case, it is not necessary to maintain information about the index of connectivity explicitly, since the time of arrival contains that information. Although this approach is conceptually more rigorous, it suffers from a number of drawbacks. It was used for the first implementation of the algorithm, written using MatLab (The MathWorks Inc, Natick, MA). This implementation was found to be extremely inefficient¹, and will not be discussed here.

The approach used for the present implementation of the FRET algorithm is to make the speed of evolution constant for all directions, and so use a constant step size (see figure 6.4). In this case, the index of connectivity is carried by each point explicitly, and the child points lie on the surface of a sphere centred on their parent point. Although using this approach removes some of the limitations of the other type of front evolution, it also has disadvantages. In particular, problems will arise when tracking a curved fibre tract, since highly connected points may be discarded in favour of less highly connected points by the surface reconstruction step. However, this is a problem that can be overcome, as described in the next section. In order to ensure its optimal efficiency, this implementation of the algorithm was written in C++.

¹MatLab is an interpreted language, and relatively inefficient for anything other than matrix operations. In addition, its syntax is such that a number of possible optimisations could not be implemented. Finally, using a variable speed of evolution is intrinsically more computationally intensive.



Figure 6.4: two possible approaches for front evolution. Left: using an ODF-dependent step size, the front evolves most rapidly along the most likely direction. The index of connectivity is then a function of the time of arrival of the front. Right: using a constant step size, the front propagates at the same rate in all directions, and the index of connectivity (represented by the grey levels in this example) has to be carried explicitly.

6.2.5 SURFACE RECONSTRUCTION

During the course of an iteration, a large number of child fronts will be generated, and many of these will inevitably overlap. When these child fronts are merged, those portions that lie behind or inside the surface of other child fronts need to be eliminated, such that the surface of the front is clearly defined.

In order to reconstruct the surface of the front, it is necessary to identify those points that lie within another child front. Since the constant speed of evolution approach was used, each child front is spherical in shape, and it is only required to test whether the distance between the point of interest and the child front's parent point is less than the child front's radius. The simplest way of reconstructing the surface would be to remove those points that lie inside other child fronts. This is however not a satisfactory solution. Consider the case of a tract with a relatively large curvature, as shown in figure 6.5. A point with a high index of connectivity, generated from a parent point inside the tract, may find itself inside a child front with an overall lower index of connectivity, originating from a parent point outside that tract, and would thus be discarded. In other words, such an approach would be biased towards low curvature tracks. This can be avoided as follows. Any point that lies inside another child front is projected along the line joining it to its parent point onto the surface of that child front, as illustrated in figure 6.6. This ensures that the surface of the front is clearly defined, since all points will now lie on the outermost surface of the front. However, points will be projected further and further away from their original position as the direction of projection tends towards the plane of the front. This problem can be minimised by introducing a constraint on the maximum angle allowed in the tracks: a threshold can be imposed on the angle between the current direction of tracking at the point of interest (given by the parent point's normal vector)

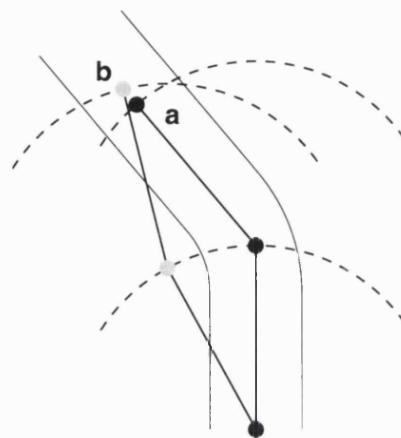


Figure 6.5: discarding points that lie inside other child fronts may introduce bias in the track generated, as illustrated here. In curved regions of the track, a point with a high index of connectivity (a) may be overtaken by a point with a low index of connectivity (b), and would thus be discarded. Curved tracks would therefore be assigned unduly low indices of connectivity.

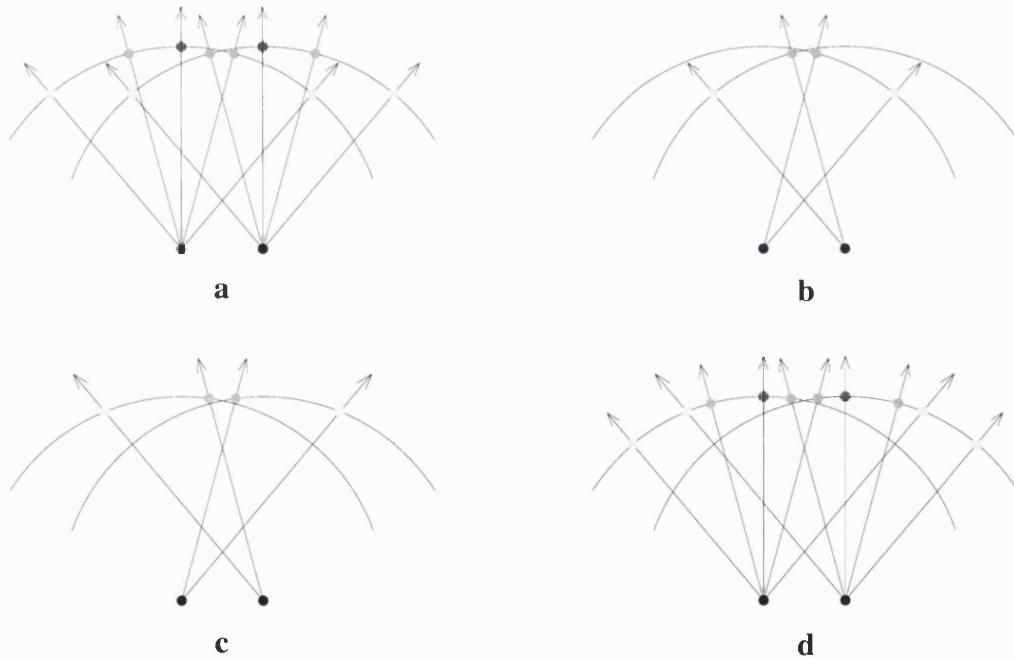


Figure 6.6: (a): there will in general be some overlap between any two child fronts (the arrows indicate the direction of tracking associated with each point). (b): those points that lie behind another child front are identified. (c): the points in (b) are then projected onto the surface of the outermost child front, retaining their original index of connectivity and direction of tracking. (d): all these points taken together define the surface of the front for the next iteration.

and the line joining the child point to its parent. Since the plane of the previous generation's front (the set of parent points) should be approximately parallel to the plane of the current generation's front, this should ensure that the directions along which the points are projected remain close to normal to the front. Points that do not satisfy this condition are discarded. This is equivalent to constraining the maximum curvature of the track generated, since the radius of curvature is a function of that angle and the step size. As mentioned in section 6.2.3, such a constraint is already used to generate the child fronts, and therefore does not need to be re-implemented.

As previously mentioned, every point on the front has an associated vector, giving the direction normal to that point's original child front at the point's location (corresponding to the direction joining that point to its parent). Points that remain after the surface reconstruction step retain their original normal vector. Thus, the set of remaining points describe the position of the reconstructed front, and the set of normal vectors at these points describe the normal to the front. A consequence of this is that it is possible for two points close to each other to have different 'normal' vectors. Although it may appear desirable to eliminate such discrepancies, it is not clear how such a smoothness constraint can be implemented whilst minimising the impact on the efficiency of the algorithm, and what its real benefits might be in practice.

6.2.6 DOWN-SAMPLING THE FRONT

If every child point was used as a parent point for the next iteration, the number of points making up the front would increase approximately exponentially, and the computation would soon become too slow to be useful, ultimately failing when the computer's memory has been filled entirely.

In order to reduce the number of points making up the front, and thus to improve the efficiency of the algorithm, the set of points that contain the most significant information must be identified. Given a set of points located very close to each other, the most significant point must be that with the highest index of connectivity, otherwise high probability tracks would be discarded arbitrarily. Therefore, the set of points that contains the most significant information consists of those points that have the highest index of connectivity within their immediate vicinity.

An appropriate subsample can be taken from the set of points currently making up the front as shown in figure 6.7. A 3-dimensional grid is used to divide the set of points into small subsets of points lying close together (i.e. within the same grid node). The subsample can then be formed by selecting the point with the highest index of connectivity from each subset. The size of the grid must be chosen to be small compared to the

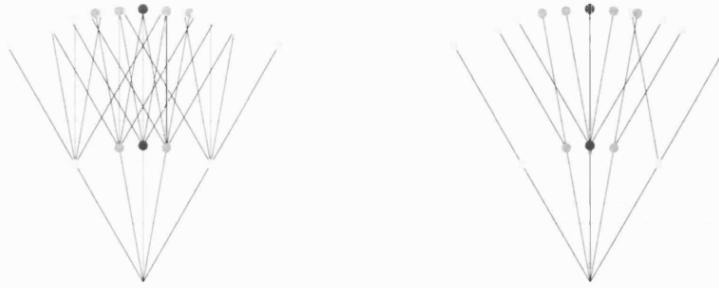


Figure 6.7: the points making up the front need to be sampled down to a number that is computationally feasible. Left: a grid is applied to the front. Right: the point with the highest index of connectivity (represented by the grey levels) within each cell is identified and all others are discarded.

step size of the algorithm, whilst ensuring that the number of points in the front remains manageable. In this study, it was found that a value for the size of each grid element of approximately 15% of the step size gave adequate results.

6.2.7 PROPAGATION OF THE INDEX OF CONNECTIVITY

Assigning an index of connectivity to each point can be done in a number of ways. If each track is seen as a series of probabilistic steps, then the probability of getting from a point p_n along the track is given by the product of the probability of getting to the previous point p_{n-1} and the probability of stepping onto p_n . Therefore, it may initially seem appropriate to assign to each newly-generated point the product of the index of connectivity of its parent point with the value returned by the ODF. However, this simple treatment assumes that there is only one path linking two points together. In the present framework, there are an infinite number of such paths, and the contribution of each one of them would need to be included, which is not possible in practice. In the present algorithm, the connectivity index assigned to each child point is derived from its parent point and the local diffusion information only; it would be very difficult to include contributions from the other parent points in the vicinity.

The method adopted in this implementation of the algorithm attempts to take both the information about the diffusion locally (via the ODF) and information about the current status of the track (via the parent point's index of connectivity) into account in order to assign an appropriate index of connectivity to each child point. A child point cannot be more highly connected with the seed region than its parent, since the connection must pass through its parent. Each child point is therefore assigned the value returned by the ODF, unless that value is higher than its parent's index of connectivity, in which case it is assigned its parent's index of connectivity. In this way, the index of connectivity of any

point on the front corresponds to the minimum value returned by the ODF at any point along the path from the seed point. Although this index is not directly related to the actual probability of the connection, it still provides useful information regarding the likelihood of a track: if any step along a track is uncertain, then any connections emanating from that point should be regarded as equally uncertain.

6.2.8 SEED POINT SELECTION

In the simplest case, a track is initiated from a single seed point. However, it is also possible to define a set of points that can be used as the first iteration of the front. This approach minimises the dependence of the tracks generated on the exact location of the initial seed point. In this study, a square grid of $N \times N$ points was used (with N ranging from 1 to 9), lying in the plane perpendicular to the initial direction of tracking. The grid spacing used was the same as that used to down-sample the front (see section 6.2.5).

6.3 METHODS

6.3.1 DATA ACQUISITION

Two data sets were acquired over two separate occasions on a healthy adult volunteer on a 1.5T Siemens Vision system, using the twice-refocused diffusion-weighted spin echo EPI sequence described in section 4.2. The parameters of the sequence were: TE = 110ms, FOV = 256×256 , matrix = 128×128 , slice thickness = 2 mm (voxel dimensions = $2 \times 2 \times 2$ mm 3), 60 contiguous slices covering the whole brain, with a b -value = 1000s/mm 2 . The diffusion encoding gradients were applied in 20 uniformly distributed directions [14], and three non-diffusion encoded ($b=0$) data sets were also acquired.

6.3.2 ASSESSMENT OF THE FRET ALGORITHM

Ideally, simulations would be used to assess the properties of the FRET algorithm. However, the simulations performed in chapter 5 were designed to evaluate streamlines algorithms that generate a single track per experiment. New measures of reliability would be needed to assess the results produced by the FRET algorithm, and these may or may not be comparable to the existing measures. For this reason, the FRET algorithm was initially assessed as described below.

The performance on the FRET algorithm was assessed first by performing tracking on real data acquired as described above, in regions with relatively well known anatomy, and

comparing the tracks obtained to an anatomical atlas. The behaviour of the algorithm with respect to partial volume effects could be appreciated by looking at regions known to contain multiple fibre populations, such as the centrum semiovale and the pons. This process was repeated in the second data set to assess the reproducibility, and hence robustness to noise, of the algorithm. The effect of changing the parameters of the algorithm (the step size, the size of the seed region, the curvature threshold, and the parameters defining the ODF) were investigated by repeating the tracking procedure with identical parameters apart from that of interest. Finally, the tracking results were compared to tracks generated using the more common streamlines technique, described below.

6.3.3 STREAMLINES TRACKING ALGORITHM

For comparison with the FRET algorithm, a commonly used streamlines algorithm (see section 3.3.3) was also used to generate tracks. Starting from a given seed point, each track is generated as follows:

- the raw diffusion-weighted intensities are calculated at the current point along the track from its 8 nearest neighbours by tri-linear interpolation
- the diffusion tensor is calculated from these by least-squares linear fitting
- the FA is evaluated from the tensor: if it falls below a threshold value of 0.1, the track is terminated
- the major eigenvector is obtained from the tensor
- the next point along the track is generated by stepping a distance of 0.1 mm along the direction of the major eigenvector
- if the angle formed between the new track segment generated and the previous one corresponds to a radius of curvature of less than 2 mm, the track is terminated
- this process is repeated until the track is terminated.

Tracks were initiated from a set of points located on a $5 \times 5 \times 5$ grid, forming a 1 mm wide cube in 3D space centred on the seed point of interest.

6.4 RESULTS AND DISCUSSION

The FRET algorithm was designed to address some of the limitations of previously published techniques, in particular susceptibility to noise (by using the ODF), problems with

fibre branching (by using a front evolution approach), and the lack of any indication of the likelihood of a connection (by introducing an index of connectivity). However, tractography techniques are difficult to validate, due to the lack of a gold standard. In this study, tracking was performed in structures with relatively well known anatomy, and the results presented show that this technique can be used to generate tracks in the major white matter tracts that correspond well to this anatomy. The examples shown were selected to demonstrate some of the characteristics of the algorithm, which are discussed below. A full assessment of the FRET algorithm will require further work, in particular using simulations on synthetic data such as those described in chapter 5.

The quality of the tracks generated can be appreciated from figures 6.8, 6.9 and 6.10, showing the results of tracking in the cortico-spinal tract and the splenium of the corpus callosum. It can be seen that in regions where the fibres are known to be highly coherently oriented (e.g. the posterior limb of the internal capsule and the splenium of the corpus callosum), the track remains well defined. As desired, subsequent possible branches are identified and assigned their corresponding index of connectivity.

An important property of any algorithm is that if it establishes a connection from one point to another, we expect it also to be able to establish a connection from the latter point to the former. This was investigated by placing the seed region in an area that had been assigned a high index of connectivity in a previous tracking experiment, the results of which are shown in figure 6.8. It can be seen in figure 6.11 that the track seeded in this area projects strongly to the structure where the seed region for the original track in figure 6.8 was placed.

6.4.1 DEPENDENCE ON THE SIGNAL TO NOISE RATIO

To investigate the reproducibility of the FRET algorithm, tracks were generated from two different data sets acquired on the same volunteer. The results are shown in figures 6.8 and 6.9, and it can be seen that both yield very similar tracks. The main differences between them are that some of the branches with low connectivity indices are present in only one of the data sets, and that the values for the connectivity indices are subtly different. This may be attributed to the fact that these data sets differ in their noise and partial volume distributions (since they were acquired on separate occasions), such that the estimated diffusion tensors differ between the two acquisitions, leading the algorithm to establish low connectivity connections in one data set but not in the other. This is most noticeable as the track enters the region of the centrum semiovale, a region known to contain crossing fibres and significant partial volume effects [15, 16].

