

Correcting for Motion between Acquisitions in Diffusion MR Imaging

BAI Yu

A dissertation submitted in partial fulfillment

of the requirements for the degree of

Doctor of Philosophy

of

University College London.

Department of Computer Science

2009

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

Abstract

The diffusion tensor (DT) and other diffusion models assume that each voxel corresponds to the same anatomical location in all the measurements. Movements and distortions violate this assumption and typically the images are realigned before model fitting. We propose a set of model-based methods to improve motion correction and avoid the errors that the traditional method introduces. The new methods are based on a three-step procedure to register DWI datasets, and use different reference images for DWIs with different gradient directions for registration, so the registrations take into account the contrast differences of measurements. Performance of the model-based registration techniques depends critically on outlier rejection. We develop new methods for fitting the diffusion tensor to diffusion MRI measurements in the presence of outliers by drawing on the RANSAC algorithm from computer vision. We compare one popularly used outlier rejection method RESTORE in the diffusion MRI literature with our new method. Then, we combine outlier rejection methods with model-based registration schemes, and compare the performance of motion correction with other methods. After aligning the dataset, we also update diffusion gradients for the registered datasets from both traditional and our methods, according to the transformations used in registrations. We develop and discuss a variety of registration evaluation methods using both synthetic and human-brain diffusion MRI datasets. Experiments demonstrate both quantitative and qualitative improvements using our new model-based methods.

Acknowledgements

First and foremost, I would like to thank Dr. Daniel Alexander for being such an excellent supervisor. Without his support I would never have completed my beloved PhD and this thesis. I would also like to thank my second supervisor Prof. Simon Arridge for his kind guidance, and Prof. Angela Sasse for being my assessor for the first year and transfer vivas.

Sincere thanks is to Prof. Gareth J Barker and Dr. David Atkinson for examining this thesis. Thanks to Geoff Parker and Karl Embleton at University of Manchester, Claudia Wheeler-Kingshott and Olga Ciccarelli in Institute of Neurology (IoN) at UCL, and Zoltan Nagy in Wellcome Trust Centre for NeuroImaging at UCL for providing the brain datasets.

I would like to thank my friends at UCL for their technical help and emotional support, Phil Cook, Matt Hall, Ivana Drobnjak, Kathleen Curran, Shahrum Nedjati-Gilani, Kiran Seunarine, Hubert Fonteijin, Laura Panagiotaki, Tony Shepherd, Gemma Morgan, Jenny Edgar, Ai Wern Chung, Zhuang Xiahai, Chou Yen-Chun, Ruan Yuan, Yang Guang and Teng Yu.

Special thanks to my dearest Aunty Patricia and Uncle David. You treat me like your own child. Aunty Mary and Uncle Nick, thank you for regarding me as part of the family. Margie, Charlie and Rowland, thank you for being my good “sister” and “brothers” in UK. Also special thanks to my family in China. I am receiving your love and support all the time across the ocean.

Mum and Dad, I feel so lucky and happy for being your child. You do not only give me chance to live, but also make my life meaningful. Love you more than words could possibly express.

Contents

1	Introduction	14
1.1	Problem statement	14
1.2	Contribution	15
1.3	Thesis overview	15
2	Diffusion MRI	17
2.1	Some basic concepts in MRI	17
2.1.1	Localisation with magnetic field gradients	18
2.1.2	Image intensity, $T1$ and $T2$	19
2.1.3	Echo Planar Imaging	19
2.2	Diffusion-weighted MRI	20
2.3	Diffusion tensor MRI	22
2.4	Diffusion reconstruction beyond tensor	25
2.5	Clinical application of Diffusion MRI	27
2.6	Motion correction motivation	28
2.6.1	Patient motion	28
2.6.2	Eddy current artifact	28
3	Registration	30
3.1	Introduction	30
3.2	Feature space (Image features)	31
3.3	Transformation models	32
3.3.1	Rigid transformation	32
3.3.2	Affine transformation	33
3.3.3	More flexible transformations	33
3.4	Similarity measures	35
3.4.1	Least squares	35

3.4.2	Normalized correlation	35
3.4.3	Correlation coefficient	36
3.4.4	Mutual information	36
3.5	Fourier methods	37
3.6	Optimisation	37
3.7	Interpolation	38
3.8	Practical issues in registration implementation	39
3.8.1	Field of view	39
3.8.2	Local minima problem	40
3.9	Registration assessment	40
3.10	Application	41
4	Registration in DW and DT MRI	42
4.1	Inter-subject registration	42
4.2	Intra-subject registration	44
4.2.1	Cost function	44
4.2.2	Pre-processing and post-processing	45
4.2.3	Reference image	45
4.2.4	Model based registration in medical imaging	46
4.2.5	Model based registration in diffusion imaging	47
4.2.6	Evaluation methods	48
4.3	Summary	49
5	Basic Model-based Registration	50
5.1	Primary model fitting method - FMAM	51
5.2	Models fitted from part of data set - FM	52
5.2.1	FM.cons	53
5.2.2	FM.a	53
5.2.3	FM.b	54
5.2.4	FM.th	55
5.3	Data sets for testing	56
5.3.1	Data sets summary	56
5.3.2	Data resizing	59
5.3.3	Data corruption	60
5.4	Experiments	60

5.4.1	Experiment	60
5.4.2	Summary	66
6	Outlier Rejection Schemes	68
6.1	Two outlier rejection schemes	68
6.1.1	Outlier	69
6.1.2	RESTORE	69
6.1.3	Estimation of sigma	70
6.1.4	RANSAC	70
6.2	Experiments	71
6.2.1	Orientation evaluation with ground truth	72
6.2.2	Experiments with RANSAC parameters	72
6.2.3	Confidence interval in RESTORE and RANSAC	78
6.2.4	Outlier rejection schemes and linear tensor fitting	82
6.3	Conclusion	85
7	Model-based Registration with Outlier Rejection	86
7.1	FMRs	86
7.1.1	FMR methods	87
7.1.2	Experiments	87
7.1.3	Conclusion	89
7.2	Comparison of FMR with other methods	90
7.2.1	Evaluation of transformation matrices	90
7.2.2	Experiments with translated synthetic datasets	92
7.2.3	Experiments on full datasets	103
7.2.4	Landmark evaluation method and experiment	107
7.3	Orientation correction	109
7.3.1	Orientation correction	109
7.3.2	Orientation evaluation without ground truth	109
7.4	Conclusion	118
8	Conclusion	119
8.1	Discussion	119
8.2	Further Outlook	121
8.2.1	Update reference image model	121

8.2.2	Reject 3D measurement using outlier histogram	121
8.2.3	Some alternative choices	121
Appendices		122
A	Tutorial: Motion Correction for Diffusion Weighted Images	123
A.1	Preparation before running mbalign	123
A.1.1	General data files and information	123
A.1.2	Sigma	123
A.1.3	Make slice to check input volume	124
A.2	Run mbalign	125
A.2.1	Run mbalign simply	125
A.2.2	Run mbalign in an advanced way	128
A.2.3	Improve performance	128
A.2.4	Update Gradient	130
A.2.5	More hints	130
A.3	Acknowledgement	131
A.4	Reference	131
B	Man Page of mbalign	132
B.1	NAME	132
B.2	SYNOPSIS	132
B.3	DESCRIPTION	132
B.4	OPTIONS	132
B.5	EXAMPLES	134
B.6	AUTHORS	134
B.7	SEE ALSO	134
Bibliography		135

List of Figures

2.1	Creation of spin echo. Spin phases are initially aligned (a) and gradually dephase due to field inhomogeneity during the first half of the sequence (b). The 180-degree pulse negates the phase of each spin (c), which then continues to dephase at the same rate so they reach alignment again at echo time (d).	18
2.2	Echo planar imaging pulse sequence diagram. The phase and readout gradients control the position in k-space that the measurement corresponds to and are designed to traverse a grid of positions.	20
2.3	Pulsed-gradient spin-echo (PGSE) sequence [3]	21
2.4	Isotropic tensor(a), anisotropic tensor(b) and two crossing anisotropic tensors(c)	23
2.5	Trace (a) and fractional anisotropy (b) images from an axial slice through a healthy human-brain	24
2.6	Principal directions of diffusion tensors in an axial slice through a healthy human-brain	25
2.7	Colour coded DT image in an axial slice through a healthy human-brain	26
2.8	Human brain pathways recovered from DT-MRI data	26
3.1	A set theory representation of the entropies involved when combining two images [128]	36
4.1	Original colour coded tensor image (a) and (c), and rotated image without re-orientation (b) and (d)	43
5.1	Comparison between Traditional Method (top) and Model Fit Correction (down)	50
5.2	Model-Based Registration Methods Flow Chart	51
5.3	Flow Chart of Step 1 in FM.cons	53
5.4	Flow Chart of Step 1 in FM.a and FM.b	54
5.5	Flow Chart of Step 1 in FM.th	55