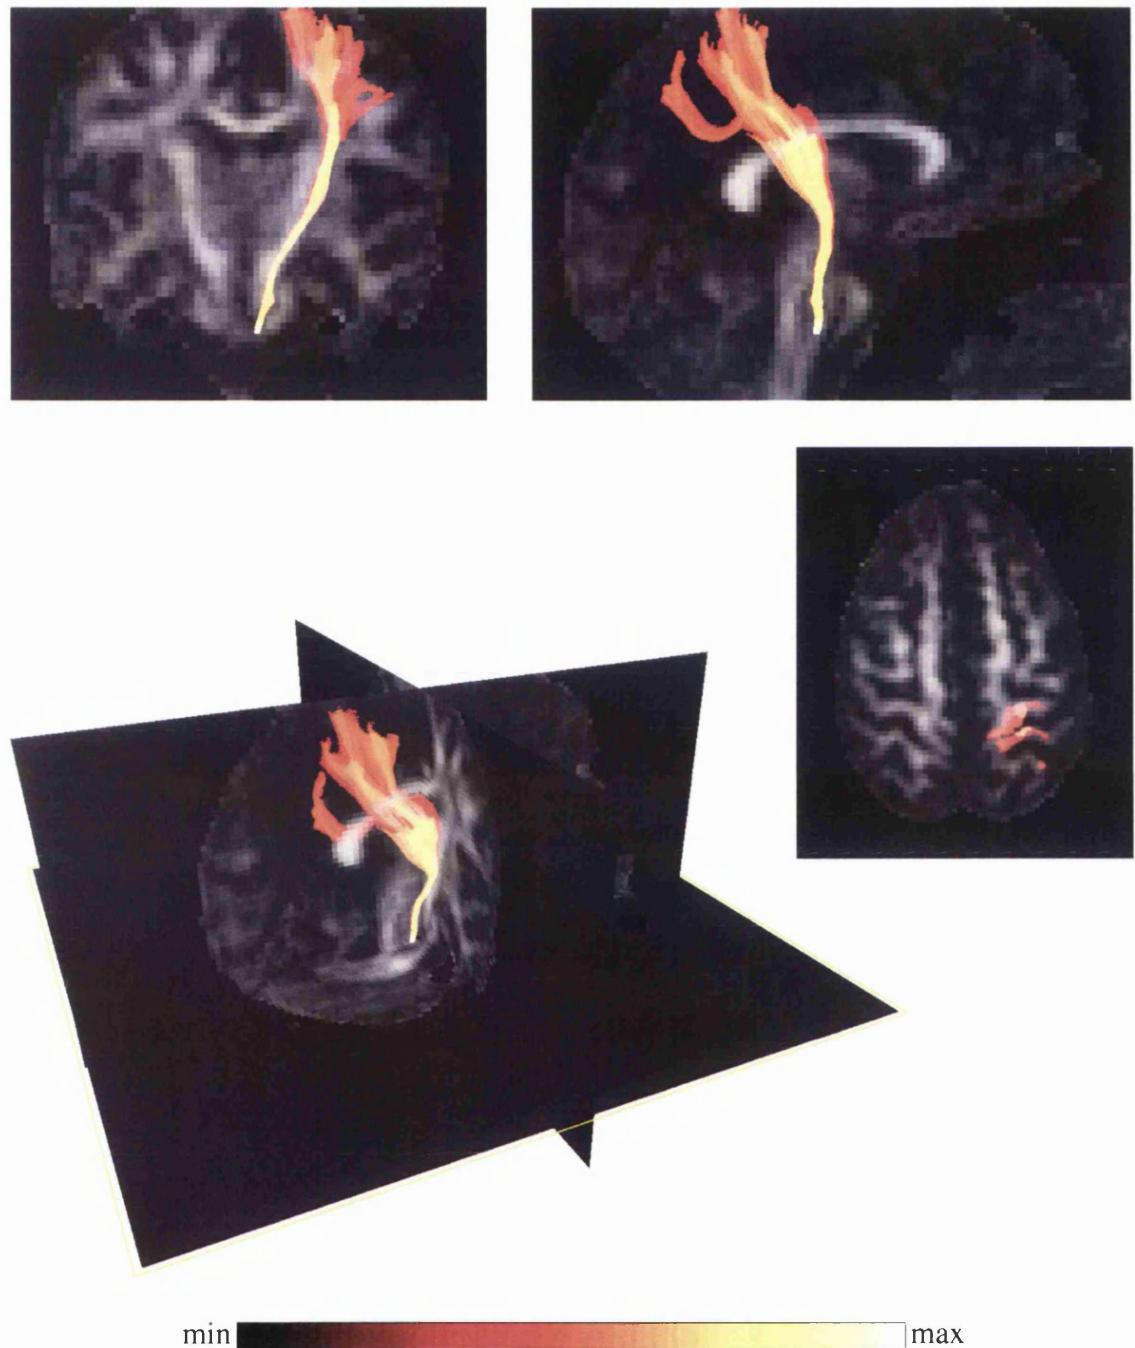


Figure 6.8: a track generated using the FRET algorithm by seeding in the right cortico-spinal tract at the level of the pons. The colour corresponds to the index of connectivity as shown on the bottom right. Top: coronal and sagittal maximum intensity projections (MIPs) of the track, overlaid on fractional anisotropy (FA) maps. Middle right: an axial MIP of the portion of the track above the level of the slice, overlaid on an axial FA map at the level of the right motor cortex, showing the region of highest connectivity. Bottom: a MIP of the track, overlaid on three orthogonal FA maps.

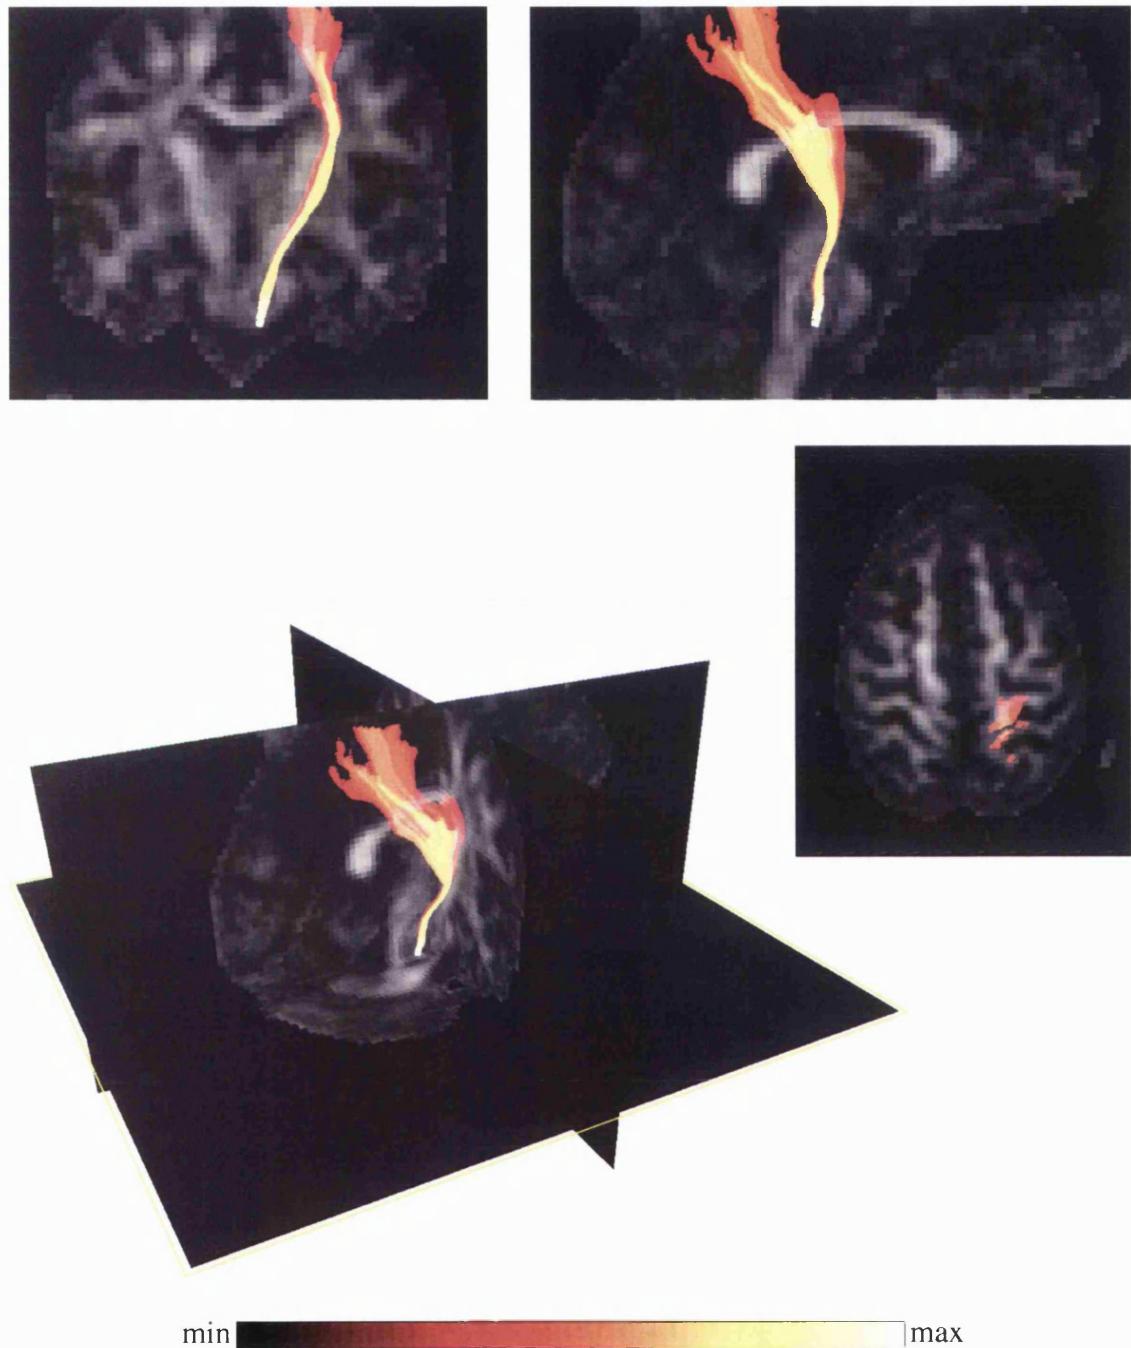


Figure 6.9: a track generated using the FRET algorithm by seeding in the right cortico-spinal tract, analogous to that displayed in figure 6.8, but obtained using a data set acquired from the same volunteer several months later. The track is displayed and colour-coded as in figure 6.8.

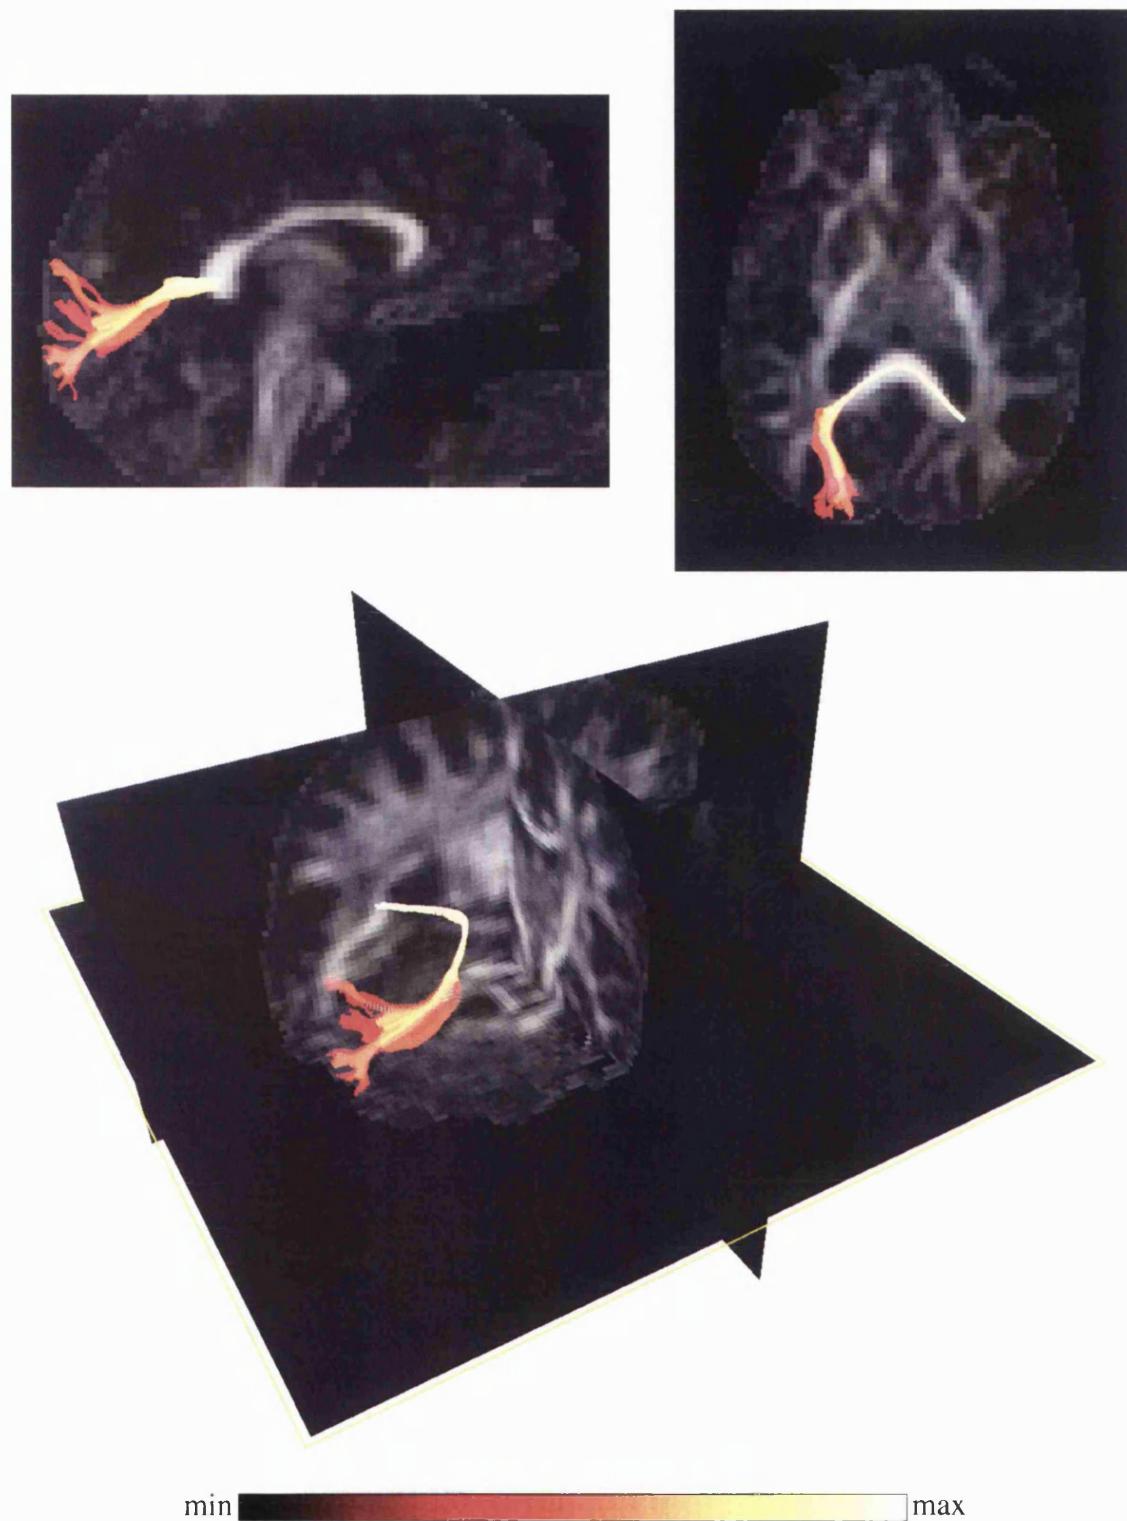


Figure 6.10: a track generated by seeding on the left side of the splenium of the corpus callosum. Top left: a sagittal MIP projected onto a sagittal FA map. Top right: an axial MIP projected onto an axial FA map, viewed from below such that left is displayed on the right. Bottom: a MIP of the track, overlaid on three orthogonal FA maps. The track is colour-coded as in figure 6.8.