5.6	Image Histograms: (a) grey-scale value histogram; (b) log of grey-scale value histogram	56
5.7	Synthetic Dataset I: last 3 of 14 measurements are translated by 8.5mm	58
5.8	Data Simulation Flow Chart	59
5.9	Rotation corruption simulation steps (a) head without rotation, (b) equivalent head contrast to the head with rotation (c) final corrupted image caused by head rotation	60
5.10	Output of Traditional Method in Experiment	62
5.11	FMAM Reference Images (a) and FM.cons Reference Images (b)	63
5.12	(a) Reference images used to identify the corrupted measurements in FM.a method; (b) Difference S_m between corrupted measurements and reference images in FM.a method	65
5.13	(a) Reference images used to identify the corrupted measurements in FM.b method; (b) Difference S_m between corrupted measurements and reference images in FM.b method	66
5.14	(a) Thresholded corrupted images; (b) Mean thresholded zero-weighted image, and the plot of differences S_m between thresholded corrupted images and mean thresholded zero-weighted image in FM.th method	67
6.1	Plots of normalised diffusion-weighted gradients on xyz-space (a) and projected onto the xy-plane (b). What the red and blue arrows pointed are the gradients for the 3rd and 9th DWIs.	77
6.2	Plots of RMSE (FA) and PD against sigma (σ) for RESTORE (a)(b) and RANSAC (c)(d) methods	78
6.3	Plots of number of rejected voxels in 12th DWI against sigma (σ) for RESTORE (a) and RANSAC (b) methods	79
6.4	Plot of number of outliers against DWI measurements using RESTORE (a), RANSAC with $f = 90\%$ (b) and $f = 75\%$ (c)	83
6.5	FA maps of ground truth tensor (a), tensor fitted directly using least-square linear tensor fitting (b), RESTORE (c), RANSAC with $f = 90\%$ (d) and $f = 75\%$ (e)	84
7.1	Synthetic reference image used in FMRestore with $\sigma = 0.05$ from dataset with 3 of 12 DWIs translated by 5 pixels	92

7.2	FAs from registered datasets using traditional method (Row A), FMAM (Row B), FMRestore (Row C) and FMRansac (Row D)	104
7.3	FA maps overlaid with a hand segmented outline of the corpus callosum from the non-diffusion-weighted image	105
7.4	FA maps overlaid with a hand segmented outline of the corpus callosum from the non-diffusion-weighted image, after (a) standard motion correction and (b) FMAM. correction (c) FMRestore correction using NC cost function	106
7.5	FA maps overlaid with a hand segmented outline of the corpus callosum from the non-diffusion-weighted image, after (a) the traditional method and (b) FMRestore correction using MI cost function	107
7.6	Landmarks of the same anatomical location in different measurement images .	108
7.7	Illustration of how rotation affects the effective gradient direction. The arrows indicate gradient directions. (a) Head without rotation. (b) Head with rotation. (c) The unrotated head (after registration) retains the signal from its rotated position.	110
7.8	Diffusion gradient vectors from the whole dataset and four subsets, mapped on unit spheres.	111
7.9	FA maps of DTs fitted from the whole dataset and four subsets	112
7.10	λ_1 image from original dataset.	112
7.11	Three specified regions (coloured in red) for comparing measurement's alignment.	113
A.1	An image volume containing 32 slices from 64 diffusion components.	125
A.2	Using MRICro to see slice volume	126
A.3	View projection of FA map generated by camino	129
A.4	View projection of mask file generated by camino	129
A.5	Illustration of how rotation affects the effective gradient direction. The arrows indicate gradient directions. (a) Head without rotation. (b) Head with rotation. (c) The unrotated head (after registration) retains the signal from its rotated position.	130

List of Tables

6.1	For various setting of k , outlier histograms produced by RANSAC. Measurement 12 is the corrupted image.	74
6.2	For various setting of f , outlier histograms produced by RANSAC, root mean square error (RMSE) of FAs and principal directions (PD) differences between fitted tensors using RANSAC and ground truth. Measurement 12 is the corrupted image.	75
6.3	Outlier histograms, RMSE and PD for various setting of n in RANSAC. Measurement 12 is the corrupted image.	76
6.4	Outlier histograms, RMSE and PD for various setting of σ in RESTORE. Measurement 12 is the corrupted image.	80
6.5	Outlier histograms, RMSE and PD for various setting of σ in RANSAC. Measurement 12 is the corrupted image.	81
6.6	Root mean square error (RMSE) of FAs and principal directions (PD) differences between fitted tensors directly using least-square linear DT fitting, RESTORE, RANSAC with $f = 90\%$ and $f = 75\%$ and ground truth.	83
7.1	RMSE between transformations used to register corrupted measurements (T) and standard matrices (T') using NC cost function	91
7.2	RMSE between transformations used to register corrupted measurements (T) and standard matrices (T') using MI cost function	91
7.3	RMSE between registered and ground truth datasets, from FMRansac with $f = 90\%$, $\sigma = 35$, using normalized correlation (NC) and mutual information (MI) cost functions. The first column N is the number of corrupted measurements. Bold means MI is bigger than NC.	95
7.4	RMSE between registered and ground truth datasets, from FMRansac with $f = 75\%$, $\sigma = 35$	95

7.5	RMSE between registered and ground truth datasets, from FMRansac with $f = 75\%$, $\sigma = 0.05$	95
7.6	RMSE between registered and ground truth datasets, from FMRestore with $\sigma = 0.05$	96
7.7	RMSE between registered and ground truth datasets, from FMRestore with $\sigma = 35$	96
7.8	RMSE between registered and ground truth datasets, from FMAM.	96
7.9	RMSE between registered and ground truth datasets, from traditional method.	97
7.10	RMSE of FAs otherwise, as Table 7.3, from FMRansac with $f = 90\%$, $\sigma = 35$	97
7.11	RMSE of FAs from FMRansac with $f = 75\%$, $\sigma = 35$	97
7.12	RMSE of FAs from FMRansac with $f = 75\%$, $\sigma = 0.05$	98
7.13	RMSE of FAs from FMRestore with $\sigma = 0.05$	98
7.14	RMSE of FAs from FMRestore with $\sigma = 35$	98
7.15	RMSE of FAs from FMAM.	99
7.16	RMSE of FAs from traditional method.	99
7.17	PDs otherwise, as Table 7.3, from FMRansac with $f = 90\%$, $\sigma = 35$	100
7.18	PDs from FMRansac with $f = 75\%$, $\sigma = 35$	100
7.19	PDs from FMRansac with $f = 75\%$, $\sigma = 0.05$	101
7.20	PDs from FMRestore with $\sigma = 0.05$	101
7.21	PDs from FMRestore with $\sigma = 35$	101
7.22	PDs from FMAM.	102
7.23	PDs from traditional method.	102
7.24	Mean of covariance matrix Trace (MCT)	108
7.25	Mean & (STD) of λ_1 's of dyadic tensors from four groups in four regions (the larger the better and 1 is perfect)	113
7.26	Mean & (STD) of $\gamma = -\log(1 - \lambda_1)$'s of dyadic tensors from four groups in four regions (the larger the better)	114
7.27	Mean & (STD) of STD of FAs from four groups through the voxels in the selected volumes/regions	114
7.28	Mean & (STD) of FAs from four groups	115
7.29	Mean & (STD) of STD of Traces	115
7.30	Mean & (STD) of Traces	116

Chapter 1

Introduction

Diffusion MRI (magnetic resonance imaging) is an innovation in MRI, [25]. It measures the local water diffusion properties in the material being imaged. In fibrous material, the diffusion in directions perpendicular to the fibre is hindered by cell walls, so water diffuses more easily along axon bundles. Thus, diffusion MRI provides fibre directions by revealing directions in which water diffuses furthest. One of the popular diffusion MRI models is diffusion tensor (DT) [20], which assumes that the molecular displacements obey a Gaussian probability density function. Brain white matter fibres connect different regions of grey matter. By following fibre directions from point to point through the image, we can recover the trajectory of white matter fibres. Diffusion MR imaging reveals the connectivity between the different regions of the brain, and provides information about the microstructure of white matter.

1.1 Problem statement

In DT-MRI, a number of diffusion-weighted (DW) images with different diffusion-weighting gradient directions are taken during scanning. During such a long time, small head movements are not easy to avoid. Diffusion-weighted MRI typically uses spin-echo sequence with echo-planar imaging (EPI) readout, which also induces displacement and distortion [70]. However, the tensor calculation assumes that each voxel corresponds to the same anatomical location in all the measurements. That means to fit the diffusion tensor, all the measurements need to align properly. A small bulk motion can cause an unmatched measurement value to be used during the tensor fitting.

The traditional correction scheme for motion within the data set selects a reference image from inside the series, and registers all the other images in turn to this fixed reference. Since the non-diffusion-weighted image does not suffer from the distortions induced by eddy current and

has the highest SNR, it is usually chosen as the reference image for registration [59]. Normally, each DWI is registered to the reference by a 3D rigid or affine transformation computed by maximizing a similarity measure such as cross-correlation (CC) or mutual information (MI).

The problem with the traditional method is that measurements with different gradient directions have different contrasts. Although MI cost function has been commonly accepted in DWI registration, the contrast differences between the DWIs and reference image have not been solved. Thus, the problem I address in this thesis is correcting motion distortion in diffusion MRI by using a model-based registration technique. It predicts separate reference images for each diffusion gradient and thus avoids the mismatching caused by the intensity differences between component images.

1.2 Contribution

In this thesis, we have several novel contributions:

- propose a set of new model-based registration method for diffusion MRIs;
- apply a computer vision technique, RANSAC, to fit diffusion tensor with outlier rejection;
- combine outlier rejection schemes with our model-based registration scheme;
- propose a new orientation correction technique to improve diffusion MRI datasets after alignment;
- discuss or provide several evaluation methods for diffusion MRI registration.

1.3 Thesis overview

Chapter 2 is a brief introduction to diffusion MRI, and explains the problem of motion during acquisition. In Chapter 3, we review the literature on image registration. Chapter 4 reviews registration algorithms used in diffusion MRI, and mainly focuses on intra-subject aspect. Chapter 5 proposes our basic model-based registration methods for motion correction in diffusion MRI, including primary model and some slightly advanced ones. Driven by the conclusion of Chapter 5, in Chapter 6 we discuss outlier rejection schemes, RESTORE and RANSAC. Then in Chapter 7, we combine outlier rejection schemes with model-based registration methods. And then run experiments on a set of synthetic and several original MRI datasets, and use various

evaluation methods to describe the results. Results prove our new methods improve alignment of images from other protocols. We draw conclusions and list some possible further work in Chapter 8 in the end.

Chapter 2

Diffusion MRI

This chapter introduces diffusion-weighted MR imaging, and focuses mainly on diffusion tensor (DT) MRI. At the end of the chapter, we summarize some outstanding problems with the technique.

2.1 Some basic concepts in MRI

The basic physical principle underlying of Magnetic resonance imaging (MRI), or Nuclear magnetic resonance imaging (NMRI), is the magnetic field associated with charged particles in motion [118]. The hydrogen nuclei in water molecules inside scanned subject aligns with the main magnetic field (B_0), and creates a net magnetic moment, M , parallel to B_0 . Then a radio-frequency (RF) pulse is applied perpendicular to B_0 , and causes M to tilt away from B_0 . When the RF pulse is turned off, the hydrogen nuclei realigns with field B_0 , which is referred as relaxation. During relaxation, the nuclei loses energy and emits its own RF signal. This signal is referred as the free-induction decay (FID) response signal. The FID signal is measured and reconstructed to MR images.

The period of the RF pulse sequence is the repetition time, TR. The spin echo signal can be measured at any time within the TR. The time between the RF pulse is applied and the spin echo signal is measured is the echo delay time, TE. The spin echo process is illustrated in Figure 2.1. After a 90° RF pulse, all the spins with different Larmor frequencies are aligned (Figure 2.1(a)); then over the first TE/2 time, the spins dephase (Figure 2.1(b)). Because of B_0 field inhomogeneities, some spins dephase faster than others; at the time of TE/2, a 180° RF pulse is applied, which negates the phase of the spins (Figure 2.1(c)). All the spins continue to precess at the same frequency; at the time of TE, the spins rephase and form an echo, and the spin echo signal is measured (Figure 2.1(d)).

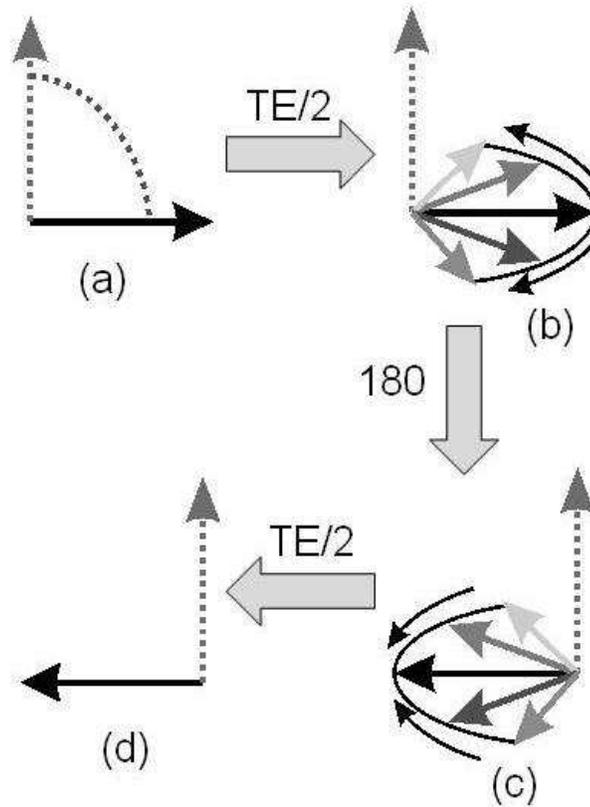


Figure 2.1: Creation of spin echo. Spin phases are initially aligned (a) and gradually dephase due to field inhomogeneity during the first half of the sequence (b). The 180-degree pulse negates the phase of each spin (c), which then continues to dephase at the same rate so they reach alignment again at echo time (d).

2.1.1 Localisation with magnetic field gradients

To reconstruct a 3D image, the spin echo signal needs to be encoded for each dimension. One way to achieve this is to incorporate a gradient into the magnetic field B_0 along a certain direction. Because the Larmor frequency

$$\omega_0 = \gamma B_0, \quad (2.1)$$

where γ is the gyromagnetic ratio and B_0 is the magnetic field strength, *the gradient causes the Larmor frequency to change linearly in the gradient direction*. Thus, for example, a slice perpendicular to B_0 , which is a transverse slice, can be selected by applying a gradient in the same direction as B_0 and choosing the frequency that corresponds to that slice.

For a transverse slice the slice gradient (G_z) is applied along the z-axis, for a coronal slice the slice gradient (G_y) is applied along the y-axis, for a sagittal slice the slice gradient

(G_x) is applied along the x -axis [118]. A slice gradient can be applied along any direction by combining three gradients G_x , G_y and G_z , to achieve any slice direction.

Frequency and phase encoding are used to reconstruct the 2D slice. If a transverse slice is taken, a gradient G_y is applied to vary the frequencies in the y -direction. G_y is then removed and another gradient G_x is applied perpendicular to G_y , so the frequencies of the nuclei vary in the x -direction and phases vary in the y -direction due to the previous gradient G_y . In this case, y -direction samples are encoded by phase and x -direction samples are by frequency. The 2D encoded image is transformed from frequency domain to spatial domain, by using 2D Fourier Transform.

2.1.2 Image intensity, $T1$ and $T2$

The image intensity of a tissue depends on its proton density. A stronger spin echo signal corresponds to a higher proton density in the tissue. While, image intensity also depends on the relaxation time $T1$ and $T2$. When people mention non-diffusion-weighted image, normally it means $T1$ -weighted or $T2$ -weighted MR Image. The image contrast of $T1$ -weighted MRI is mainly influenced by $T1$, although it is also sensitive to $T2$; likewise for $T2$ -weighted MRI [118]. $T1$ -weighted image is created typically by using short TE and TR, because $T1$ controls how long it takes for spins to relax between excitations. $T1$ images show better anatomical detail and better differentiation between solid and fluid filled structures. Comparing with $T1$ -weighted image, $T2$ -weighted image is generated using longer TE and TR times, because $T2$ controls how quickly the signal decays after excitations. $T2$ -weighted images often show local pathology more clearly and have high signal intensity from water, haematoma, tumours, inflammation, oedema and proteinaceous fluid.

2.1.3 Echo Planar Imaging

Echo planar imaging (EPI) is a fast scan imaging technique based on gradient echo [90]. In diffusion MRI (which will be introduced in the next section), EPI technique is widely used, since a large number of images are required. EPI's speed comes from the fact that a single slice only uses a single RF pulse, whereas the standard sequence applies an RF pulse for each line of k -space. In EPI, a single excitation is followed by a series of gradient echoes, which is formed by the rapid switching of a strong gradient, as shown in Figure 2.2. Each of the echoes gives a different degree of phase encoding, and together they can be reconstructed to form a slice.

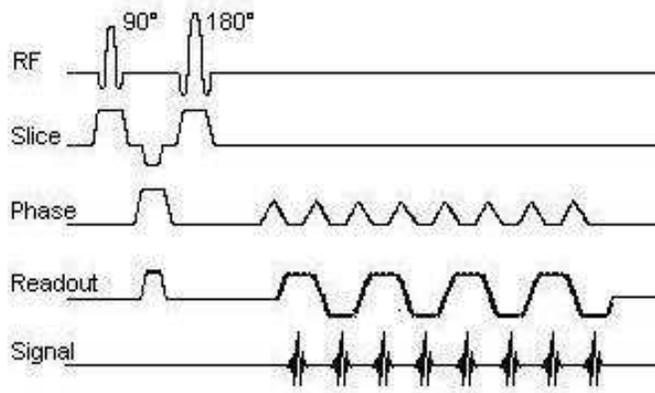


Figure 2.2: Echo planar imaging pulse sequence diagram. The phase and readout gradients control the position in k-space that the measurement corresponds to and are designed to traverse a grid of positions.

Water resonates at a different frequency than fat. EPI has chemical shift artifacts, because the bandwidth per pixel in the phase encoding direction is relatively small, leading to a chemical shift between water and fat. The MRI scanner mistakes the phase difference as a spatial position difference [118] causing the position of tissue with high fat content, for example, to appear at the wrong position relative to other bits of tissue. The differences between water and air also cause distortions due to changes in magnetic susceptibility that affect the local magnetic field.

The rapid changes of the gradient pulses can generate eddy currents in the surrounding conducting surfaces around the gradient coils. The eddy currents generate additional magnetic fields causing inhomogeneity in the net magnetic field and corresponding distortions of the image. More about eddy current artifacts are discussed in §2.6.

2.2 Diffusion-weighted MRI

DW-MRI measures water diffusion in vivo, over approximately 0.01 second. During this time, the average distance water molecules move is several micrometres. The voxel size of typical in-vivo human-brain scans is approximately 8mm^3 , for example $1.7\text{mm} \times 1.7\text{mm}$ pixels in plane with slices separated by 2.3mm.