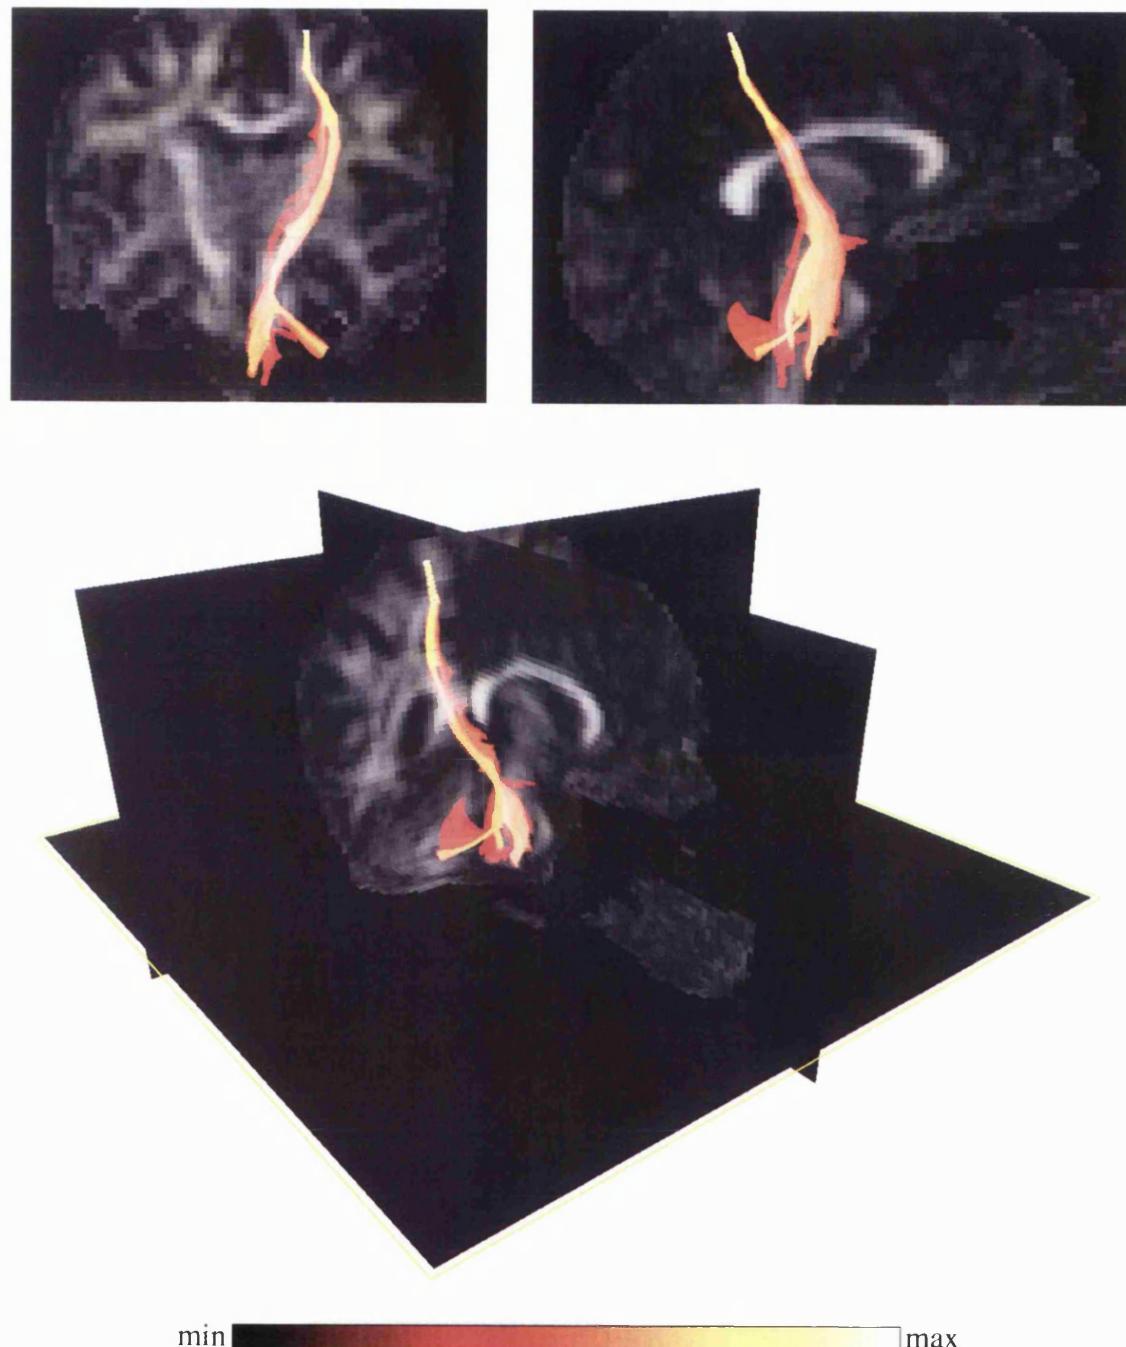


Figure 6.11: a track generated using the FRET algorithm by seeding in the white matter adjacent to the right motor cortex, in a region indicated as having a high index of connectivity when seeding in the cortico-spinal tract (see figure 6.8). Top left: a coronal MIP projected onto a coronal FA map. Top right: a sagittal MIP projected onto a sagittal FA map. Bottom: a MIP of the track, overlaid on three orthogonal FA maps. The track is colour-coded as in figure 6.8.

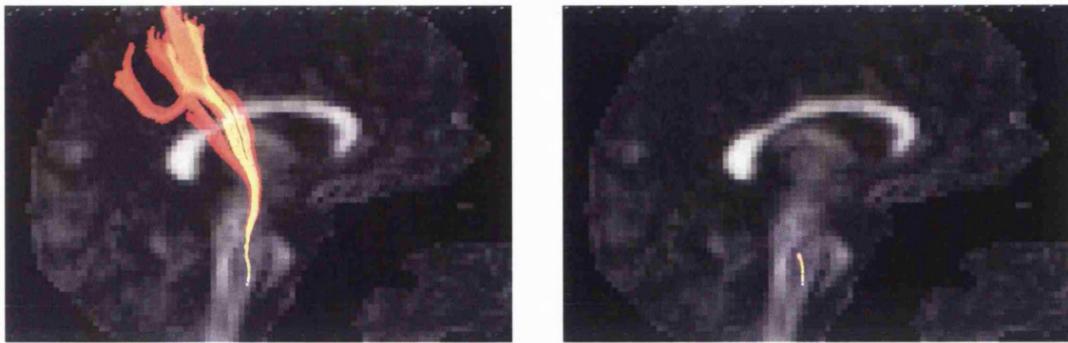


Figure 6.12: tracks generated using the FRET algorithm to show the sensitivity of the track to the exact position of the seed point. Left: a MIP of the track generated by placing a single seed point in the centre of the seed region used to generate figure 6.8. Right: a track generated by placing a single seed point 1 mm away from the previous seed point. As can be seen, the track is terminated at the level of the pontocerebellar fibres.

6.4.2 DEPENDENCE ON INITIAL SEED REGION

Figure 6.12 was produced to illustrate the dependence of the tracks generated on the exact location of the initial seed point or region (see section 6.2.8). At the level of the pons, the corticospinal tract runs very close to a number of other white matter tracts, in particular the pontocerebellar fibres that run across to form the middle cerebellar peduncle. This region is therefore very prone to partial volume effects, making the tracking process unreliable. In figure 6.12, the seed region was placed in the corticospinal tract, just below the level of the pontocerebellar fibres. When seeding from a *single* seed point in this region, the tracks generated by the FRET algorithm are very sensitive to its exact location: starting the FRET algorithm from a point only a millimetre away from a previously successful seed point results in the front being terminated at the level of the pons. When using a grid of seed points for the seed region (see section 6.2.8), this dependence is markedly reduced, since the seed region is more likely to contain an appropriate starting point. However, the size of the seed grid should be kept small enough to fit within the structure of interest to avoid tracking adjacent white matter pathways and producing erroneous results. For the results shown here, the maximum extent of the grid was 1.2 mm.

6.4.3 DEPENDENCE ON ALGORITHM PARAMETERS

The properties of a track, in terms of its dispersion, connectivity index, or the number of branches, depend on the particular model used and its assumptions. Figure 6.13 shows tracks generated using different values for the parameters $\{a_i\}$ used to define the ODF, resulting in a wider ODF ($\{a_i\} = \{14, 0.25, 100, 16, 0.3\}$). As expected, the track gener-



Figure 6.13: a track generated using the FRET algorithm from seed points in the cortico-spinal tract, analogous to that in figure 6.8, but obtained using a modified ODF (as described in section 6.4.3). Left: a coronal MIP of the track overlaid on a coronal FA map. Right: a sagittal MIP of the track overlaid on a sagittal FA map. The track is colour-coded as in figure 6.8. This track contains more branching and dispersion than that in figure 6.8. However, the regions that have been assigned the highest index of connectivity are the same as those in figure 6.8.

ated exhibited a higher degree of dispersion and contained more branches, and these were assigned a higher index of connectivity. The region of highest connectivity was however the same as when using the unmodified ODF. Using such a broad ODF thus increased the amount of computation performed by the algorithm, without adding a significant amount of information.

The effect of each a_i parameter on the ODF and the resulting track may be appreciated by looking at equation 6.1. The peak intensity for the ODF is determined by α , and its ‘width’ by β . Both α and β are simple sigmoid functions that decrease as a function of anisotropy: a_1 and a_2 correspond respectively to the sharpness and crossing point of the α curve, and a_3 , a_4 , and a_5 correspond respectively to the amplitude, sharpness and crossing point of the β curve. This allows fine control over the properties of the ODF used in the present study. As mentioned above, the $\{a_i\}$ parameters used in this work were chosen empirically. However, it should be possible to determine the values to use for these parameters more rigorously by performing simulations on synthetic data, where the true path of the fibre tract is known. Such simulations would also help to determine appropriate values for the curvature and index of connectivity thresholds, and to validate the FRET algorithm, in the absence of any true ‘gold standard’. Simulations to determine the optimal values for these parameters and to assess the limitations and requirements of the FRET algorithm more rigorously will be the subject of further work (cf. chapter 5 [17]).

Figure 6.14 illustrates the effect of modifying the threshold for the minimum radius of curvature. This threshold is a statement about the expected properties of fibres in the

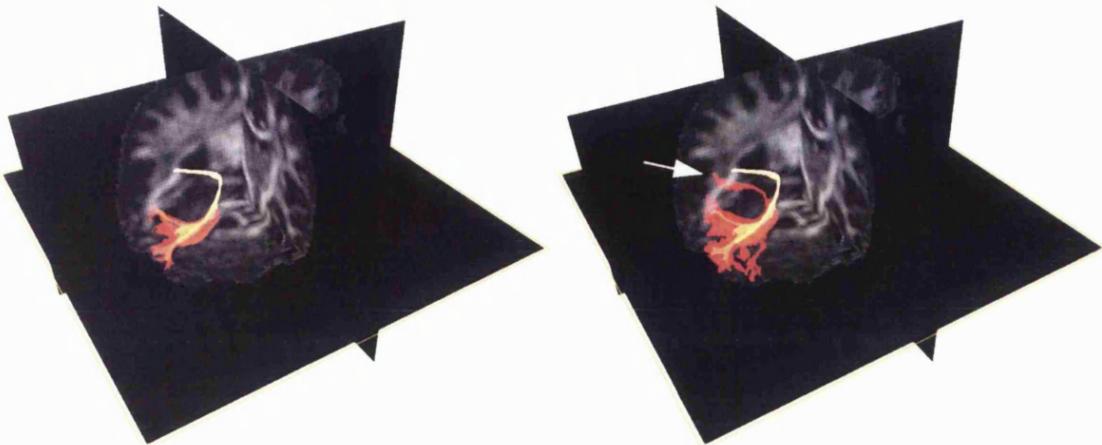


Figure 6.14: tracks generated using the FRET algorithm from seed points in the left side of the splenium of the corpus callosum, analogous to that in figure 6.10. Left: using a radius of curvature threshold of 2 mm. Right: using a radius of curvature threshold of 1 mm. Note the presence of an additional branch curving away from the original direction of tracking in an improbable way. The tracks are colour-coded as in figure 6.8.

brain, and changing it allows a different degree of branching and dispersion to occur in the track. When small radii of curvature are allowed, the track exhibits more dispersion, and branches are present that would not have been allowed using a higher threshold. With a minimum radius of curvature of 1 mm, a branch (highlighted by the arrow) is generated that curves almost 90° away from the original direction of evolution of the front, in an improbable way. This branch is removed using a threshold of 2 mm or higher. It should be noted that the voxel size in the data set was 2 mm, and that structures with a radius of curvature lower than this value cannot be expected to be represented accurately. The use of a curvature threshold can therefore prevent this kind of error. It may also be used to state prior assumptions about the tract of interest, for example in a major white matter pathway that is not expected to contain high curvature.

6.4.4 DEPENDENCE ON PARTIAL VOLUME EFFECTS

The effect of partial voluming between different fibre populations can be appreciated from figures 6.8 and 6.11. The anatomical regions where most dispersion in the tracks occurs correspond to areas of known fibre crossing, which are likely to suffer from a high degree of partial volume contamination [15, 16]. In particular, connections are seen in figure 6.8 posterior to the sensorimotor cortex that are anatomically implausible. These are likely to have been caused by partial volume effects in the region of the centrum semiovale, causing the track to branch into an adjacent, yet unconnected tract. Since the ODF used in this study is based on the diffusion tensor, it is expected to behave relatively poorly

with respect to partial volume effects (see later). However, these contentious connections have been assigned a relatively low index of connectivity, indicating that this measure is indeed a useful marker of connectivity.

The track in figure 6.11 also establishes some implausible connections. First, a connection is established to the medial lemniscus, which is known to project to the thalamus. Within the midbrain, the medial lemniscus runs close to the cortico-spinal tract, and it is likely that this connection was established due to partial volume effects. Second, the track also establishes a connection to the ipsilateral cerebellum via the middle cerebellar peduncle, which is not anatomically connected to the motor cortex. This artefact is again most likely due to partial volume effects between the cortico-spinal tract and the adjacent pontocerebellar fibres, and is aggravated by the fact that the angle between the fibres in these structures is relatively small. The diffusion tensor in such regions will remain fairly anisotropic and prolate, such that these regions can only be assumed from the data available to contain a single fibre population.

6.4.5 COMPARISON WITH STREAMLINES

The tracking results obtained using the FRET algorithm were compared with results obtained from the streamlines technique, as shown in figure 6.15. The tracks generated using the streamlines technique agree well with those generated using the FRET algorithm. In particular, the track obtained using the FRET algorithm contains all the tracks generated from the streamlines technique. Moreover, the streamlines-generated tracks correspond to the portion of the front that has been assigned the highest index of connectivity: this is expected since streamlines only generates the most likely track emanating from each seed point given the data. A limitation of the streamlines techniques is the lack of any indication of the reliability of the tracks thus generated, since all are effectively given the same weighting. In contrast, it can be seen that the FRET algorithm does provide additional information about the degree of connectivity of these and other possible tracks.

6.4.6 THE INDEX OF CONNECTIVITY

The algorithm presented in this study defines an index of connectivity (see sections 6.2.2 and 6.2.7) as a measure of the reliability of the connections established by the algorithm. Ideally, the actual probability of the connection given the diffusion-weighted data would be produced. However, this is very difficult to achieve in practice, primarily for the following two reasons. First, the method chosen to propagate the index of connectivity (see section 6.2.7) is intrinsically unable to produce suitable estimates of the probability of a connection. Second, the probability of a connection is directly dependent on the ODF



Figure 6.15: a sagittal MIP of the same track as in figure 6.8, together with the tracks (shown in blue) generated using the streamlines algorithm (see section 6.3.3), by seeding in the same region. Both are overlaid on a midline sagittal FA map. The track generated by the FRET algorithm is colour-coded as in figure 6.8.

used and the model from which it is derived. As mentioned in the introduction, there is currently no model that would take the effects of both noise and partial volume into account to provide an adequate estimate of the probability density function for the underlying fibre orientation.

It should be emphasised that the FRET algorithm is not restricted to a particular ODF, and the latter can easily be replaced if a more suitable model is developed. The current implementation of the ODF used here is a function of the diffusion tensor. Although inadequate in many cases, the tensor formalism is well accepted and its properties can be derived in a relatively straightforward fashion. Given this model and its assumptions, the major eigenvector is expected to correspond to the most likely orientation for the underlying fibre bundle. However, the exact relationship between the diffusion tensor and the fibre orientation is still not clear. Ways of accounting for the effects of noise have recently been presented [5, 18], but it is still not possible to account for partial volume effects in a rigorous way. In the present study, the ODF was therefore empirically chosen to reflect its expected behaviour. It should also be mentioned that since all images used in this study contain similar amounts of noise, the actual value of the SNR was not used as a parameter in the derivation of the ODF. It is thus clear that a more rigorous ODF

is required in order to obtain better estimates of connectivity, especially when tracking through regions containing partial volume effects.