In diffusion MRI, symmetric diffusion-weighting gradient-pulses are introduced to the

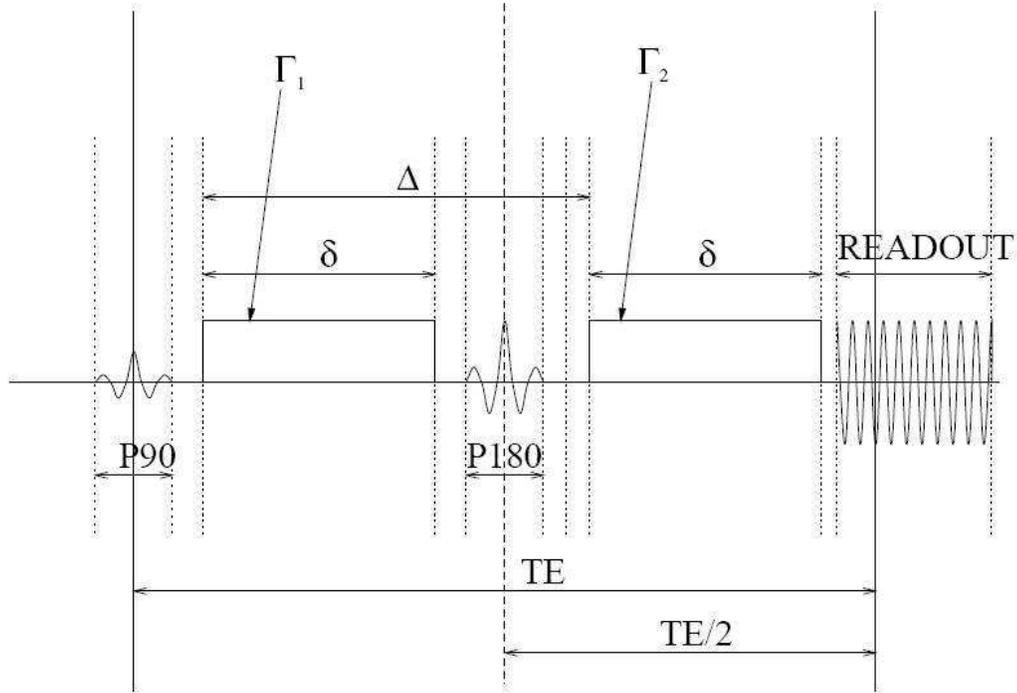


Figure 2.3: Pulsed-gradient spin-echo (PGSE) sequence [3]

standard spin-echo sequence, both sides of the 180° refocusing pulse [127]. Figure 2.3 shows the pulsed-gradient spin-echo (PGSE) sequence [127], where Γ_1 and Γ_2 are the gradient-pulses with the same duration δ , Δ is the pulse separation time, and TE is the echo time. Γ_1 offsets the phase of molecules' spins, and Γ_2 will provide equal but opposite rephasing. If the spins do not move, the second pulse rephases them perfectly, so their magnetizations are all in phase at the time of the echo when the measurement is made. However, if the spins move between the pulses, they are not perfectly rephased and have a distribution of phases at echo time depending on the displacement. That means only spins that move in the direction of the gradients are not fully rephased. The net signal is reduced due to the spins within a voxel having different phases. In this way, the measurements acquired in DW-MRI are sensitive to the motion of water molecules and the more motion, the greater the attenuation in the signal, ie the lower the signal.

DW-MRI measures the probability density function p of particle displacements \mathbf{x} over a fixed time t . p is the scatter pattern of molecules during the diffusion time, and its features provide information about the microstructure. If we can assume that δ is negligible, ie $\delta \ll \Delta$, then the summed magnetization from all the spins is the MR signal $A^*(\mathbf{q})$,

$$A^*(\mathbf{q}) = A^*(\mathbf{0}) \int p(\mathbf{x}) \exp(i\mathbf{q} \cdot \mathbf{x}) d\mathbf{x}, \quad (2.2)$$

where, $A^*(\mathbf{0})$ is the MR signal without diffusion-weighted gradients, $\mathbf{q} = \gamma\delta\mathbf{g}$ is the wavenumber, γ is the position where the gradient pulse offsets the phase of spin and \mathbf{g} is the diffusion weighting gradient vector. The normalised signal is

$$A(\mathbf{q}) = (A^*(\mathbf{0}))^{-1}A^*(\mathbf{q}). \quad (2.3)$$

If the average displacement is zero,

$$A(\mathbf{q}) = \int p(\mathbf{x}) \cos(\mathbf{q} \cdot \mathbf{x}) d\mathbf{x}. \quad (2.4)$$

2.3 Diffusion tensor MRI

Diffusion tensor (DT) MRI [20] computes the apparent diffusion tensor. It assumes that the molecular displacements obey a zero-mean trivariate Gaussian probability p , so that

$$p(\mathbf{x}) = G(\mathbf{x}; D, t), \quad (2.5)$$

where

$$G(\mathbf{x}; D, t) = \frac{1}{\sqrt{(4\pi t)^3 \det(D)}} \exp\left(-\frac{\mathbf{x}^T D^{-1} \mathbf{x}}{4t}\right), \quad (2.6)$$

D is the diffusion tensor and t is diffusion time. Substituting equation (2.6) into (2.4), we get

$$A(\mathbf{q}) = \exp(-t\mathbf{q}^T D \mathbf{q}) = \exp(-b\hat{\mathbf{q}}^T D \hat{\mathbf{q}}), \quad (2.7)$$

where $b = t|\mathbf{q}|^2$ is the diffusion weighting factor and $\hat{\mathbf{q}}$ is a unit vector in the direction of \mathbf{q} . For Gaussian p , we can use the effective diffusion time $t = \Delta - \delta/3$ to account perfectly for finite δ . The diffusion tensor

$$D = \begin{pmatrix} D_{xx} & D_{xy} & D_{xz} \\ D_{xy} & D_{yy} & D_{yz} \\ D_{xz} & D_{yx} & D_{zz} \end{pmatrix} \quad (2.8)$$

is a symmetric three-by-three matrix, where D_{xx} , D_{yy} and D_{zz} are diffusion coefficients along x , y and z axes and the D_{xy} , D_{xz} and D_{yz} are correlation coefficients between the axes. The eigenvalues λ_1 , λ_2 and λ_3 of D determine the shape of p . The eigenvectors of D , \mathbf{e}_1 , \mathbf{e}_2 and \mathbf{e}_3 , give the orientation.

When water moves in all directions equally, the function p is isotropic (Figure 2.4(a)). Typically, grey-matter contains microstructure with no dominant orientation so water movement is hindered equally in all directions and on average the diffusion is isotropic, so $\lambda_1 \approx \lambda_2 \approx \lambda_3$. In regions like brain white-matter, there are bundles of parallel axon fibres connecting different regions of the brain. In the white-matter, water diffusion is hindered more across the fibres

than along them, so the function p is anisotropic with prolate shape (Figure 2.4(b)), [109], with $\lambda_1 \gg \lambda_2 \approx \lambda_3$. Oblate diffusion tensors (Figure 2.4(c)) with $\lambda_1 \approx \lambda_2 \gg \lambda_3$ can occur in fibre crossing regions or regions of sharp bending or fanning.

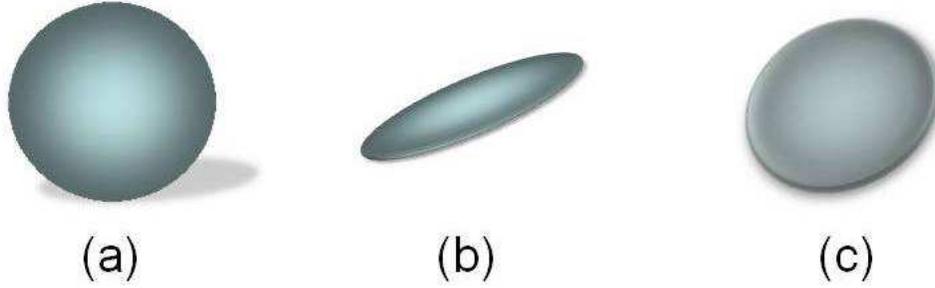


Figure 2.4: Isotropic tensor(a), anisotropic tensor(b) and two crossing anisotropic tensors(c)

To fit the six free parameters in D , a minimum of six measurements $A(\mathbf{q})$ are required with independent \mathbf{q} . Most often, the linear least-squares algorithm is used to fit the tensor D to the log of the measurement via equation (2.7) [21]. However, it is common in DTI to obtain more than the basic 6 measurements to give better orientation invariance. Spherical acquisitions protocols [72] are now standard, which acquire some number $N \geq 6$ measurements with a fixed b value greater than zero and gradient directions spread evenly over the hemisphere, together with some smaller number M measurements with $b=0$. Alexander and Barker [4] suggest that N should be between 6 and 9 times larger than M . From equation (2.7), each $\log A(\mathbf{q}_i), i = 1, \dots, N$ gives a linear constraint on the six elements of D , so we can write all the constraints as a matrix equation

$$\mathbf{A} = \mathbf{B} \cdot \mathbf{D}, \quad (2.9)$$

where $\mathbf{A} = (\ln(A(\mathbf{q}_1)), \dots, \ln(A(\mathbf{q}_N)))^T$, \mathbf{D} contains the elements of the diffusion tensor: $\mathbf{D} = (D_{xx}, D_{xy}, D_{xz}, D_{yy}, D_{yz}, D_{zz})^T$, and \mathbf{B} has $i - th$ row $(-t\mathbf{q}_{i1}^2, -2t\mathbf{q}_{i1}\mathbf{q}_{i2}, -2t\mathbf{q}_{i1}\mathbf{q}_{i3}, -t\mathbf{q}_{i2}^2, -2t\mathbf{q}_{i2}\mathbf{q}_{i3}, -t\mathbf{q}_{i3}^2)$. To estimate D , we solve equation (2.9) using the pseudo inverse of \mathbf{B} : $(\mathbf{B}^T \mathbf{B})^{-1} \mathbf{B}^T \mathbf{A} = \mathbf{D}$. Also the singular value decomposition (SVD) can be used to compute the pseudo-inverse of matrix \mathbf{B} . If matrix \mathbf{B} with SVD is $\mathbf{B} = \mathbf{U} \mathbf{\Sigma} \mathbf{V}^T$, where \mathbf{U} is an m -by- m unitary matrix, the matrix $\mathbf{\Sigma}$ is m -by- n diagonal matrix with nonnegative real numbers on the diagonal, and \mathbf{V}^T denotes the conjugate transpose of \mathbf{V} . The pseudo-inverse of \mathbf{B} with SVD is $\mathbf{B}^\dagger = \mathbf{V} \mathbf{\Sigma}^\dagger \mathbf{U}^T$, where $\mathbf{\Sigma}^\dagger$ is the inverse matrix of diagonal matrix $\mathbf{\Sigma}$, \mathbf{V}^T and \mathbf{U}^T is the transpose of \mathbf{V} and \mathbf{U} .

Some useful scalar values can be derived from D . Two of the most commonly used are trace

$$\text{Tr}(D) = \lambda_1 + \lambda_2 + \lambda_3 \quad (2.10)$$

and the fractional anisotropy (FA) [67]

$$FA = \left(\frac{3}{2} \sum_{i=1}^3 \left(\lambda_i - \frac{1}{3} \text{Tr}(D) \right)^2 \right)^{\frac{1}{2}} \left(\sum_{i=1}^3 \lambda_i^2 \right)^{-\frac{1}{2}}, \quad (2.11)$$

and Figure 2.5 shows images of the scalars over an axial human-brain slice, which are from the dataset ‘‘Fcontrol’’ (More information about the dataset will be introduced in §5.3.1). The trace of the DT is the sum of eigenvalues. It is proportional to the mean squared displacement, and this indicates the mobility of the water molecules, which reflects tissue density. FA is the normalised standard deviation of the eigenvalues, so high values of FA appear in areas in which the local fibres have consistent orientation[67]. It is often used as a marker for white matter integrity.

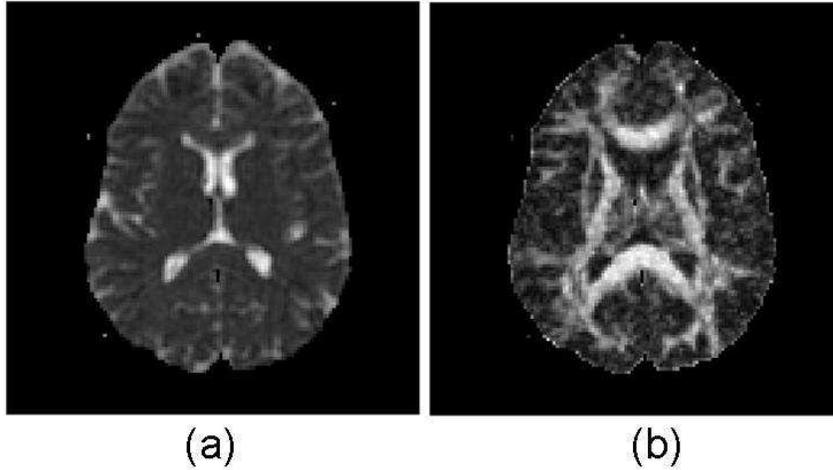


Figure 2.5: Trace (a) and fractional anisotropy (b) images from an axial slice through a healthy human-brain

The principal direction (PD) e_1 of the tensor provides an estimate of the directions of fibres. Figure 2.6 shows a slice of PD image scaled by FA, in which PDs are projected onto the image plane. We can use red, green and blue colours to represent the x, y and z values to the direction vector of tensor image, [120]. Then a colour coded image can be derived as in Figure 2.7.

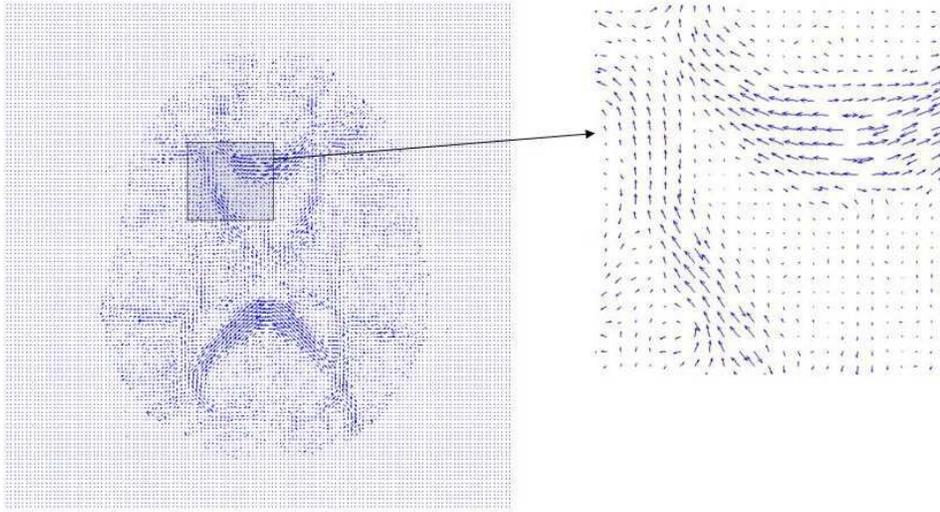


Figure 2.6: Principal directions of diffusion tensors in an axial slice through a healthy human-brain

Tractography algorithms reconstruct the three dimensional structures of bundles of axon fibres by following fibre directions through the image. Generally speaking, the approaches can be divided into two types: line propagation techniques and energy minimization techniques, [100]. Line propagation techniques use local tensor information for fibre reconstruction, in which the information from neighboring pixels is considered [98, 97, 145, 34, 22]. The energy minimization techniques are using global information to find the best path between two pre-determined pixels [108, 136]. An example of tractography images (provided by P. Cook [35]) using streamline algorithm is shown in Figure 2.8.

2.4 Diffusion reconstruction beyond tensor

The DT model quantifies diffusion anisotropy and provides an estimation of a single fibre principal direction, but it cannot reconstruct fibre crossings. Thus various multiple-fibre models and reconstruction algorithms have been developed to recover more information from diffusion MRI measurements. Seunarine and Alexander [122] give a conceptual overview.

The multi-tensor model replaces the Gaussian model in DTI with a mixture of Gaussian densities. It cannot be expressed as a linear function of the measurements so the model fitting requires non-linear optimization. The “ball and stick” model [23, 64] assumes that water molecules belong to either a restricted or a free population. Behrens et al. [23] use an isotropic Gaussian model (“ball”) for the free population, and a Gaussian model which has only one

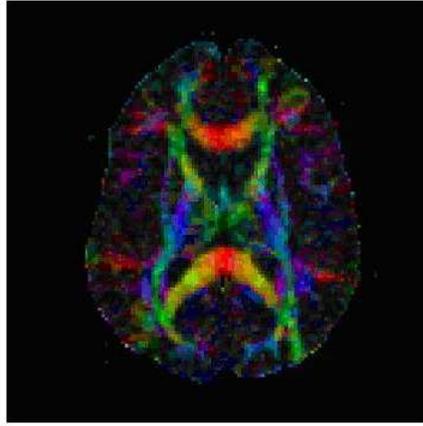


Figure 2.7: Colour coded DT image in an axial slice through a healthy human-brain

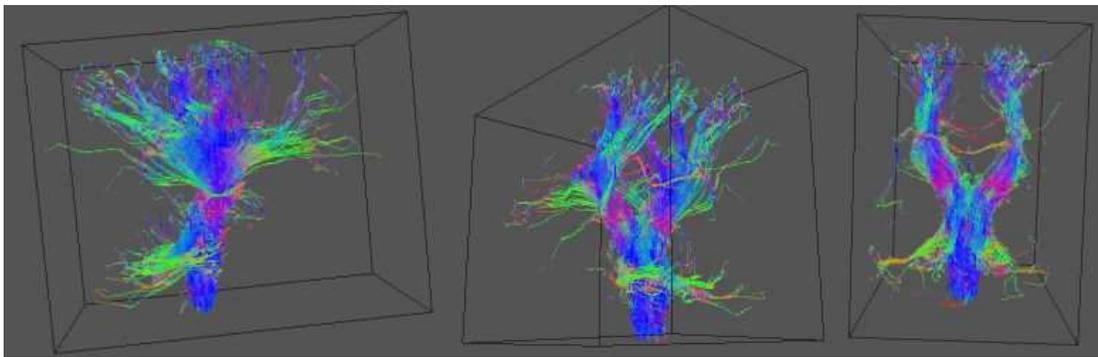


Figure 2.8: Human brain pathways recovered from DT-MRI data

non-zero eigenvalue (“stick”) for the restricted population. Later, Assaf et al’s Composite hindered and restricted model of diffusion (CHARMED) [12] used a cylinder [103] to model the restricted population. Both the “ball and stick” and CHARMED models extend to multiple fibre directions by including multiple anisotropic components. A two-tensor model could provide poorer estimates for a single fibre-orientation, so ideally we would fit a one-tensor model to single fibre and a two-tensor model to cross fibre regions. Alexander et al. [6] propose a method to classify isotropic, one-fibre and two-fibre regions.