Despite these limitations, the index of connectivity used here provides useful information. Connections that are not anatomically plausible have been assigned low values for their index of connectivity, whereas the most likely connections have been assigned a high index of connectivity.

6.5 CONCLUSION

As has been shown in this chapter, the FRET algorithm possesses many of the properties that are desirable in a tracking algorithm: it produces results consistent with known anatomy, and provides additional information regarding other possible branches and their degree of connectivity with the seed region. It attempts to address many of the issues previously raised, in particular fibre branching and susceptibility to noise and partial volume effects. One of the main limitations of the current implementation of the FRET algorithm is that it is still susceptible to partial volume effects. However, if a more adequate model for diffusion is developed, the improved ODF may be used as a replacement. Finally, further work is needed to assess its properties and limitations more thoroughly, using simulations such as those presented in chapter 5. In this case, new measures of reliability applicable to FRET would need to be defined, since the current measures apply only to streamlines.

REFERENCES

- [1] MORI, S., CRAIN, B. J., CHACKO, V. P., AND VAN ZIJL, P. C. M. Three-dimensional tracking of axonal projections in the brain by magnetic resonance imaging. *Annals of Neurology* 45 (1999), 265–269.
- [2] CONTURO, T. E., LORI, N. F., CULL, T. S., AKBUDAK, E., SNYDER, A. Z., SHIMONY, J. S., MCKINSTRY, R. C., BURTON, H., AND RAICHLE, M. E. Tracking neuronal fiber pathways in the living human brain. *Proceedings of the National Academy of Sciences* 96 (1999), 10422–10427.
- [3] BASSER, P. J., PAJEVIC, S., PIERPAOLI, C., DUDA, J., AND ALDROUBI, A. In vivo fiber tractography using DT-MRI data. *Magnetic Resonance in Medicine* 44 (2000), 625–632.
- [4] TUCH, D. S., BELLIVEAU, J. W., AND J., W. V. A path integral approach to white matter tractography. In *Proceedings of the ISMRM* (2000), vol. 8, p. 791.
- [5] JONES, D. K. Determining and visualizing uncertainty in estimates of fiber orientation from diffusion tensor MRI. *Magnetic Resonance in Medicine* 49 (2003), 7–12.
- [6] PARKER, G. J. M., WHEELER-KINGSHOTT, C. A. M., AND BARKER, G. J. Estimating distributed anatomical connectivity using fast marching methods and diffusion tensor imaging. *IEEE Transactions on Medical Imaging* 21 (2002), 505–512.
- [7] KOCH, M. A., NORRIS, D. G., AND HUND-GEORGIADIS, M. An investigation of functional and anatomical connectivity using magnetic resonance imaging. *NeuroImage* 16 (2002), 241–250.

- [8] LAZAR, M., AND ALEXANDER, A. L. White matter tractography using random vector (RAVE) perturbation. In *Proceedings of the ISMRM* (2002), vol. 10, p. 539.
- [9] PARKER, G. J. M., BARKER, G. J., THACKER, N. A., AND JACKSON, A. A framework for a streamline-based probabilistic index of connectivity (PiCo) using a structural interpretation of anisotropic diffusion. In *Proceedings of the ISMRM* (2002), vol. 10, p. 1165.
- [10] BEHRENS, T. E. J., JENKINSON, M., BRADY, J. M., AND SMITH, S. M. A probabilistic framework for estimation neural connectivity from diffusion weighted MRI. In *Proceedings of the ISMRM* (2002), vol. 10, p. 1142.
- [11] CORY, D. G., AND GARROWAY, A. N. Measurement of translational displacement probabilities by NMR: an indicator of compartmentation. *Magnetic Resonance in Medicine* 14 (1990), 435–444.
- [12] TUCH, D. S., REESE, T. G., WIEGELL, M. R., MAKRIS, N., BELLIVEAU, J. W., AND WEDEEN, V. J. High angular resolution diffusion imaging reveals intravoxel white matter fiber heterogeneity. *Magnetic Resonance in Medicine* 48 (2002), 577–582.
- [13] BASSER, P. J., AND PIERPAOLI, C. Microstructural and physiological features of tissues elucidated by quantitative-diffusion-tensor MRI. *Journal of Magnetic Resonance B* 111 (1996), 209–219.
- [14] JONES, D. K., HORSFIELD, M. A., AND SIMMONS, A. Optimal strategies for measuring diffusion in anisotropic systems by magnetic resonance imaging. *Magnetic Resonance in Medicine* 42 (1999), 515–525.
- [15] ALEXANDER, D. C., BARKER, G. J., AND ARRIDGE, S. R. Detection and modelling of non-Gaussian apparent diffusion coefficient profiles in human brain data. *Magnetic Resonance in Medicine* 48 (2002), 331–340.
- [16] FRANK, L. R. Characterisation of anisotropy in high angular resolution diffusion-weighted MRI. *Magnetic Resonance in Medicine* 47 (2002), 1083–1099.
- [17] TOURNIER, J.-D., CALAMANTE, F., KING, M. D., GADIAN, D. G., AND CONNELLY, A. Limitations and requirements of diffusion tensor fiber tracking: an assessment using simulations. *Magnetic Resonance in Medicine* 47 (2002), 701–708.
- [18] BEHRENS, T. E. J., WOOLRICH, M., JENKINSON, M., BRADY, J. M., AND SMITH, S. M. Bayesian parameter estimation in diffusion weighted MRI. In *Proceedings of the ISMRM* (2002), vol. 10, p. 1160.

IMAGING ABNORMALITIES ASSOCIATED WITH LANGUAGE DEFICITS IN PATIENTS WITH ACQUIRED UNILATERAL BASAL GANGLIA INFARCTION

CONTENTS

7.1	Introduction	143
7.2	Methods	146
7.2.1	Study population	146
7.2.2	Language assessment	146
7.2.3	Voxel-based morphometry	147
7.2.4	MR imaging and analysis	148
7.3	Results	151
7.3.1	Language performance	151
7.3.2	Conventional T_1 and T_2 weighted imaging	152
7.3.3	Voxel-based morphometry: structural imaging	152
7.3.4	Voxel-based morphometry: diffusion tensor imaging	156
7.4	Discussion	158
7.5	Conclusion	161
	References	161

7.1 INTRODUCTION

It has already been established that the basal ganglia¹ are important for intact motor control [1], and for the motor control of speech functions [2]. It is also known that most speech and language processing in the human brain usually takes place in the left hemisphere [3], in particular in the inferior frontal gyrus, also known as Broca's area, and in the posterior superior temporal cortex, also known as Wernicke's area. However, the relationship between aphasia² and basal ganglia damage is poorly understood. Several theories have been put forward to explain the inconsistencies observed between speech and language function and abnormalities found using various imaging techniques. These include the following possibilities.

- **The direct involvement of specific basal ganglia nuclei:** specific nuclei of the basal ganglia may play a direct role in speech and language functions. A number of studies suggest that the basal ganglia are involved in normal language processing either directly or indirectly through their connections with cortical language sites [4, 5, 6, 7]. The variability in severity and type of aphasia following subcortical damage may therefore be explained by the size or site of the lesion. However, several studies have shown an absence of a correlation between lesion size and severity [8, 9] or type of aphasic impairment [10] following basal ganglia lesions. On the other hand, other studies have shown an association between aphasia and damage to the anterior portion of the head of the caudate nucleus and to the anterior limb of the internal capsule [4, 5]. However, although the caudate head may be directly involved in language processing, the observed deficits may also be due to the disruption of connections between this structure and other language sites.
- **Damage to the white matter tracts in and around the lesion:** the basal ganglia are surrounded by white matter pathways that connect the basal ganglia to the thalamus, much of the cerebral cortex and the brain stem [11]. White matter tracts surrounding the basal ganglia also contain extensive projections between the thalamus, cerebral cortex and brainstem that do not involve connections with the basal ganglia. Lesions affecting the basal ganglia frequently include damage to these white matter tracts, particularly affecting the internal capsule and sometimes extending into the external capsule, extreme capsule and corona radiata. It may therefore be that speech and language difficulties associated with basal ganglia injury are not

¹The basal ganglia are a collection of five large nuclei located beneath the cerebral cortex in the forebrain, midbrain and diencephalon: the caudate nucleus, putamen, globus pallidus (consisting of internal and external segments), substantia nigra (pars compacta and pars reticulata), and subthalamic nucleus.

²Impairment or loss of the faculty of using or understanding spoken or written language

necessarily a direct result of injury to the basal ganglia nuclei, but of damage to white matter connections between the basal ganglia and the cortex or even to white matter tracts that do not involve the basal ganglia. A number of studies of basal ganglia infarction in adults have concluded that white matter tract damage is crucial in the manifestation of speech and language deficits [8, 12, 13] and that the type of linguistic disturbance exhibited is dependent on the specific region of white matter tract damage [8].

- **Reduced input from the basal ganglia to the cortex:** speech and language impairments observed in adult patients with left hemisphere basal ganglia infarctions have been linked to the presence of abnormalities beyond the basal ganglia and surrounding white matter. Such abnormalities might be caused by deafferentiation from subcortical inputs. The basal ganglia are part of a complex circuit involving large regions of the cerebral cortex and the brain stem. Virtually all areas of the cerebral cortex, including the frontal and temporal regions [14] involved in speech and language processing, project ipsilaterally to the basal ganglia and back to the same cortical regions via the thalamus. It might therefore be expected that damage to the basal ganglia nuclei would deprive the language cortices of neuronal input. This may lead to inactivity of the connected regions of the cortex and may even progress to Wallerian degeneration³ in the long-term. The cortical disconnection hypothesis also gains support from SPECT perfusion studies [16, 17], where the presence of aphasia in a group of patients with subcortical infarctions was associated with hypoperfusion in perisylvian regions including Broca's and Wernicke's areas. However, since the majority of those patients did not have occlusion of the internal carotid artery (ICA), the perfusion deficit was thought to be a result of reduced cortical activity caused by deafferentiation from the subcortical inputs.
- **Chronic hypoperfusion in cortical speech and language areas:** an alternative explanation for abnormalities remote from the core lesion site affecting speech and language following basal ganglia injury is that of cortical hypoperfusion in regions that appear structurally normal [18]. Since infarctions of the basal ganglia usually involve the ICA or middle cerebral artery (MCA), the entire MCA territory may be at risk. In the left hemisphere, this territory includes the perisylvian cortex, which is essential for language functions [18]. In a recent report, using highly sensitive diffusion and perfusion-weighted MRI techniques, Hillis *et al* [19] showed that aphasia or neglect in 44 patients with acute subcortical infarctions was consistently related (in 100% of cases) to cortical hypoperfusion in the MCA territory. More-

³the anterograde degeneration of axons and myelin sheaths following proximal neuronal injury [15]

over, intervention to induce reversal of the cortical hypoperfusion was associated with a resolution of aphasia in all 6 of the patients who underwent intervention.

The above theories may help in the current study to relate language impairments with imaging findings obtained from a group of children and adolescents with acquired unilateral basal ganglia infarction. These investigations were undertaken as part of a study performed in collaboration with the Developmental Cognitive Neuroscience Unit at ICH, in particular with Alison Rowan (PhD student). The respective contributions to this study have already been outlined in the introduction (page 13). The data collected included behavioural tests of speech and language performance, together with structural, diffusion and perfusion MR imaging. Since speech and language function are known to be mostly left-lateralised, significant differences in speech and language performance were expected between the group of patients with a left-sided infarct and that with a right-sided infarct. However, although both speech and language function in both patient groups were found to be impaired relative to age-matched controls, no significant differences were found between the two patient groups. On the other hand, a considerably greater variation in language performance was observed in the group of patients with a left-sided lesion: some patients with left-sided injuries demonstrated marked difficulties in language function, with three patients having scores lower than two standard deviations below the population mean, while others with apparently similar lesions showed no evidence for language difficulties. In contrast, all individuals with right-sided injuries performed within the normal range on tests of language function.

Until recently, many studies have been limited by the lack of imaging resolution or sensitivity to different properties of tissue. In many cases, this has prevented the relative importance of the different mechanisms involved in the manifestation of aphasia from becoming apparent. Conventional MR imaging was used to select the patients participating in this study. However, additional more subtle brain abnormalities that are not seen on conventional MRI might perhaps explain the variation in language performance mentioned above. Such abnormalities might be detectable using a number of more recently developed MRI image acquisition and analysis techniques. Such techniques were therefore employed in the present study, with the following aims:

- to characterise the site and extent of basal ganglia infarctions, including the specific basal ganglia nuclei involved
- to characterise the site and extent of any white matter tract damage surrounding and beyond the basal ganglia infarctions
- to identify the presence of cortical abnormalities that are not visible on clinical images, specifically involving the cortical speech and language regions

subject group	group size	mean age	age range
left-sided infarct	10	12:4	6:10 – 19:11
right-sided infarct	7	9:5	7:5 – 12:0
control	17	11:4	7:2 – 21:5

Table 7.1: description of the patients used in this study. All age values are written using the notation [years:months].

- to assess for the presence of perfusion abnormalities beyond the basal ganglia that have not resulted in structural abnormalities but that may produce functional impairments
- to identify relationships between arterial abnormalities, deficits in perfusion and functional abnormalities

This chapter will describe only a section of the whole study in order to demonstrate a practical application of diffusion tensor imaging, namely that related to aims 2 & 3 above. These data enabled the investigation of the correlation between language performance and focal changes in both morphology and diffusion anisotropy (identified using voxel-based morphometry) in a group of patients with acquired unilateral basal ganglia infarction.

7.2 METHODS

7.2.1 STUDY POPULATION

The patient group consisted of 17 children and adolescents with acquired unilateral basal ganglia infarction. This group was split up into a group of 10 patients with a left-sided lesion, and a group of 7 patients with a right-sided lesion. The age at assessment and time elapsed since infarction are shown in table 7.1. Tests were also performed on a group of age-matched controls (N = 17).

7.2.2 LANGUAGE ASSESSMENT

Language function was assessed using the clinical evaluation of language fundamentals (CELF-III) test [20]. The CELF-III is designed to assess performance in selected aspects of language, with its focus on language form and content rather than use in conversational or communicative contexts. It assesses both receptive and expressive language skills, with subtests evaluating word meaning (semantics), sentence structure (syntax), and recall and retrieval (auditory memory). In the normal population, receptive, expressive and total language scores have a mean of 100 and a standard deviation of 15.

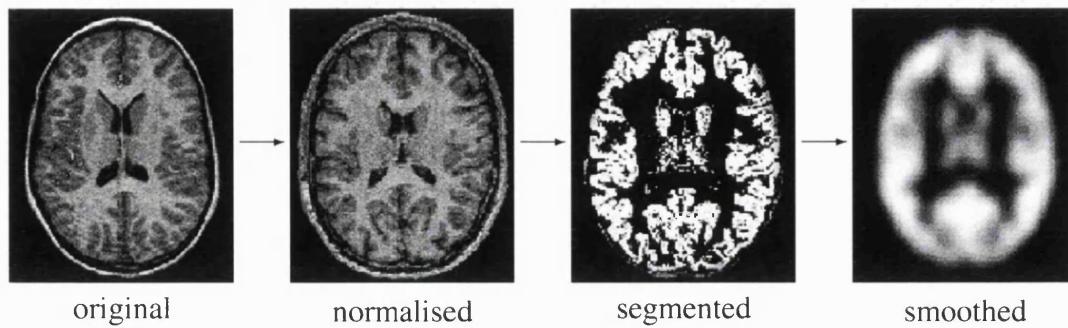


Figure 7.1: processing stages of voxel-based morphometry. The original 3D data are first normalised into the same stereotactic space. They are then segmented into grey matter, white matter, and CSF. Each of these maps is then smoothed to increase the validity of the parametric statistical test.

7.2.3 VOXEL-BASED MORPHOMETRY

Voxel-based morphometry (VBM) is a method of characterising subtle differences between groups of images that may not be evident on visual inspection. VBM provides quantitative information through statistical analyses that test hypotheses (with *t*-tests and *F*-tests) by making comparisons between brain images on a voxel-by-voxel basis. It is often used to identify focal differences between images acquired on a particular group of patients and images acquired on a different group (usually control subjects). It can also be used to find correlations within a group of patients between changes in the images and a measure of performance (such as IQ or language performance). It is not biased towards a given structure or regional difference and can provide a comprehensive analysis of the entire brain. VBM has the additional advantage of not being dependent on subjective judgements, since it uses an automatic procedure and standardised parametric statistics to provide an objective, quantitative measurement of grey and white matter changes [21, 22].