The model-based methods recover a finite number of dominant fibre-orientations and do not naturally distinguish fanning or bending configuration from parallel fibre populations. To solve this problem, some other methods try to reconstruct the fibre orientation distribution function (fODF), which is a probability distribution on the sphere. Diffusion spectrum imaging (DSI) [138, 134] and QBall [135] imaging reconstruct the diffusion orientation distribution

function (dODF). The dODF describes the probability of water molecules' movements in each direction, which is different from fibre orientation probability in fODF. DSI attempts to measure p directly from a 3D grid of measurements in \mathbf{q} -space, which means the acquisition requires an order of magnitude more measurements than typical DTI. QBall approximates the dODF using Funk Radon Transform (FRT), which is a transformation that maps one function of the sphere to another, so requires only the spherical sampling pattern of DTI. Spherical deconvolution (SD) algorithm [133] recovers a more direct estimate of the fODF, by deconvolving the measurements with a response function R . Thus it assumes that R is the same for all fibre populations, but different cell sizes, densities, permeability and packing configurations make the assumption limited. Another limitation is that SD is more sensitive to noise, thus filtering step is often used after model fitting. PASMRI [68] recover different functions containing similar structure to the dODF. The PAS (persistent angular structure) is a property of p rather than the true fODF.

2.5 Clinical application of Diffusion MRI

Diffusion MRI provides complementary information for assessing brain information to conventional MRI. In particular, DT-MRI is able to characterize anisotropy and estimate fibre directions [72]. FA is often used to investigate changes in white matter microstructure with disease [18]. The tractography technique can be used to determine the effect of brain tumours on white matter pathways for radiotherapy treatment and planning prior to surgery [99], and to assess differences in connectivity between different population groups [33].

Current clinical application mainly includes brain maturation and aging, cerebral ischemia, multiple sclerosis, epilepsy, metabolic disorders and brain tumor [42].

T1- and T2-weighted MRIs may reflect a decrease in brain water content and an increase in white matter myelination during brain maturation in children, but DTI investigates movement of water molecules and microstructures of the cerebral tissue. Thus it provides a sensitive imaging modality to assess brain maturation in children, newborns, or premature infants [66]. DTI has also been used in detecting age-related degeneration [129].

Cerebral ischemia occurs when blood vessels are occluded or damaged. In conventional MRI and CT, it could only be demonstrated at a later stage, when disruption of the blood-brain

barrier has already led to water and macromolecule exudation from the vascular circulation, leading to accumulation of fluid in the damaged tissue. Diffusion MRI can detect ischemic when conventional MR imaging is still normal [101, 137, 117].

MRI helps to diagnose and monitor disease progression in patients with multiple sclerosis, but inflammation, edema, demyelination, gliosis and axonal loss have identical high signal on conventional T2-weighted images [94], but ADC and FA offer complementary quantitative information to conventional MR techniques in evaluating the structural damage in multiple sclerosis lesions [131, 31].

Despite the information provided by conventional MR, the specification and grading of brain tumors is still limited. Diffusion MRI is increasingly used to investigate various tumor components and assess tumoral invasion from normal tissue or edema [74].

2.6 Motion correction motivation

For fitting the diffusion tensor or other diffusion models, voxels in different diffusion-weighted images must correspond to the same anatomical location. Thus, all the measurement images need to be well aligned, but the misalignment can be caused by patients' motion and eddy current artifact.

2.6.1 Patient motion

As we have mentioned in the previous section, to fit the six free parameters in D , a minimum of six measurements are required, which must be acquired from the patient in one scan. In fact, in order to obtain less noisy DT-MR images, around 50 measurements are typically acquired. The whole scan usually lasts around 20 minutes. During such a long time, some effects such as muscle relaxation or squashing of cushions on the scanner table can cause small head movements, and translations and rotations of the head are not easy to avoid. Baby patients make larger and more frequent movement than adults. For people whose brains are damaged seriously, such as Parkinson's patients, the movements are also more likely to happen.

2.6.2 Eddy current artifact

The varying orientations of the diffusion gradients cause eddy-current-induced 2-D geometric distortions, which often remain in diffusion-weighted images. The effect of residual eddy

currents is to cause image shearing, scaling and shifting. Shearing and scaling are induced by read-direction and phase encode-direction eddy current gradient. Shifting is induced by a B_0 eddy current [70]. In order to get the different diffusion-weighted gradient directions, the strength and location of the residual fields change for each measurement. So each measurement image is distorted individually [59]. Distortions of the individual diffusion-weighted images produce misregistration, which reduces accuracy in computed DT images and derived scalar images [70].

William S. Price [112] summarises approaches for minimizing or coping with the effects of eddy currents, which mainly include hardware solutions, pulse sequence and postprocessing. The postprocessing does not reduce the eddy current distortions of the gradient pulses themselves, but reduce distortion in the acquired images. In image domain, DWIs are often registered to a non-diffusion MR image obtained in the same acquisition [59], and we will discuss more about intra-subject registration in §4.2.

Chapter 3

Registration

Image registration plays an important role in medical image post-processing. Most of the current post-processing methods can be summarised as follows: 1) A set of medical images is acquired and reconstructed using standard methods, which include denoising and intensity correction. 2) Images are registered to ensure that a fixed image coordinate corresponds to the same structure or anatomical coordinate.

This chapter begins with an general introduction to image registration, and then discusses its essential components in details.

3.1 Introduction

Image registration determines a transformation that warps a source image so that its features are in the same position as corresponding features in a target or reference image. Registration seeks the transformation that minimises the difference between the two images, or maximises the correspondence.

Brown [27] gives a theoretical definition of image registration. If we define two 2D images, $I_1(x, y)$ and $I_2(x, y)$, the mapping between them can be expressed as

$$I_2(x, y) = g(I_1(f(x, y))), \quad (3.1)$$

where f is a spatial-coordinate transformation, which maps spatial coordinates x and y to new ones x' and y' ,

$$(x', y') = f(x, y), \quad (3.2)$$

and g is an intensity transformation. The intensity transformation is not always necessary. An example when it is needed is for highly specula objects, where the viewpoint or surface orientation relative to the light source. Or in MRI when different field inhomogeneities cause

different low frequency variations in intensity across the image. Often however, we can regard the expression of registration as

$$I_2(x, y) = I_1(f(x, y)) \quad (3.3)$$

in medical imaging. However, warping diffusion tensor MR images requires a step analogous to grey level correction, which corrects the tensor direction after warping [2], which we can represent similarly to the intensity transformation g . Details of diffusion registration will be introduced later in Chapter 4.

We can break image registration into four key parts following by Brown [27]:

1. Feature space

2. The similarity metric

It defines the optimal transformation for the selected feature set.

3. Search space

It is the space of all possible transformations.

4. Search strategy

Crum et al. [37] do a similar division:

1. The similarity measures.

It measures how well two images match.

2. Transformation model.

It defines how one image can be deformed to match another. It characterises the type of deformations.

3. Optimization process.

It varies the parameters of the transformation model to maximize the matching criterion.

3.2 Feature space (Image features)

When complex distortions are present, selecting a feature space instead of matching on the raw intensities can be more suitable, because it generally reduces the search space and removes irrelevant information. Image features can be **image intensity**, or selected anatomical **landmarks/control points**, segmented **boundaries or surfaces** and **curvatures** [80].

Control points can either be **intrinsic** or **extrinsic** [27]. Intrinsic control points are markers not derived from the data themselves. They are placed in easily identified positions. For example, identifiable structures are placed in known positions in the patients as reference points. When intrinsic points are not available, extrinsic control points can be used, which are determined either manually or automatically. People with knowledge in the domain can make control points according to anatomical structures. Goshtasby [53] summarizes the features used to locate control points automatically, which include corners, line intersections, points of locally maximum curvature on contour lines, centers of windows having locally maximum curvature, and centers of gravity of closed boundary regions.

Boundaries or surfaces in medical images tend to be more distinct than control points, and various segmentation algorithms can locate such high-contrast surfaces and match them between images to guide alignment [41]. When the images are differentiable up to third-order, an alternative way of surface matching is use curvatures/crestlines [95]. That means image can be registered through aligning the crest lines.

3.3 Transformation models

To transform source image I_1 to match the reference image I_2 , a geometric transformation T needs to be determined. Several transformation types are available: rigid, affine and other non-rigid transformations. Rigid transformation can be described by 6 parameters, 3 translations and 3 rotations; affine transformation can be described by 12 parameters, 6 of rigid plus 3 scalings and 3 shears.

For some registration application, T is the most useful outcome. In some study of brain development, T provides the rate of growth of different brain structures from children [130].