Before MR data can be statistically analysed, they must first undergo several stages of processing. Figure 7.1 illustrates the processing stages of VBM. The MR images are first normalised into the same stereotactic space. Following this, the normalised data are usually segmented into grey matter, white matter and CSF. This step generates probability maps for each voxel being either grey matter, white matter or CSF, depending on its spatial position and signal intensity in comparison to prior probability information derived from control subjects. These maps are then smoothed using an isotropic Gaussian kernel, to ensure that the data are more normally distributed, thereby increasing the validity of the parametric statistical analyses. Furthermore, using a smoothing kernel of a similar size to that of the expected differences sensitises the analysis to the spatial scale of the structure of interest. Statistical parametric maps are then created for the entire brain

by performing voxel-by-voxel statistical tests. For a more detailed explanation of these processing stages, see Ashburner and Friston (2000).

These statistical parametric maps are formed from a large number of statistical tests. In order to limit the generation of false positive results, it is necessary to correct for multiple comparisons [21]. The correction used here is based on Gaussian Random Field Theory. It can be applied over the whole brain, to yield so-called *corrected p*-values, or if a prior hypothesis exists, over a small volume of interest containing the structure thought to be relevant, to yield so-called *small volume corrected* (SVC) *p*-values. It is also acceptable to interpret significant results that are not corrected for multiple comparisons in cases where a prior hypothesis exists about specific voxels.

In this study, all VBM analyses were performed using Statistical Parametric Mapping software (SPM99, Wellcome Department of Imaging Neuroscience, London, UK).

7.2.4 MR IMAGING AND ANALYSIS

All MR images were collected using a 1.5T Siemens Vision System. The protocol consisted of conventional T_1 and T_2 weighted imaging, high resolution 3D structural imaging, and diffusion tensor imaging. Although they will not be reported here, perfusion imaging and MR angiography were also performed when possible. Note that in this chapter, all images are displayed in neurological convention (i.e. the left hemisphere is on the left, right is on the right), since this is the SPM convention.

CONVENTIONAL T_1 AND T_2 WEIGHTED IMAGING

T_1 weighted images were acquired using a sagittal multi-slice sequence (TR = 570 ms, TE = 14 ms, flip angle = 60°, FOV = 200 mm, matrix size = 256×512, slice thickness = 4 mm, 19 slices). Axial T_2 weighted images were acquired using a multi-slice turbo spin-echo (TSE) sequence (TR = 3458 ms, TE = 96 ms, 7 echoes, FOV = 220 mm, matrix size = 196×512, slice thickness = 5 mm, 19 slices). The T_1 and T_2 weighted images obtained from all patients and control subjects were visually inspected by a neuroradiologist (Dr. K. Chong). The presence of lesions, the number of lesions, and the specific structures affected were recorded. In addition, the presence of atrophy and other abnormalities were noted.

HIGH RESOLUTION 3D STRUCTURAL IMAGING

High resolution images were acquired using a T_1 weighted 3D fast low angle shot (FLASH) sequence [23] (TR = 16.8 ms, TE = 5.70 ms, flip angle = 21°, FOV = 200 mm, matrix size = 200×256, 160 partitions, voxel size = 1.00×0.78×0.78 mm). A VBM analysis was

then performed on these data. The MR data were first normalised by global grey matter to a T_1 weighted template. They were then segmented into grey matter, white matter and CSF, and the resulting grey matter and white matter maps were smoothed with a 12 mm and an 8 mm isotropic Gaussian kernel respectively since these dimensions correspond approximately to the dimensions of the cortical regions of interest and the cross-sectional dimensions of white matter. A series of correlation analyses was then carried out to examine the relationship between structural abnormalities on the 3D FLASH scans and the CELF-III measure of language, with particular focus on the following areas: the basal ganglia nuclei of the injured hemisphere, the cortical language areas and white matter tracts surrounding the basal ganglia and connecting the cortical language areas with the basal ganglia, in particular the internal capsule.

Inferences from statistical parametric maps were made at three different threshold levels. Firstly, correlations throughout the entire brain were reported if they reached the significance value $p = 0.05$, corrected for multiple comparisons across the entire brain. Secondly, correlations involving the regions predicted to be affected were reported if they reached significance with the application of small volume corrections to correct for multiple comparisons within the regions of predicted differences (i.e. $p = 0.05$, SVC). The principal regions of predicted difference were the basal ganglia and surrounding white matter, and the left hemisphere inferior frontal and superior temporal regions (corresponding to the cortical language areas). The regions used for the SVC are shown in figure 7.2. These regions were highly conservative since they encompassed regions of greater volume than the structures of interest. Finally, correlations involving the regions predicted to be affected were reported if they reached a conservative level of significance without correction for multiple comparisons (i.e. $p = 0.001$, uncorrected).

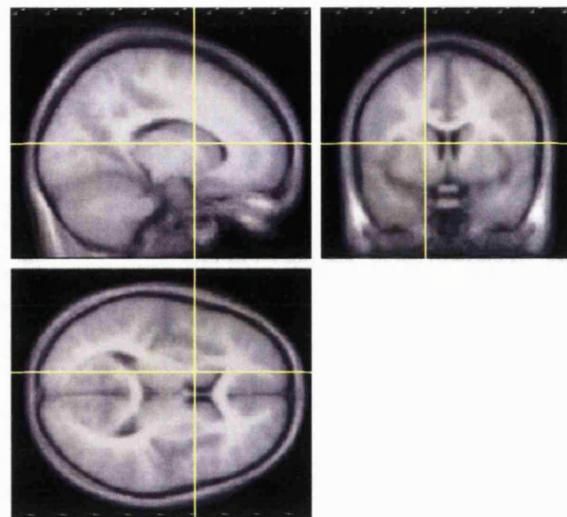
The statistical parametric maps are shown superimposed onto the mean normalised image of the group data in order to aid anatomical localisation. Identification of the anatomical location of regions of abnormality was carried out in reference to Durvernoy's atlas [24]. The maps are shown at an uncorrected threshold $p = 0.005$, and cross hairs indicate the location of the maximal peak.

DIFFUSION TENSOR IMAGING

Diffusion tensor data sets were acquired using a twice-refocused diffusion-weighted spin-echo EPI sequence (TE = 110 ms, matrix size = 128×128 , zero-filled to 256×256 , pixel size = 1.5×1.5 mm after zero-filling, slice thickness = 3 mm, 40 contiguous slices covering the whole brain). For each slice location, a total of 23 images were acquired, including 3 $b=0$ images, and the remainder with the diffusion-encoding gradients (b -value = 1000 s/mm 2) applied along 20 non-collinear directions uniformly distributed over

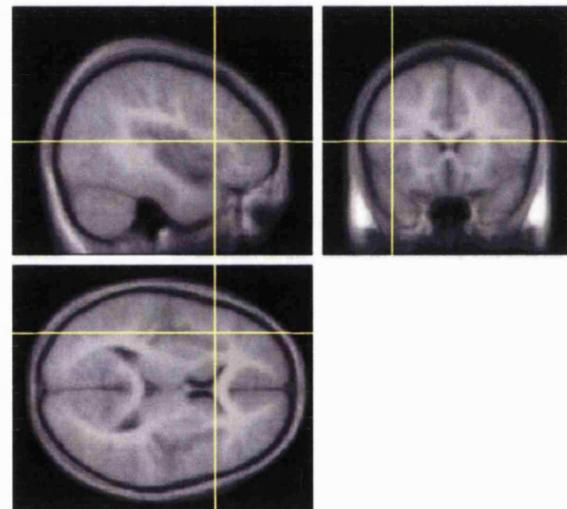
SVC around basal ganglia and surrounding white matter:

30x42x30 mm box,
centre: [±15 6 10]



SVC around left inferior frontal region (Broca's area):

sphere, 20 mm radius,
centre: [-40 20 10]



SVC around left superior temporal region (Wernicke's area):

sphere, 20 mm radius,
centre: [-60 -35 20]

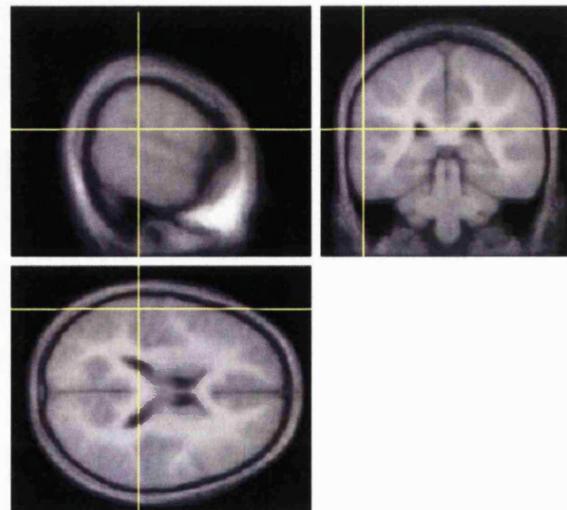


Figure 7.2: regions where a small volume correction (SVC) analysis was performed in this study. The centre of each SVC region is indicated by cross-hairs overlaid on axial, sagittal and coronal group-averaged structural images.

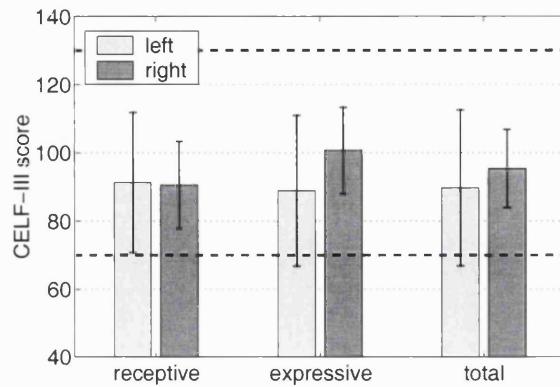


Figure 7.3: average receptive, expressive, and total language scores for both left and right-sided lesion groups. The thick dashed lines represent the normal range (mean \pm 2 SD).

the surface of a sphere. This sequence is described in more detail in section 4.2.

The raw diffusion-weighted images underwent a number of pre-processing stages using in-house software. First, all slices were visually inspected to identify those corrupted by pulsation or movement. These were not taken into account in the subsequent calculation of the tensor images. A subject's entire data set was rejected if more than five diffusion-weighted images from any one slice needed to be rejected. Maps of the fractional anisotropy (FA) [25] were then calculated from the resulting diffusion tensor maps.

The FA maps thus generated were analysed using VBM. In order to normalise the FA maps, the corresponding $b=0$ (T_2 weighted) images were normalised to a standard EPI template, and the same transformation was applied to the corresponding FA data. In contrast to the analysis performed for the structural data, no segmentation was performed for the normalised FA maps, since the intensity in these images is mainly from white matter. The normalised FA data were then smoothed and analysed in the same way as described above for the structural data.

7.3 RESULTS

7.3.1 LANGUAGE PERFORMANCE

The group average CELF-III scores of the left and right-injured groups are shown on figure 7.3. No significant differences in the mean performance were observed between the left-sided and right-sided lesion groups for all three measures of language. Individual scores for the left-sided and right-sided lesion groups are shown in figure 7.4. Language performance for all patients in the right-sided lesion group is within the normal range. However, the variance in language performance in the left-sided lesion group is much

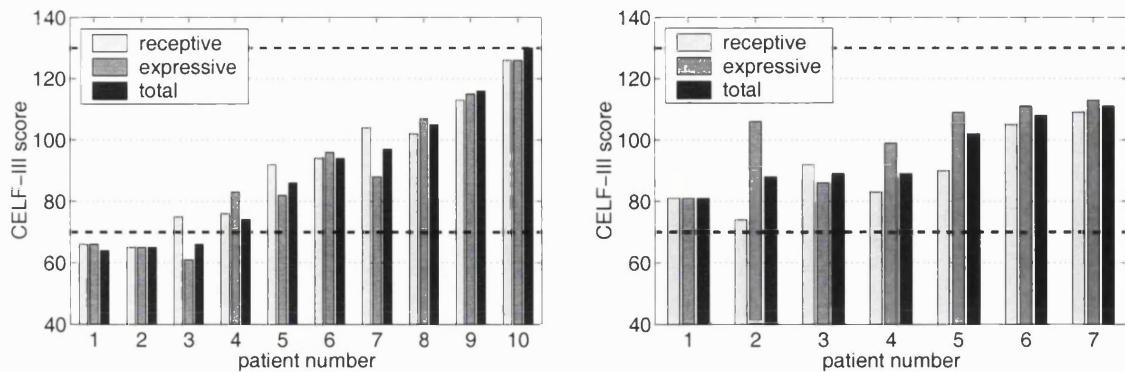


Figure 7.4: receptive, expressive and total language scores for individual patients in the left-sided (left) and right-sided (right) lesion groups. The thick dashed lines represent the normal range (mean \pm 2 SD).

greater, with three patients having scores lower than two standard deviations below the population mean.

No significant differences were observed between receptive and expressive language scores measured using this test. The total language score was therefore considered to be representative of overall language function, and was used as a correlate in subsequent analyses.

7.3.2 CONVENTIONAL T_1 AND T_2 WEIGHTED IMAGING

The T2-weighted clinical MR images are shown in figures 7.5 and 7.6. No common site of damage was observed in the basal ganglia in relation to language performance. In particular, the head of the left hemisphere caudate nucleus was not damaged in all patients with language impairments, but was damaged in some patients without language difficulties.

7.3.3 VOXEL-BASED MORPHOMETRY: STRUCTURAL IMAGING

LEFT-SIDED LESION GROUP

Figures 7.7 to 7.11 and table 7.2 show the results of the correlation between language function (as measured by the CELF-III) and changes on structural scans in patients with a left-sided injury. In this group, grey and white matter density in and around the basal ganglia were not found to be significantly correlated with language function. However, a significant correlation between language function and grey matter density was found in the left cortical language areas, including the inferior frontal gyrus, middle frontal gyrus, insular cortex and superior temporal gyrus (all smoothed to 12 mm). A significant correlation between white matter density and language scores was also observed in left-

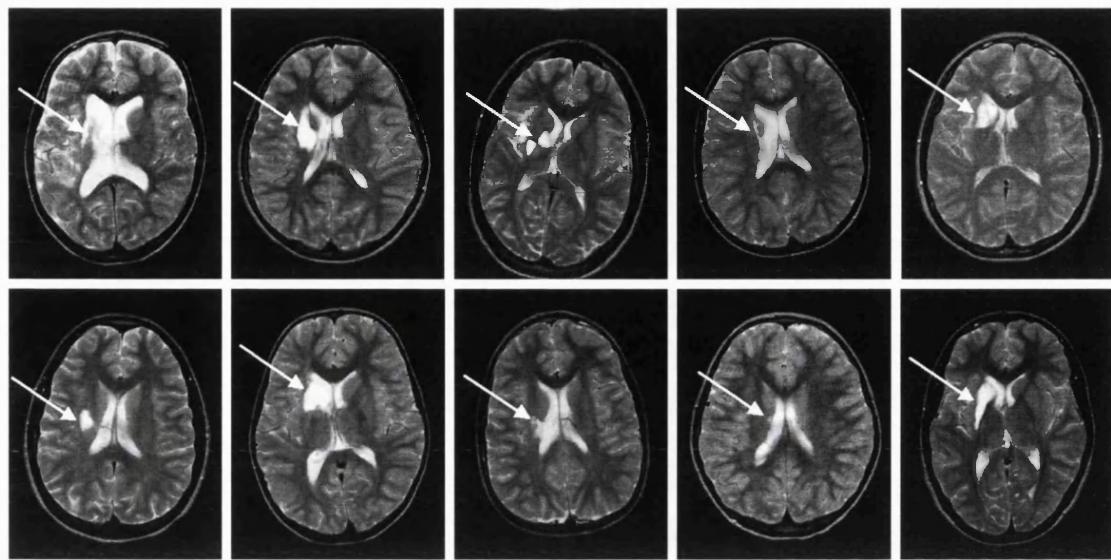


Figure 7.5: T_2 weighted images showing the location of the lesion for each patient in the left-sided lesion group. The arrows indicate the site of the lesion. Top, left to right: patients 1 to 5 (numbered as in figure 7.4). Bottom, left to right: patients 6 to 10.