3.3.1 Rigid transformation

A rigid transformation preserves relative distances. For example, if P and Q are transformed to P' and Q' then the distance from P to Q is the same as that from P' to Q' . A 3D rigid transformation has six parameters to specify translation in the three orthogonal directions (the rows or columns are perpendicular to each other) and a rotation about any axis. We can write the transformation

$$T(\mathbf{x}) = \mathbf{t} + \mathbf{R}\mathbf{x}, \quad (3.4)$$

where \mathbf{R} is the rotation matrix, \mathbf{x} is the coordinate vector and \mathbf{t} is the translation vector. In rigid transformation, objects in the images retain their shape and size.

3.3.2 Affine transformation

An affine transformation is an extension of the rigid transformation, and allows more complicated distortions. As well as all rigid transformations, the affine transformation also includes shears and scaling. An affine transformation is given by

$$T(\mathbf{x}) = \mathbf{t} + \mathbf{R} \cdot \mathbf{M} \cdot \mathbf{S}\mathbf{x}, \quad (3.5)$$

where in 3D scaling matrix

$$\mathbf{M} = \begin{pmatrix} m_x & 0 & 0 \\ 0 & m_y & 0 \\ 0 & 0 & m_z \end{pmatrix}, \quad (3.6)$$

and shear matrix

$$\mathbf{S} = \begin{pmatrix} 1 & s_3 & s_2 \\ 0 & 1 & s_1 \\ 0 & 0 & 1 \end{pmatrix}. \quad (3.7)$$

It is often written in the form which combines \mathbf{t} and $\mathbf{R} \cdot \mathbf{W} \cdot \mathbf{S}$ together into a single matrix,

$$T(\mathbf{x}) = \mathbf{A}\mathbf{x}, \quad (3.8)$$

where

$$\mathbf{A} = \begin{pmatrix} a_{11} & a_{12} & a_{13} & t_x \\ a_{21} & a_{22} & a_{23} & t_y \\ a_{31} & a_{32} & a_{33} & t_z \\ 0 & 0 & 0 & 1 \end{pmatrix}, \quad (3.9)$$

and to match the dimension of \mathbf{A} ,

$$\mathbf{x} = \begin{pmatrix} x \\ y \\ z \\ 1 \end{pmatrix}. \quad (3.10)$$

3.3.3 More flexible transformations

By adding more degrees of freedom (DOF), rigid and affine transformations can be extended to more general polynomial transformation. For example, the second-order polynomial is defined

as:

$$\begin{pmatrix} x' \\ y' \\ z' \\ 1 \end{pmatrix} = \begin{pmatrix} a_{00} & \dots & a_{08} & a_{09} \\ a_{10} & \dots & a_{18} & a_{19} \\ a_{20} & \dots & a_{28} & a_{29} \\ 0 & \dots & 0 & 1 \end{pmatrix} \begin{pmatrix} x & y & z & x^2 & xy & xz & y^2 & yz & z^2 & 1 \end{pmatrix}^T, \quad (3.11)$$

which has 30 degrees of freedom. However, polynomial transformation can only model global shape changes, not local shape changes [58].

An alternative way to define transformations that can capture more local shape changes is through basis functions. A basis function transformation can be written as

$$\begin{pmatrix} x' \\ y' \\ z' \\ 1 \end{pmatrix} = \begin{pmatrix} a_{00} & \dots & a_{0n} \\ a_{10} & \dots & a_{1n} \\ a_{20} & \dots & a_{2n} \\ 0 & \dots & 1 \end{pmatrix} \begin{pmatrix} \theta_1(x, y, z) \\ \dots \\ \theta_n(x, y, z) \\ 1 \end{pmatrix}. \quad (3.12)$$

The choices of basis function commonly used to represent the deformation field are trigonometric [11] and wavelet bases [7].

Splines registration technique is based on the assumption that a set of corresponding points or landmarks can be identified in the source and target images. The ways of control point determination have been mentioned in §3.2. At these control points, spline-based transformations either interpolate or approximate the displacements to map the location of the control points in the target image into its corresponding point location in the source image.

Thin-plate spline is part of the splines family based on radial basis functions. The transformation is

$$x' = a_1 + a_2x + a_3y + a_4z + \sum_{j=1}^n b_j \theta(|\phi_j - (x, y, z)|), \quad (3.13)$$

with similar expressions for y' and z' , where θ is radial basis function, and $\phi_j, j = 1 \dots n$, are fixed locations in the image. So thin-plate spline gives the displacements between landmarks by finding coefficients of radial basis functions. In the 2D thin-plate spline (TPS), selected points (x, y) are independently displaced within the plane, and the displacements have both x - and y -directional components, thus two TPS models are required to warp one 2D image [80]. There are a number of alternative choices for radial basis functions including multiquadrics and

Gaussians [40, 9].

3.4 Similarity measures

Image registration warps one image volume X into alignment with another one Y . More precisely, the registration problem is to find the geometric transformation of image Y , that maximizes the similarity between X and Y . To do this, we construct a cost function that quantifies the dissimilarity between two images and search for the transformation T^* which gives the minimum cost [69]:

$$T^* = \arg \min C(Y, T(X)), \quad (3.14)$$

where $C(I_1, I_2)$ is the cost function, and $T(X)$ is the image X transformed by T .

Many intensity-based cost functions are used widely, which include least squares (LS), normalized correlation (NC), correlation ratio (CR) and mutual information (MI).

3.4.1 Least squares

Least squares is defined as

$$C^{LS} = \sum_{i=1}^N (Y_i - X_i)^2, \quad (3.15)$$

where Y_i and X_i are the intensities in i -th voxels of images Y and X respectively and the sum is over the N foreground voxels of the images. Least squares is used for intra-modality registration using voxel similarity measures [41]. It is fast to compute and its derivative is easy to obtain. As it can not assume the tissues to be matched have different intensities, it is not suitable for matching images from different modalities.

3.4.2 Normalized correlation

Normalized correlation (NC) measure assumes there is a linear relationship between the intensities in the two image [41], and is defined as

$$C^{NC} = \frac{\sum_{i=1}^N (X_i \cdot Y_i)}{\sqrt{\sum_{i=1}^N X_i^2} \sqrt{\sum_{i=1}^N Y_i^2}}. \quad (3.16)$$

It overcomes some limitations of the least squares cost, in particular, it is not sensitive to intensity scale differences which are common in MRI. However, this cost function is still not appropriate for the registration of different modalities, where the mapping between correspond-

ing intensities is often highly nonlinear.

3.4.3 Correlation coefficient

Similarly to NC, correlation coefficient (CC) similarity is also a linear relationship between the intensity values in two images, and is defined as

$$C^{CC} = \frac{\sum_{i=1}^N (X_i - \bar{X})(Y_i - \bar{Y})}{\sqrt{\sum_{i=1}^N (X_i - \bar{X})^2 \sum_{i=1}^N (Y_i - \bar{Y})^2}}, \quad (3.17)$$

where \bar{X} and \bar{Y} are mean intensity values of images X and Y . CC is not suited for multi-modality registration either.

3.4.4 Mutual information

Mutual information (MI) is particularly suitable for registration of images from different modalities [147], and it is defined as

$$C^{MI} = H(X, Y) - H(X) - H(Y), \quad (3.18)$$

where H is the standard entropy [116]. An image with many different intensities has high entropy, since it contains much information; an image with a single intensity has a low entropy value. MI measures the statistical dependence of one image on another. The lower of joint entropy ($H(X, Y)$), the more similar between the two images are. Figure 3.1 visualises the relationship between entropies, where the size of the circles represents the value of the particular entropy, and the overlapping areas represent the mutual information [128].

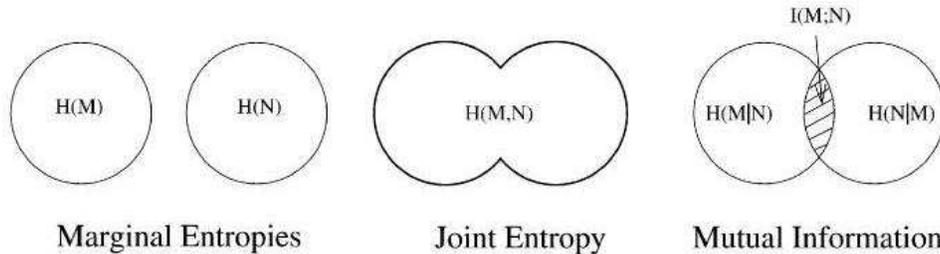


Figure 3.1: A set theory representation of the entropies involved when combining two images [128]

3.5 Fourier methods

Fourier registration methods compute the optimised objective function from a frequency domain. For white noise, registration in time-domain is robust; for frequency-dependent noise due to illumination or changes in sensors, similarity metrics based on the invariant properties of the Fourier transform are good candidates [27]. The basic Fourier registration method, which is called phase correlation, is given by Kugin and Hines [76], but it can only correct for translation. In later work, De Castro and Morandi [30] gave an extension method, which uses a full rigid transformation.

Phase correlation relies on the translation property of the Fourier transform. For two images I_1 and I_2 , they have the same Fourier magnitude, and their shift displacement directly relate to a phase difference, which can be denoted as

$$F[f(x - d_x, y - d_y)] = e^{-j(w_x d_x + w_y d_y)} F[f(x, y)], \quad (3.19)$$

where f is transformation in space domain, d_x and d_y are the displacement on each axis. When the image noise is limited to a narrow bandwidth, the phase difference contributes equally on the frequency information. Thus the phase correlation is well suited to images with this type of noise. Because this method is insensitive to changes in spectral energy, it is also useful for images taken from different sensors. Meanwhile, the significant image white noise, which spreads across the whole frequency domain, can make the method inaccurate.