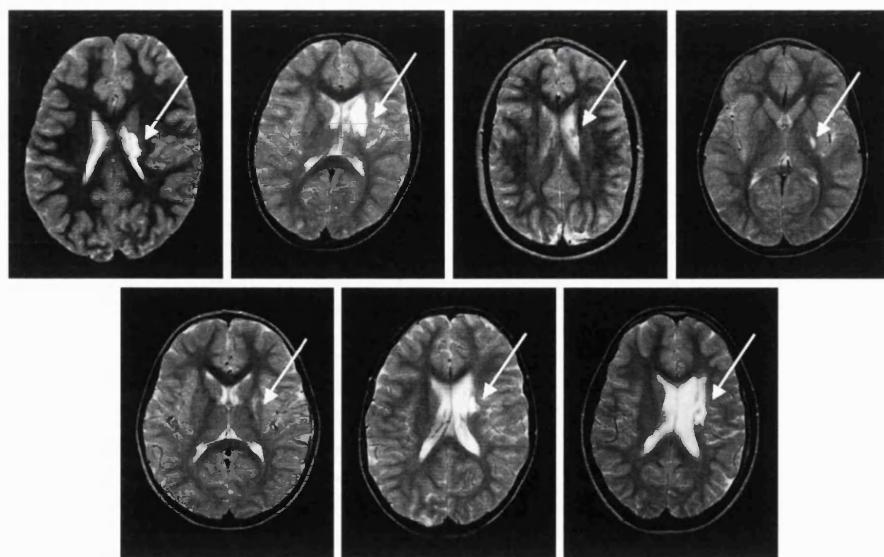


Figure 7.6: T_2 weighted images showing the location of the lesion for each patient in the right-sided lesion group. The arrows indicate the site of the lesion. Top, left to right: patients 1 to 4 (numbered as in figure 7.4). Bottom, left to right: patients 5 to 7.

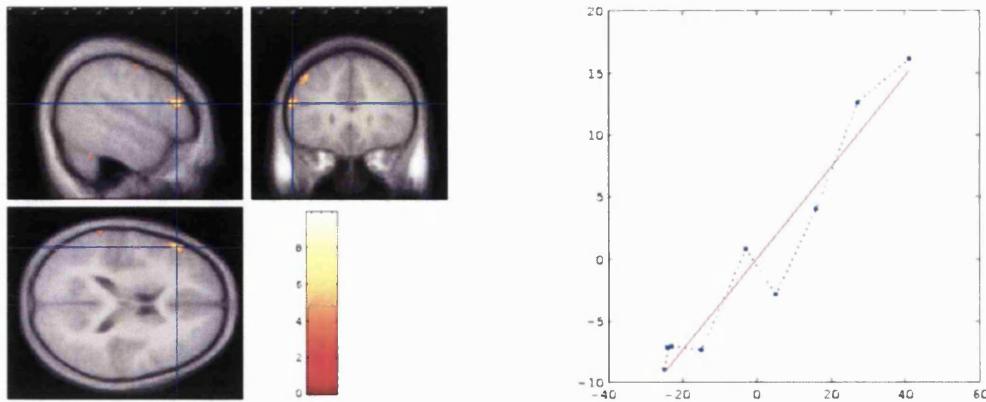


Figure 7.7: correlation between CELF-III language scores and grey matter density in the left inferior frontal gyrus (pars triangularis), identified by VBM in the group of patients with a left-sided lesion. Right: The centre of the significant region is indicated by cross-hairs overlaid on axial, sagittal and coronal group-averaged structural images. Left: CELF-III total language score (x -axis) plotted against grey matter density (y -axis). The red line represents the best fit through the data.

hemisphere deep frontal white matter (smoothed to 8 mm). No other predicted regions of grey or white matter density were found to be significantly correlated with language function.

RIGHT-SIDED LESION GROUP

In the right-injured group, the same analysis did not highlight any cortical language area where grey matter density was significantly correlated with language score. However, a significant correlation was observed between white matter density and language score in deep frontal white matter, as shown in table 7.2.

location	position	Z-score	corrected	p-value	smoothing
left inferior frontal gyrus	[-52, 30, 18]	4.23	SVC	0.05	12 mm
left middle frontal gyrus	[-40, 26, 45]	3.80	no	<0.001	12 mm
left insular cortex	[-42, -4, -3]	3.40	no	<0.001	12 mm
left superior temporal gyrus	[-69, -40, 15]	3.04	no	0.001	12 mm
left deep frontal white matter	[-34, 27, 16]	3.83	no	<0.001	8 mm
right deep frontal white matter	[38, 32, 8]	3.55	no	<0.001	8 mm

Table 7.2: results of correlation analyses between language function (CELF-III) and structural damage on 3D-FLASH images in patients with left-sided (first 5 results) and right-sided (last result) infarctions

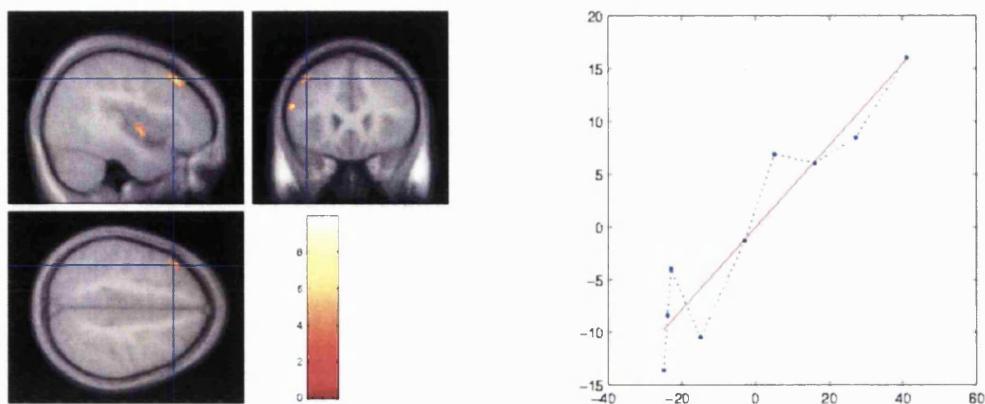


Figure 7.8: correlation between CELF-III language scores and grey matter density in the left middle frontal gyrus, identified by VBM in the group of patients with a left-sided lesion (cf. figure 7.7).

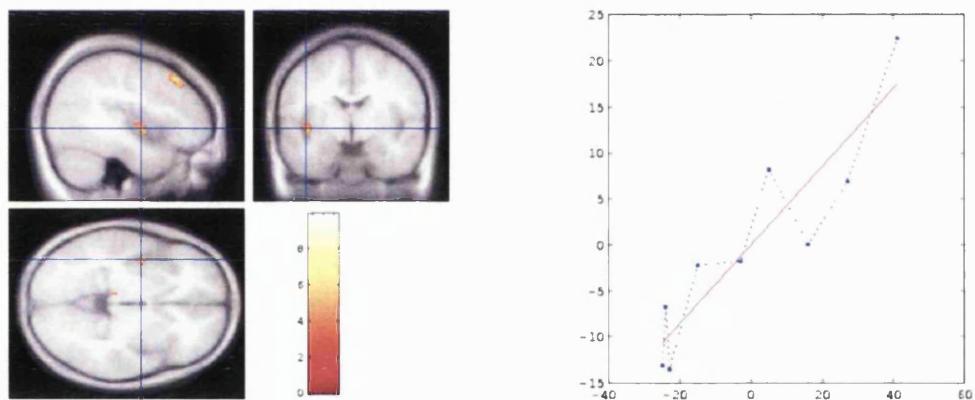


Figure 7.9: correlation between CELF-III language scores and grey matter density in the left insular cortex, identified by VBM in the group of patients with a left-sided lesion (cf. figure 7.7).

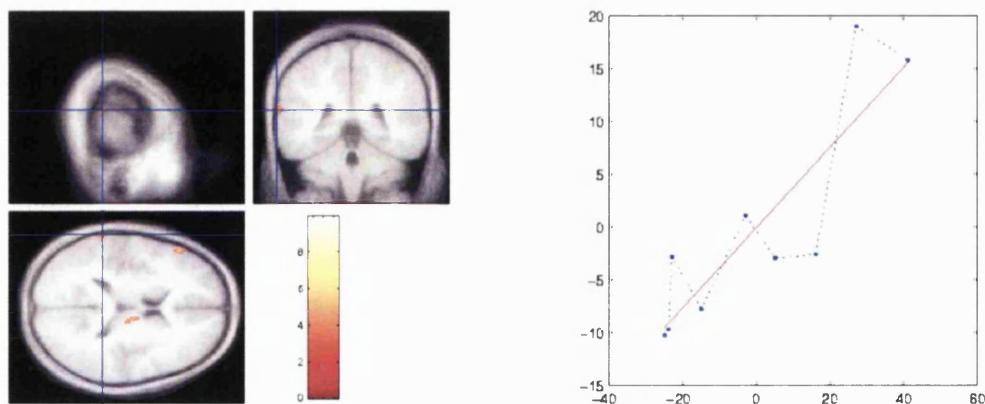


Figure 7.10: correlation between CELF-III language scores and grey matter density in the left superior temporal gyrus, identified by VBM in the group of patients with a left-sided lesion (cf. figure 7.7).

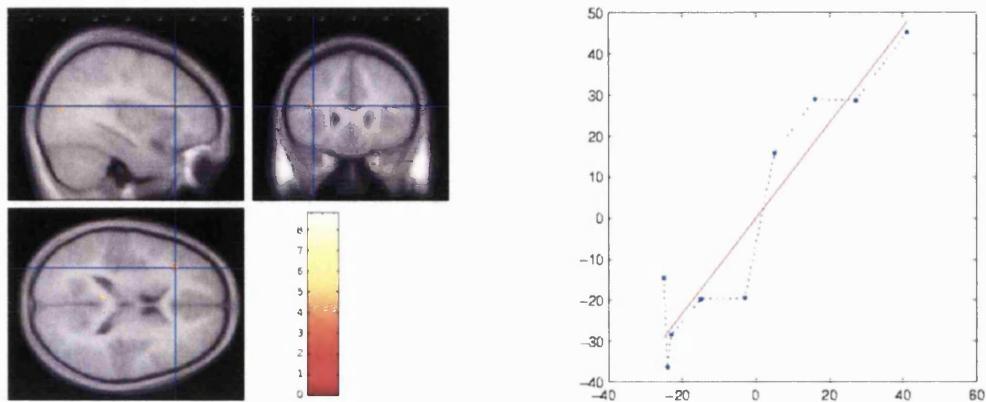


Figure 7.11: correlation between CELF-III language scores and white matter density in the left deep frontal white matter, identified by VBM in the group of patients with a left-sided lesion (cf. figure 7.7).

TOTAL CONTROL GROUP

In the control group, no region was found where either grey or white matter density were significantly correlated with language score.

7.3.4 VOXEL-BASED MORPHOMETRY: DIFFUSION TENSOR IMAGING

LEFT-SIDED LESION GROUP

Figures 7.12 to 7.14 and table 7.3 show the results of the correlation between language function (as measured by the CELF-III) and changes on FA maps in patients with left-hemispheric injuries. Anisotropy in the left internal capsule and deep frontal white matter (all smoothed to 8 mm) was found to be significantly correlated with language function in patients with left-hemispheric infarctions. No other predicted regions of white matter anisotropy were found to be significantly correlated with language function.

RIGHT-SIDED LESION GROUP

In the right-sided lesion group, no significant correlation was found between fractional anisotropy and language score.

location	position	Z-score	corrected	p-value	smoothing
left gyrus rectus	[-7, 1, -15]	4.32	no	<0.001	8 mm
left internal capsule, anterior limb	[-14, 4, 15]	3.50	no	<0.001	8 mm
left deep frontal white matter	[-27, 15, 18]	4.46	no	<0.001	8 mm

Table 7.3: results of correlation analyses between language function (CELF-III) and fractional anisotropy in patients with left-sided infarctions

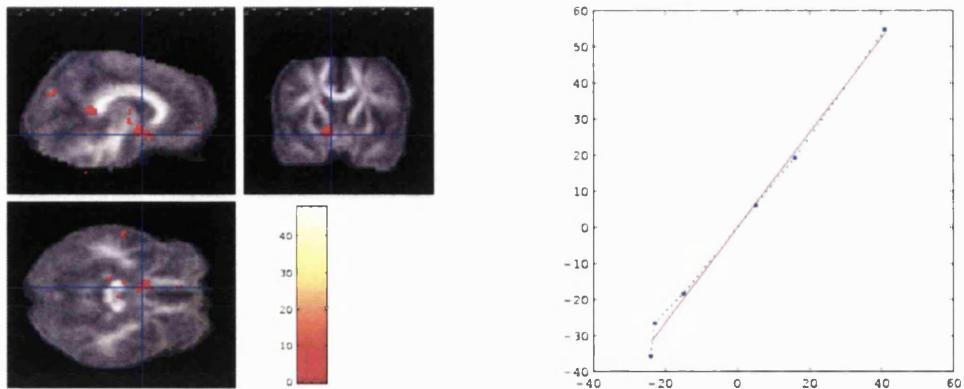


Figure 7.12: correlation between CELF-III language scores and diffusion anisotropy in the left gyrus rectus, identified by VBM in the group of patients with a left-sided lesion. Right: The centre of the significant region is indicated by cross-hairs overlaid on axial, sagittal and coronal group-averaged fractional anisotropy maps. Left: CELF-III total language score (x -axis) plotted against fractional anisotropy (y -axis). The red line represents the best fit through the data.

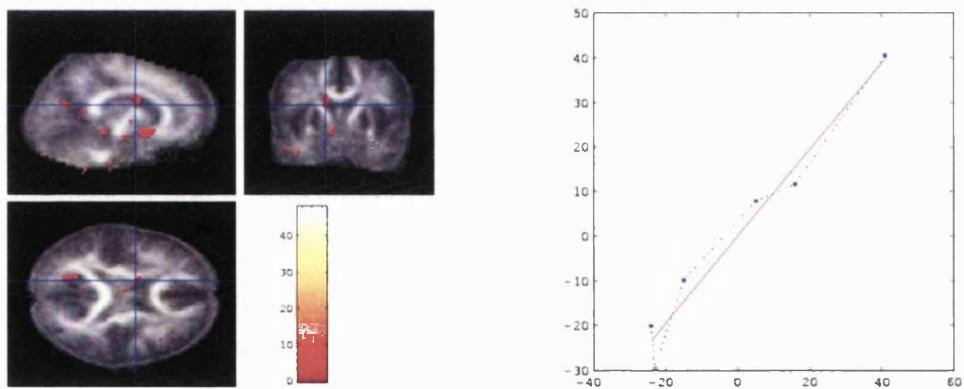


Figure 7.13: correlation between CELF-III language scores and diffusion anisotropy in the anterior limb of the left internal capsule, identified by VBM in the group of patients with a left-sided lesion (cf. figure 7.12).

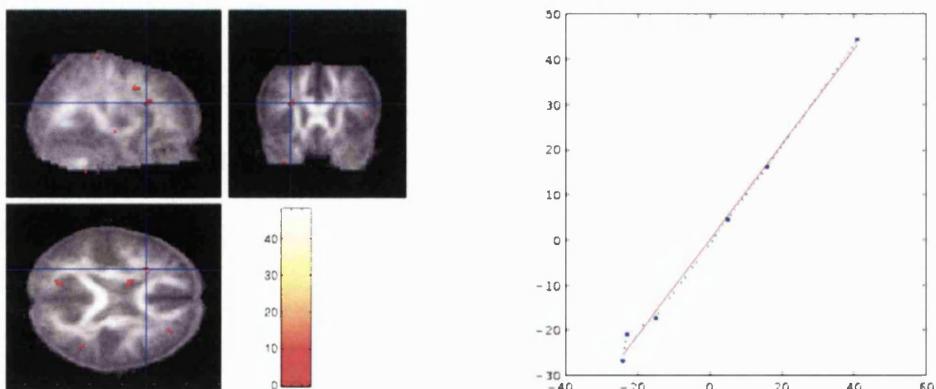


Figure 7.14: correlation between CELF-III language scores and diffusion anisotropy in the left deep frontal white matter, identified by VBM in the group of patients with a left-sided lesion (cf. figure 7.12).

TOTAL CONTROL GROUP

In the control group, no region was found where the fractional anisotropy was significantly correlated with language score.

7.4 DISCUSSION

This study highlights some of the limitations of relying on conventional imaging to investigate structure-function relationships. Abnormalities may exist that are not apparent using conventional imaging techniques. Such a lack of evidence may be misleading and could result in the wrong structure/function conclusions being drawn, as pointed out by Incisa della Rocchetta *et al* [26]. There is considerable controversy surrounding the role of the basal ganglia in speech and language function. Using advanced MR image acquisition and analysis techniques, the current study has identified further abnormalities remote from the site of the lesion that were not apparent using conventional imaging.

Visual inspection of the conventional T_1 and T_2 weighted imaging provided no evidence for a relationship between language performance and lesion site or size. Damage to specific basal ganglia nuclei did not explain the wide variance in language performance observed in the left-sided lesion group. In particular, the head of the left caudate nucleus was not damaged in all patients with language impairments and was damaged in some patients without language difficulties, contrary to the suggestion that damage to the head of the caudate nucleus is crucial in the mediation of language deficits in such patients [4, 5]. These results are consistent with a number of studies demonstrating an absence of a clear relationship between the type of speech and language deficit and the involvement of specific basal ganglia nuclei [7, 18, 27].

The VBM analysis of the high resolution structural imaging provided evidence for a positive correlation in the left-injured group between language performance and grey matter density in a number of conventional language areas, including Broca's area and Wernicke's area. No such correlation was observed in the other groups, suggesting that this structure-function relationship does not extend to the normal population.

The VBM analyses of both the structural and diffusion tensor data independently identified a positive correlation in the left-injured group between language performance and white matter density (for the structural data) and fractional anisotropy (for the diffusion tensor data) in a similar location in the left deep frontal white matter. The proximity of this region to the inferior frontal gyrus suggests that the white matter pathway connecting the basal ganglia to the inferior frontal gyrus may be involved. However, this region contains many white matter pathways and it is not possible from this study to determine which is

affected. This issue may be resolved using tractography techniques (see section 3.3.3), although this would require the acquisition of higher quality diffusion tensor data, leading to much increased scan times. The structural VBM analysis also identified a similar correlation in the right-injured group in a homologous area in the right deep frontal white matter. Again, no such correlation was found in the control group, suggesting that it is not due to a normal structure-function relationship.

In the left-sided lesion group, the VBM analysis of the diffusion tensor data identified a positive correlation between language score and fractional anisotropy in the anterior limb of the left internal capsule. Since this region is close to the lesion site, this relationship might be due to the more extensive lesions resulting in poorer functional performance. On the other hand, the anterior limb of the internal capsule is known to contain fibre projections from the basal ganglia and thalamus to all parts of the frontal lobes [11].

In summary, abnormalities remote from the core site of the lesion were identified in key language areas and in white matter regions likely to contain connections between some of these areas and the basal ganglia. These results are consistent with the theory that language deficits following basal ganglia infarctions are related to white matter tract damage [8, 13], and/or cortical deafferentiation from basal ganglia inputs [16].

The findings reported here should at this stage be regarded with caution since most regions do not reach significance when corrected for multiple comparisons. The correlation identified in the left-injured group in the left inferior frontal gyrus (Broca's area) between language performance and grey matter density did however reach significance with a small volume correction, and can be regarded with more confidence. Also, the analyses of both the structural and diffusion tensor data independently identified a region of significance in the left deep frontal white matter. This convergence of evidence suggests that the findings in this region are less likely to represent false positive results.

Diffusion anisotropy values must be interpreted with care, since they have been shown to be dependent on the degree of myelination [28], which increases most rapidly during the first two years of life, but continues to rise into early adulthood [29]. However, the control subjects were age-matched, effectively removing this confounding factor. In addition, since language performance was not found to be significantly correlated to age in these patients, the correlation found between the imaging data and language performance cannot be attributed to age. Partial volume effects also complicate the interpretation of anisotropy data. As described in section 3.3.4, it is possible for several highly anisotropic fibre populations to be present within the same voxel, arranged such that the overall anisotropy is reduced. However, this should not influence the significance of the results, since these effects are not expected to be group specific.

The normalisation of fractional anisotropy maps into a common stereotactic space for

use in the VBM analyses is not straightforward, since there are a number of ways to perform this operation. First, the anisotropy maps themselves could be directly normalised. However, this implies the existence of a template image with identical contrast properties, with which each patient's image can be compared, and such a template is difficult to generate⁴. A second approach would be to normalise the diffusion tensor maps prior to calculating the corresponding anisotropy maps. However, the normalisation process involves linear and non-linear transformations that may adversely affect the diffusion tensor. Although it has been shown that such transformations can be done robustly [30], performing the normalisation requires considerably more processing [31]. Moreover, this approach would also require the creation of a template. The approach that was used in this study was to normalise each subject's T_2 weighted $b=0$ image to the standard EPI template supplied with SPM99, and apply the same transformation parameters to the anisotropy maps. This is possible since the anisotropy map is calculated voxel-by-voxel from the base images, such that a transformation that is valid for one image should also be valid for the other.

It is interesting to note that the correlations observed between language scores and fractional anisotropy exhibit better linearity than those observed using the structural data. However, to understand this relationship, further work is needed. As previously mentioned, changes in diffusion anisotropy are difficult to interpret, but the diffusion tensor contains more information that may prove to be useful. There may for instance be a relationship between the trace of the diffusion tensor and language performance, that would be revealed by performing additional VBM analyses. Considerable insight may also be gained by looking at the properties of the diffusion tensor in the areas highlighted by this study. For example, there may be a relationship between language scores and the direction of the major eigenvector in the deep frontal white matter region: if that region contains several white matter pathways, it is possible that the degeneration of one pathway (in this case, that connecting Broca's area to the basal ganglia) allows the orientation of the others to dominate, as has been shown in the pons by Pierpaoli *et al* [32]. This effect may also be evident on inspection of the profile of the apparent diffusion coefficient as a function of orientation: it may be possible to observe a reduction in the contribution from the pathway that is thought to be involved [33].

⁴Generating a template involves averaging several normal brain images normalised to the same standard space. However, the problem of normalising these images still remains.

7.5 CONCLUSION

The results presented here provide evidence for abnormalities remote from the core site of the lesion that are not detectable using conventional MR image acquisition and analysis techniques in patients with acquired unilateral basal ganglia infarction. The findings are consistent with the theory that language deficits in this group of patients are associated with the loss of grey matter in left hemisphere cortical language areas and the interruption of the white matter tracts connected to these areas. This study highlights the importance of considering abnormalities beyond those observed on conventional MR imaging when investigating structure-function relationships. Advanced MR imaging techniques such as diffusion tensor imaging can provide additional valuable information that may influence the interpretation of the observed findings, and should be used whenever possible.

REFERENCES

- [1] MINK, J. W. *Basal ganglia*. In Zigmond et al. [34], 1999, pp. 951–972.
- [2] ALEXANDER, M. P. *Clinico-anatomical correlations of aphasia following predominantly subcortical lesions*. In Boller and Grafman [35], 1989, pp. 47–66.
- [3] PRICE, C. J. The anatomy of language: contributions from functional neuroimaging. *Journal of Anatomy* 197 (2000).
- [4] DAMASIO, A. R., DAMASIO, H., RIZZO, M., VARNEY, N., AND GERSH, F. Aphasia with nonhemorrhagic lesions in the basal ganglia and internal capsule. *Archives of Neurology* 39 (1982), 15–20.
- [5] GURD, J. M., AND BAMFORD, J. M. Striato-capsular aphasia: Contrasting cases. *Journal of Neurolinguistics* 10 (1997), 325–346.
- [6] KERTESZ, A., HARLOCK, W., AND COATES, R. Computer tomographic localisation, lesion size, and prognosis in aphasia and nonverbal impairment. *Brain and Language* 8 (1979), 34–50.
- [7] WEILLER, C., WILLMES, K., REICHE, W., THRON, A., ISENSEE, C., BUELL, U., AND RINGELSTEIN, E. B. The case of aphasia or neglect after striatocapsular infarction. *Brain* 116 (1993), 1509–1525.
- [8] NAESER, M. A., ALEXANDER, M. P., HELM, E. N., LEVINE, H. L., LAUGHLIN, S. A., AND GESCHWIND, N. Aphasia with predominantly subcortical lesion sites: Description of three capsular/putamal aphasia syndromes. *Archives of Neurology* 39 (1982), 2–14.
- [9] DEMEURISSE, G., CAPON, A., AND VERHAS, M. Prognostic value of computed tomography in aphasic stroke patients. *European Neurology* 24 (1985), 134–139.
- [10] WILLMES, K., AND POECK, K. To what extent can aphasic syndromes be localised? *Brain* 116 (1993), 1527–1540.
- [11] CARPENTER, M. *Core text of neuroanatomy*. Williams and Wilkins, London, 1991.
- [12] LIEBERMAN, R. R., ELLENBERG, M., AND RESTUM, W. H. Aphasia associated with verified subcortical lesions: three case reports. *Archives of Physical Medicine and Rehabilitation* 67 (1986), 410–414.
- [13] ALEXANDER, M. P., NAESER, M. A., AND PALUMBO, C. L. Correlations of subcortical CT lesion sites and aphasia profiles. *Brain* 110 (1987), 961–991.

[14] ROLLS, E. T., AND JOHNSTONE, S. *Neurophysiological analysis of striatal function*. In Vallar et al. [36], 1992, pp. 61–97.

[15] WERRING, D. J., TOOSY, A. T., CLARK, C. A., PARKER, G. J. M., BARKER, G. J., MILLER, D. H., AND THOMPSON, A. J. Diffusion tensor imaging can detect and quantify corticospinal tract degeneration after stroke. *Journal of Neurology, Neurosurgery, and Psychiatry* 69 (2000), 269–272.

[16] PERANI, D., VALLAR, G., CAPPA, S., MESSA, C., AND FAZIO, F. Aphasia and neglect after subcortical stroke. a clinical/cerebral perfusion correlation study. *Brain* 110 (1987), 1211–1229.

[17] OKUDA, B., TANAKA, H., TACHIBANA, H., KAWABATA, K., AND SUGITA, M. Cerebral blood flow in subcortical aphasia: perisylvian cortical hypoperfusion as a crucial role. *Stroke* 25 (1994), 1495–1499.

[18] NADEAU, S. E., AND CROSSON, B. Subcortical aphasia: Response to reviews. *Brain and Language* 58 (1997), 436–458.

[19] HILLIS, A. E., WITYK, R. J., BARKER, P. B., BEAUCHAMP, N. J., GAILLOUD, P., MURPHY, K., COOPER, O., AND METTER, E. J. Subcortical aphasia and neglect in acute stroke: the role of cortical hypoperfusion. *Brain* 125 (2002), 1094–1104.

[20] SEMEL, E., WIIG, E. H., AND SECORD, W. A. *Clinical evaluation of language fundamentals*, third ed. The Psychological Corporation, London, 1995.

[21] ASHBURNER, J., AND FRISTON, K. J. Voxel-based morphometry – the methods. *NeuroImage* 11 (2000), 805–821.

[22] GITELMAN, D. R., ASHBURNER, J., FRISTON, K. J., TYLER, L. K., AND PRICE, C. J. Voxel-based morphometry of herpes simplex encephalitis. *NeuroImage* 13 (2001), 623–631.

[23] HAASE, A., FRAHM, J., MATTHAEI, D., HÄNICKE, W., AND MERBOLDT, K.-D. FLASH imaging. rapid NMR imaging using low flip-angle pulses. *Journal of Magnetic Resonance* 67 (1986), 258–266.

[24] DUVERNOY, H. M. *The human brain: surface, three dimensional sectional anatomy and MRI*. Springer-Verlag, New York, 1991.

[25] BASSER, P. J., AND PIERPAOLI, C. Microstructural and physiological features of tissues elucidated by quantitative-diffusion-tensor MRI. *Journal of Magnetic Resonance B* 111 (1996), 209–219.

[26] INCISA DELLA ROCCHETTA, A., GADIAN, D. G., CONNELLY, A., POLKEY, C. E., JACKSON, G. D., WATKINS, K. E., JOHNSON, C. L., MISHKIN, M., AND VARGHA-KHADEM, F. Verbal memory impairment after right temporal lobe surgery: role of contralateral damage as revealed by ¹H magnetic resonance spectroscopy and T₂ relaxometry. *Neurology* 45 (1995), 797–802.

[27] KERTESZ, A. *Subcortical agraphia*. In Vallar et al. [36], 1992, pp. 334–356.

[28] BEAULIEU, C. The basis of anisotropic water diffusion in the nervous system – a technical review. *NMR in Biomedicine* 15 (2002), 435–455.

[29] DOBBING, J., AND SANDS, J. Quantitative growth and development of human brain. *Archives of Disease in Childhood* 49 (1973), 757–767.

[30] ALEXANDER, D. C., PIERPAOLI, C., BASSER, P. J., AND GEE, J. C. Spatial transformations of diffusion tensor magnetic resonance images. *IEEE Transactions on Medical Imaging* 20 (2001), 1131–1139.

[31] JONES, D. K., GRIFFIN, L. D., ALEXANDER, D. C., CATANI, M., HORSFIELD, M. A., HOWARD, R., AND WILLIAMS, S. C. R. Spatial normalisation and averaging of diffusion tensor MRI data sets. *NeuroImage* 17 (2002), 592–617.

[32] PIERPAOLI, C., BARNETT, A., PAJEVIC, S., CHEN, R., PENIX, L., VIRTA, A., AND BASSER, P. Water diffusion changes in Wallerian degeneration and their dependence on white matter architecture. *NeuroImage* 13 (2001), 1174–1185.

[33] FRANK, L. R. Characterisation of anisotropy in high angular resolution diffusion-weighted MRI. *Magnetic Resonance in Medicine* 47 (2002), 1083–1099.

- [34] ZIGMOND, M. J., BLOOM, F. E., LANDIS, S. C., ROBERTS, J. L., AND SQUIRE, L. R., Eds. *Fundamental neuroscience*. Academic Press, San Diego, 1999.
- [35] BOLLER, F., AND GRAFMAN, J., Eds. *Handbook of neuropsychology*. Elsevier, Amsterdam, 1989.
- [36] VALLAR, G., CAPPA, S. F., AND WALLESCHE, C. W., Eds. *Neuropsychological disorders associated with subcortical lesions*. Oxford University Press, Oxford, 1992.

CHAPTER 8

GENERAL DISCUSSION

CONTENTS

8.1 Diffusion anisotropy	165
8.2 Diffusion tensor tractography	166
8.3 Remaining issues	168
References	169

Diffusion tensor imaging is emerging as a powerful tool to probe tissue microstructure *in vivo*, and provides quantitative information that cannot be obtained using conventional methods. Maps of diffusion anisotropy have been used successfully in a number of patient studies, and practical applications of tractography are beginning to be identified. However, diffusion tensor data acquisition sequences and analysis techniques are technically challenging, and a number of issues need to be taken into consideration. Some of these issues have been the subject of this thesis.

Optimal diffusion tensor imaging requires the acquisition of a large number of diffusion-weighted images, with the diffusion-encoding gradients applied along a large number of directions. The arrangement of these directions needs to be chosen with care, since it will have an impact on the quality of the data (see section 3.3.2). In addition, diffusion-weighted images are very sensitive to motion. For this reason, rapid single-shot imaging techniques such as EPI are frequently used. As detailed in section 2.3, EPI images are prone to local susceptibility-induced distortions and signal dropout, as well as global eddy-current-induced distortions. Furthermore, the large gradients needed to induce a suitable diffusion weighting introduce further problems, described in section 3.2.3, notably increased eddy-current distortions and a significant reduction in the signal-to-noise ratio (SNR). The eddy-current-induced distortions will in general be different depending

on the direction of application of the diffusion-encoding gradients, such that the different diffusion-weighted images may not align with each other correctly. Finally, the combination of low SNR with the need to acquire a large number of images means that a compromise must be made between total imaging time, image resolution, and SNR. This is most acute in the clinical setting, where the scan time must be kept as short as possible. It is therefore important to rigorously assess the diffusion tensor acquisition sequence that is to be used to ensure that the data produced are of sufficient quality. The optimal compromise between scan time, image resolution and SNR will depend on the application for which the data is intended. There are two main applications for diffusion tensor data (apart from trace mapping): diffusion anisotropy mapping and tractography. Their particular requirements are discussed below.

8.1 DIFFUSION ANISOTROPY

Diffusion tensor data are often acquired for visual inspection or to perform region of interest analyses on anisotropy maps. These types of analyses have been performed to investigate acute [1] or chronic stroke [2, 3], multiple sclerosis [4, 5], epilepsy [6], schizophrenia [7], and CADASIL [8], amongst others. In this case relatively few slices are needed, the slice thickness can be made relatively large to increase the SNR, and the scan time can be kept down to a few minutes whilst still producing images of acceptable quality. A more rigorous analysis of diffusion anisotropy maps can be performed using statistical techniques such as voxel-based morphometry (VBM). Such studies have previously been performed to investigate patients with epilepsy [9], schizophrenia [10], or poor reading ability [11]. For this type of analysis, the data should ideally cover the whole brain with no gaps between the slices. This will tend to increase the number of slices required, and consequently also the total scan time.

For diffusion anisotropy-related applications, the SNR and scan time are often improved by increasing the slice thickness and/or reducing the image resolution. Therefore, the data produced typically have large voxels and contain a relatively large amount of partial volume averaging (see section 3.3.4). However, this type of application is relatively insensitive to these effects, and suitable data can be acquired within clinically acceptable scan times.

Chapter 7 presented results obtained using DTI in conjunction with VBM to investigate the relationship between language deficits and imaging abnormalities remote from the lesion site in patients with unilateral basal ganglia infarctions. Correlations were found in several cortical language areas remote from the infarct site between language performance and grey matter density, calculated from high resolution structural data. A

correlation was also observed in a region in the left deep frontal white matter between language performance and both white matter density and diffusion anisotropy. These abnormalities were not detected using conventional MR image acquisition and analysis techniques. This highlights the importance of considering abnormalities beyond those observed on conventional MR imaging when investigating structure-function relationships. It can also be seen from this study that the DTI protocol used was of sufficient quality to perform a rigorous quantitative analysis of anisotropy maps.

8.2 DIFFUSION TENSOR TRACTOGRAPHY

Fibre-tracking or tractography is a methodology that is still under development, but has the potential to provide new applications for diffusion tensor data. Applications to date include studies of band heterotopia [12], Wallerian degeneration [2], brain tumours [13], multiple sclerosis [14] and periventricular leukomalacia [15]. However, as demonstrated in chapter 5, these techniques are highly dependent on SNR, partial volume effects, and specific anatomical properties of the structure being tracked, such as anisotropy. This dependence stems from the fact that errors in the major eigenvector may cause the track to follow structures adjacent but unrelated to that intended. This places more stringent constraints on the quality of the data required for successful fibre-tracking.

First, the voxel size should be kept small in order to minimise partial volume effects. Although the major tracts may initially seem relatively unaffected by these problems, there are many regions where these pathways cross or run immediately past each other. In these areas, erroneous connections may be inferred due to partial volume effects [2].

Second, due to the intrinsically 3-dimensional nature of the fibre-tracking process, it is essential to ensure that the diffusion tensor data set used to perform the tracking is also 3-dimensional: it should consist of contiguous slices covering the entire structure of interest, and the voxels should be as isotropic as possible to avoid directional bias.

In order to satisfy both of these constraints, the acquisition should consist of a large number of contiguous thin high-resolution slices. This implies an increase in the scan time due to the increased number of slices to be acquired, as well as a marked reduction in SNR due to the reduced slice thickness. This drop in SNR is usually addressed by acquiring multiple averages of the data (or more directions for the diffusion weighting), which results in a large additional increase in the scan time. It may then become necessary to correct for patient movement during the scan. Realigning diffusion-weighted images with different contrasts is a difficult problem, although various techniques have been proposed [16, 17]. It may also be necessary to adjust the b -matrix image-by-image if large rotations are needed in the correction, since the direction of the diffusion weighting with

respect to the sample would then be incorrect (this effect has not been investigated so far).

For these reasons, diffusion tensor data sets intended to be used for fibre-tracking will generally require longer scan times than data intended for the analysis of anisotropy maps. In order to assess the requirements that diffusion tensor tractography imposes on the quality of the data, the simulations presented in chapter 5 were performed. Image acquisition sequences can then be designed to meet those requirements.

The simulations presented in chapter 5 also made it possible to assess the limitations and requirements of the tractography algorithm investigated, in terms of its robustness to the various problems previously outlined. In particular, as might be expected, the tracking reliability was demonstrated to be dependent on SNR, step size and the interpolation method. The effects of partial volume averaging were of particular interest, since they are expected to have an increasingly detrimental effect as the size of the structures tracked decreases. The simulations could therefore be used to provide an estimate of the size of the structures that could be reliably tracked, depending on the image acquisition parameters. Conversely, they could be used to determine the image quality required to successfully track white matter structures with particular characteristics.

The most commonly used tracking algorithms are based on the streamlines approach pioneered by Mori *et al* [18] and Conturo *et al* [19], and have a number of inherent limitations. First, only one track is generated per seed point. Due to the presence of noise or partial volume effects, its path will to some degree deviate from the true underlying fibre tract, and may follow entirely unrelated structures if those small deviations lead the track outside the original structure. Furthermore, these algorithms give no indication of the level of confidence that can be placed in the resulting track. Possible alternative paths for the underlying tract from that seed point are not investigated, although this is addressed to some extent by using multiple seed points. In order to perform reliable tractography, algorithms must be capable of producing more informative results.

A more fundamental problem common to most tractography algorithms to date is the use of the diffusion tensor model to relate the measured diffusion-weighted signal to the estimated orientation of the underlying fibre population. A number of studies [20, 21, 22] have shown the inadequacy of this model in regions containing partial volume averaging, where several fibre populations are present that have distinct orientations. It is unfortunately not clear at this time how this problem can be resolved (see section 3.3.5).

The FRET algorithm, described in chapter 6, was developed in order to address some of the issues raised above. It was designed to follow as many possible alternative paths as was computationally feasible, taking into account the effects of noise and partial volume effects, and assign an index of the connectivity with the seed region to each point reached by the algorithm. The current implementation was designed to be as independent as pos-

sible of the model used to relate the diffusion-weighted signal to the orientation of the underlying fibre tract, so that new models could easily be incorporated into the algorithm. Preliminary results obtained using this algorithm are encouraging, although susceptibility to partial volume effects remains an issue. However, since this implementation uses the diffusion tensor model, such problems were anticipated, and may be alleviated by using a more adequate model. For a more complete assessment of the FRET algorithm, simulations would need to be performed on synthetic data where the true outcome is known, as has been done in chapter 5 for the streamlines algorithm.

8.3 REMAINING ISSUES

Diffusion tensor imaging is an emerging technique that holds considerable promise as a means of probing tissue micro-architecture. Maps of diffusion anisotropy have been shown to be useful in identifying abnormalities not visible using other imaging techniques. There are however a number of issues that need to be addressed before diffusion tensor imaging can be used robustly for fibre-tracking applications.

First, depending on the intended application, the required SNR and image resolution may imply impractical scan times, although with careful sequence design, DTI data suitable for most applications can be acquired within a clinically acceptable scan time [23]. The scan time may also be reduced by improving the MR hardware. In particular:

- using larger and faster gradients allows a reduction of the echo time, resulting in increased SNR. In addition, the EPI echo train length can be reduced, resulting in reduced image distortions.
- parallel imaging techniques such as SMASH [24] or SENSE [25] can be used to substantially reduce susceptibility-induced distortions in the image, although the spatial non-uniformity of the noise may become a problem.
- Imaging at higher field strengths may also bring considerable benefits, due to the significant increase in SNR. However, at these field strengths, increased susceptibility artefacts make the use of EPI problematic. Furthermore, dielectric resonance effects in the sample degrade the B_1 field homogeneity, introducing intensity distortions in the image. In addition, T_2 relaxation times are shorter at higher field strengths, reducing the gain in SNR. Since the diffusion effect is not dependent on B_0 (as opposed to the BOLD effect [26]), it is not clear at present whether DTI will benefit significantly from higher field systems.

Second, as previously mentioned, the diffusion tensor model is not adequate in regions affected by partial volume effects. Unfortunately, there is currently no viable alternative

model relating the measured diffusion-weighted signal to the underlying fibre orientation¹. This places restrictions on the practical use of diffusion tractography, since the results obtained may be corrupted. There is therefore a need for a more generally applicable model that would be able to give an estimate of the distribution of fibre orientations present. Such a model will most likely be based on high angular resolution acquisition methods (see section 3.3.5), which may require longer scan times. If such a model was developed, further work would be required to optimise the acquisition schemes for its use.

The third limitation, also related to partial volume effects, is that it may not be possible to distinguish between crossing and 'kissing' fibres in voxels containing distinct fibre populations using diffusion-weighted MRI (see section 3.3.4). Regardless of the model used, the distribution of spin displacements in these two cases may be too similar to allow a distinction to be made.

Despite these limitations, diffusion tensor imaging can provide valuable information that cannot be obtained by conventional imaging methods. Maps of diffusion anisotropy are increasingly being used, and are sufficiently robust for clinical use. Sequences for the acquisition of diffusion tensor data are currently being implemented by the major scanner manufacturers, and clinicians will soon be able to generate maps of anisotropy online routinely. On the other hand, current tractography techniques produce reliable results only if given sufficient prior information. They are therefore currently not sufficiently robust for clinical use. However, research in this area is progressing rapidly, and algorithms suitable for implementation on clinical systems may soon become available.

REFERENCES

- [1] YANG, Q., TRESS, B. M., BARBER, P. A., DESMOND, P. M., DARBY, D. G., GERRATY, R. P., LI, T., AND DAVIS, S. M. Serial study of apparent diffusion coefficient and anisotropy in patients with acute stroke. *Stroke* 30 (1999), 2382–2390.
- [2] PIERPAOLI, C., BARNETT, A., PAJEVIC, S., CHEN, R., PENIX, L., VIRTÀ, A., AND BASSER, P. Water diffusion changes in Wallerian degeneration and their dependence on white matter architecture. *NeuroImage* 13 (2001), 1174–1185.
- [3] WERRING, D. J., TOOSY, A. T., CLARK, C. A., PARKER, G. J. M., BARKER, G. J., MILLER, D. H., AND THOMPSON, A. J. Diffusion tensor imaging can detect and quantify corticospinal tract degeneration after stroke. *Journal of Neurology, Neurosurgery, and Psychiatry* 69 (2000), 269–272.
- [4] BAMMER, R., AUGUSTIN, M., STRASSER-FUCHS, S., SIEFERT, T., KAPELLER, P., STOLL-BERGER, R., EBNER, F., HARTUNG, H.-P., AND FAZEKAS, F. Magnetic resonance diffusion tensor imaging for characterizing diffuse and focal white matter abnormalities in multiple sclerosis. *Magnetic Resonance in Medicine* 44 (2000), 583–591.
- [5] FILIPPI, M., CERCIGNANI, M., INGLESE, M., HORSFIELD, M. A., AND COMI, G. Diffusion tensor magnetic resonance imaging in multiple sclerosis. *Neurology* 56 (2001), 304–311.

¹Although techniques such as diffusion spectrum imaging [27] may provide a more robust estimate of the distribution of the fibres, they require wholly impractical scan times.

- [6] ARFANAKIS, K., HERMANN, B. P., ROGERS, B. P., CAREW, J. D., SEIDENBERG, M., AND MEYERAND, M. E. Diffusion tensor MRI in temporal lobe epilepsy. *Magnetic Resonance Imaging* 20 (2002), 511–519.
- [7] FOONG, J., MAIER, M., CLARK, C. A., BARKER, G. J., MILLER, D. H., AND RON, M. A. Neuropathological abnormalities of the corpus callosum in schizophrenia: a diffusion tensor imaging study. *Journal of Neurology, Neurosurgery, and Psychiatry* 68 (2000), 242–244.
- [8] CHABRIAT, H., PAPPATA, S., POUPEON, C., CLARK, C. A., VAHEDI, K., POUPEON, F., MANGIN, J. F., PACHOT-CLOUARD, M., JOBERT, A., LE BIHAN, D., AND BOUSSER, M.-G. Clinical severity in CADASIL related to ultrastructural damage in white matter: in vivo study with diffusion tensor MRI. *Stroke* 30 (1999), 2637–2643.
- [9] ERIKSSON, S. H., RUGG-GUNN, F. J., SYMMS, M. R., BARKER, G. J., AND DUNCAN, J. S. Diffusion tensor imaging in patients with epilepsy and malformations of cortical development. *Brain* 124 (2001), 617–626.
- [10] AGARTZ, I., ANDERSSON, J. L. R., AND SKARE, S. Abnormal brain white matter in schizophrenia: a diffusion tensor imaging study. *NeuroReport* 12 (2001), 2251–2254.
- [11] KLINGBERG, T., HEDEHUS, M., TEMPLE, E., SALZ, T., GABRIELI, J. D. E., MOSELEY, M. E., AND POLDRACK, R. A. Microstructure of temporo-parietal white matter as a basis for reading ability: evidence from diffusion tensor magnetic resonance imaging. *Neuron* 25 (2000), 493–500.
- [12] ERIKSSON, S. H., SYMMS, M. R., RUGG-GUNN, F. J., BOULBY, P. A., WHEELER-KINGSHOTT, C. A. M., BARKER, G. J., DUNCAN, J. S., AND PARKER, G. J. M. Exploring white matter tracts in band heterotopia using diffusion tractography. *Annals of Neurology* 52 (2002), 327–334.
- [13] STIELTJES, B., KAUFMANN, W. E., VAN ZIJL, P. C. M., FREDERICKSEN, K., PEARLSON, G. D., SOLAIYAPPAN, M., AND MORI, S. Diffusion tensor imaging and axonal tracking in the human brainstem. *NeuroImage* 14 (2001), 723–735.
- [14] TENCH, C. R., MORGAN, P. S., WILSON, M., AND BLUMHARDT, L. D. White matter mapping using diffusion tensor MRI. *Magnetic Resonance in Medicine* 47 (2002), 967–972.
- [15] HOON, A. H., LAWRIE, W. T., MELHEM, E. R., REINHARDT, E. M., VAN ZIJL, P. C. M., SOLAIYAPPAN, M., JIANG, H., V., J. M., AND MORI, S. Diffusion tensor imaging of periventricular leukomalacia shows affected sensory cortex white matter pathways. *Neurology* 59 (2002), 752–756.
- [16] MANGIN, J.-F., POUPEON, C., CLARK, C., LE BIHAN, D., AND BLOCH, I. Distortion correction and robust tensor estimation for MR diffusion imaging. *Medical Image Analysis* 6 (2002), 191–198.
- [17] ANDERSSON, J. L. R., AND SKARE, S. A model-based method for retrospective correction of geometric distortions in diffusion-weighted EPI. *NeuroImage* 16 (2002), 177–199.
- [18] MORI, S., CRAIN, B. J., CHACKO, V. P., AND VAN ZIJL, P. C. M. Three-dimensional tracking of axonal projections in the brain by magnetic resonance imaging. *Annals of Neurology* 45 (1999), 265–269.
- [19] CONTURO, T. E., LORI, N. F., CULL, T. S., AKBUDAK, E., SNYDER, A. Z., SHIMONY, J. S., MCKINSTRY, R. C., BURTON, H., AND RAICHLE, M. E. Tracking neuronal fiber pathways in the living human brain. *Proceedings of the National Academy of Sciences* 96 (1999), 10422–10427.
- [20] TUCH, D. S., REESE, T. G., WIEGELL, M. R., MAKRIS, N., BELLIVEAU, J. W., AND WEDEEN, V. J. High angular resolution diffusion imaging reveals intravoxel white matter fiber heterogeneity. *Magnetic Resonance in Medicine* 48 (2002), 577–582.
- [21] FRANK, L. R. Anisotropy in high angular resolution diffusion weighted MRI. *Magnetic Resonance in Medicine* 45 (2001), 935–939.
- [22] ALEXANDER, D. C., BARKER, G. J., AND ARRIDGE, S. R. Detection and modelling of non-Gaussian apparent diffusion coefficient profiles in human brain data. *Magnetic Resonance in Medicine* 48 (2002), 331–340.

- [23] JONES, D. K., WILLIAMS, S. C. R., GASSTON, D., HORSFIELD, M. A., SIMMONS, A., AND HOWARD, R. Isotropic resolution diffusion tensor imaging with whole brain acquisition in a clinically acceptable time. *Human Brain Mapping* 15 (2002), 216–230.
- [24] SODICKSON, D. K., AND MANNING, W. J. Simultaneous acquisition of spatial harmonics (SMASH): fast imaging with radiofrequency coil arrays. *Magnetic Resonance in Medicine* 38 (1997), 591–603.
- [25] PRUESSMANN, K. P., WEIGER, M., SCHEIDECKER, M. B., AND BOESIGER, P. SENSE: sensitivity encoding for fast MRI. *Magnetic Resonance in Medicine* 42 (1999), 952–962.
- [26] GATI, J. S., MENON, R. S., UĞURBIL, K., AND RUTT, B. K. Experimental determination of the BOLD field strength dependence in vessels and tissue. *Magnetic Resonance in Medicine* 38 (1997), 296–302.
- [27] WEDEEN, V. J., REESE, T. G., TUCH, D. S., WEIGEL, M. R., DOU, R. M., WEISKOFF, R. M., AND CHESSLER, D. Mapping fibre orientation spectra in cerebral white matter with Fourier-transform diffusion MRI. In *Proceedings of the ISMRM* (2000), vol. 8, p. 82.