3.6 Optimisation

When a cost function is chosen, an optimization approach is used to search the parameter space of the transformation. The rigid transformations have 6 parameters (3 rotations and 3 translations), and the affine transformations have 12 parameters, as shown in Eq. (3.9). The optimization algorithm finds the combination of those parameter settings that maximize the similarity measure. Transformations with more degrees of freedom require more sophisticated optimization algorithms.

Search strategy in optimisation could include hierarchical/multiresolution techniques [16, 24, 39, 106], decision sequencings [19], relaxations [65, 111, 114, 124], linear programmings [15], tree and graph matchings [52, 121], dynamic programmings [54, 86, 93, 104], and heuristic searches. More general optimization algorithms in high dimensional image registration may have multiple phases and either include an initial rigid or affine translation, or use

rigid or affine registration to provide a starting estimate for a more complex transformation [41].

One common problem in optimisation is that local minima can cause the algorithm to fail to find the global minimum. That means the returned transformation from optimisation corresponds to a local minimum of the cost function, rather than the desired global minimum. This is a major cause of registration failure. Jenkinson et al. [69] approach the local-minimum problem in two ways:

- (1) A smoothing method is formulated to eliminate discontinuous changes in images;
- (2) A multiresolution framework, which is called hybrid global-local optimization, is used for optimization.

This means the local optimization works progressively from low resolution images to higher resolution images, which reduces the effects of local minima traps.

Rajwade et al. [113] propose a probability density estimation to smooth the noisy signal, and Shams et al. [123] adapted it to histogram for covering the local minimum, which is using “uniform volume histogram” for the optimisation when initialising the rotation parameters.

Gradient descent [13, 126] is an commonly used optimisation algorithm. It takes steps proportional to the negative of the gradient at the current point to find a local minimum of a function. Gauss-Newton algorithm [47, 26] and Levenberg-Marquardt algorithm [115] are used to solve nonlinear optimisation problems. Powell method [110] is used in many optimization procedures, because it does not need second derivatives, and more easily adapted to the implementation.

3.7 Interpolation

The intensity is defined on a grid in digital discrete data, but after geometric mapping, a point in the source image volume and the corresponding point in the target image volume will not fall on a voxel. For example, a geometric transform function F maps the pixel (x, y, z) to a new position (x', y', z') , but the position of the transformed point does not in general lie at a precise 3D voxel location in the output image. The results of the transformation F are a collection of non-integer co-ordinates. To determine the intensity at intermediate geometric locations, interpolation is required.

Interpolates could use alternative neighbourhood methods: nearest neighbour, trilinear, spline and windowed sinc. The nearest neighbour interpolation assigns to a voxel in the target image the value from the closest voxel in the source image; trilinear, spline and windowed sinc are all based on a convolution with kernels, for instance, trilinear interpolation uses the distance-weighted value between eight neighbouring voxels. Normally, the more accurate interpolation costs the longer computational time [105].

Interpolation causes some loss of high frequency information in the images, but can also reduce the noise. Thus, in general, the SNR can increase slightly after interpolation, although simple interpolation algorithms, such as nearest-neighbour interpolation, do not affect the SNR at all. Although the nearest neighbour interpolation used to reduce the image resolution does not change the noise level of the original data, the model fitting and resampling reduces the noise level significantly compared to the original registration problem, because the tensor model fits the data exactly before any corrupting transformations are applied. The noise level in the simulation therefore does not reflect the level of noise in a real clinical application. The role of the simulation is as a proof of concept and comparison of the candidate approaches in ideal conditions. Further experiments with more realistic noise in the images would be required to evaluate expected performance on clinical data.

3.8 Practical issues in registration implementation

There are quite a few good registration softwares, such like Automated Image Registration (AIR) [143], SPM [49], FMRIB's Linear Image Registration Tool (FLIRT) [69, 125], vtkCmic, ect. Here we discuss some of the featured approaches used mainly in AIR and FLIRT.

3.8.1 Field of view

The changing amount of overlap of the reference and target image causes discontinuities. To overcome this, there are two ways to treat voxels outside the field of view (FOV):

1. Treat all values as zero (AIR),
2. Do all calculations strictly in the overlapping region (FLIRT).

Treating all values as zero creates artificial intensity boundaries when the object is not wholly contained within the FOV. When all the calculations are done in the overlapping region, the number of points counted in the overlap region varies, so both the numerator and the denominator of the cost functions (except least squares) will change discontinuously [69].

The loss of information outside FOV is usually seen in the end slices. AIR and SPM assume that all the affected voxels are either zero or can be excluded from further calculations; but FLIRT pads the end-slices, such as increasing the extent of each volume by two slices, which means interpolation from outside the FOV will take on sensible values.

Thresholding can exclude voxels from outside the body. AIR applies a threshold to source and target images to exclude voxels, for example outside the head.

3.8.2 Local minima problem

FLIRT deals with the local minima problem in terms of “apodization of the cost function” and “hybrid optimization technique”. Apodization is to remove or smooth a sharp discontinuity in a mathematical function. In order to apodize the cost function by removing the discontinuities, the weights are added to make the pixel or voxel values near the edge to continuously drop to zero at the edge. Thus, the contributions near the overlapping edge are de-weighted. In hybrid optimization, FLIRT uses a local optimization method with a multiresolution framework. To avoid unnecessary evaluation at low resolutions, different parameters are set according to the scaling changes. AIR also uses multiresolution technique during optimisation, although the benefit they focus is “to improve speed”, but local minima problem is reduced in fact as well.

3.9 Registration assessment

Estimation of accuracy of the registration algorithms is a substantial part of registration process. Zitova and Flusser [147] review basic error classes and methods for measuring the registration accuracy, which includes **localization error**, **matching error** and **alignment error**. Localization error means the displacement of the control point coordinates; matching error is measured by the number of false matches between control points; alignment error is the difference between the mapping model and the distortion image. Alignment error is most commonly used, and the other two are only mentioned by few people. Alignment error can be evaluated in several ways, including mean square error at the control points (CPE), test point error (TPE) and consistency check using multiple cues. CPE is commonly used.

Validation of a registration embodies more than the accuracy verification, but includes: precision, accuracy, robustness, reliability, resource requirements, algorithm complexity, as-

sumption verification and clinical use [85].

3.10 Application

The motivation behind rigid and affine registrations are either combining complementary information about the same patient from different imaging modalities (intermodality registration), or aligning the same imaged with the same modality at different times (intramodality registration). Maguire et al. [73], Peters et al. [132], Schad et al. [82], Levin et al. [81], Faber et al. [45], Evans et al. [44] and Hill et al. [63] all give examples of intermodality registration, and Hajnal et al. [57], Freeborough et al. [1], Woods et al. [142] and Holden et al. [83] talk about the different aspects at intramodality registration.

Intermodality registration is used because patients are commonly to be imaged with more than one modality, such like MRI, CT and PET. Registration of MR and CT images is often applied to head images to help surgery and radiotherapy planning, and a rigid transformation is usually determined [139]. When there are scaling or skew errors in the dataset, an affine transformation can be used occasionally to generate more accurate registration. Because of the low resolution of PET images, registration of MR or CT images with PET images is desirable to make use of anatomical detail from MR or CT images [140].

The same subject can be imaged by the same modality, but separated in time. For example, to monitor disease progression or response to treatment, multiple MR images are acquired at different times. These images can be viewed side-by-side, but small changes between scans are difficult to identify. Intramodality registrations help to identify changes easier.

Compared with intermodality registration, intramodality registration might be easier at first sight, as the images are very similar to one another. However, when high registration accuracy is demanded, great effort must be taken.

Chapter 4

Registration in DW and DT MRI

Registration in diffusion MRI can be used to compare the datasets from different patients (inter-subject), or to align the different images in the same dataset from one patient (intra-subject). For inter-subject registration, the transformations change the orientation of tensors after moving to a new location, thus the tensor orientation needs to be handled properly. In intra-subject registration, the source images are taken at different times, and with different diffusion weighted gradients. Because intra-subject registration is to the focus of this thesis, we only give a general survey on inter-subject registration in §4.1, and concentrate more on intra-subject registration in §4.2.

4.1 Inter-subject registration

We discuss two issues in this section: how to warp a DTI and what similarity measure to drive registration.

Diffusion-tensor images contain orientational information, which reflects the orientation of fibres in the tissue. If we simply transform the voxel values from the original image to the transformed image, then fit the tensor, each DT will retain the same shape and orientation as before the transformation, which means the principal directions remain the same as they were in the original image. This is illustrated in Figure 4.1. Thus, to remain consistent with anatomical structures after warping, **each tensor must be reoriented**. The Finite Strain (FS) [87] and Preservation of Principal Direction (PPD) [2] are commonly used reorientation strategies. The FS uses only the rotation component of deformation field, and the rotation matrix is extracted from the local affine approximation to the warp at each point. Using FS strategy, the deformation component of the transformation, which includes shearing and nonuniform scaling or stretching, is discarded. Alexander et al. [2] propose the PPD method to account for the effect

of the deformation component which also affects the reorientation. Xu. et al. [144] consider not only the principal axis of individual tensor, but also the neighbourhood information contributes to the calculation.

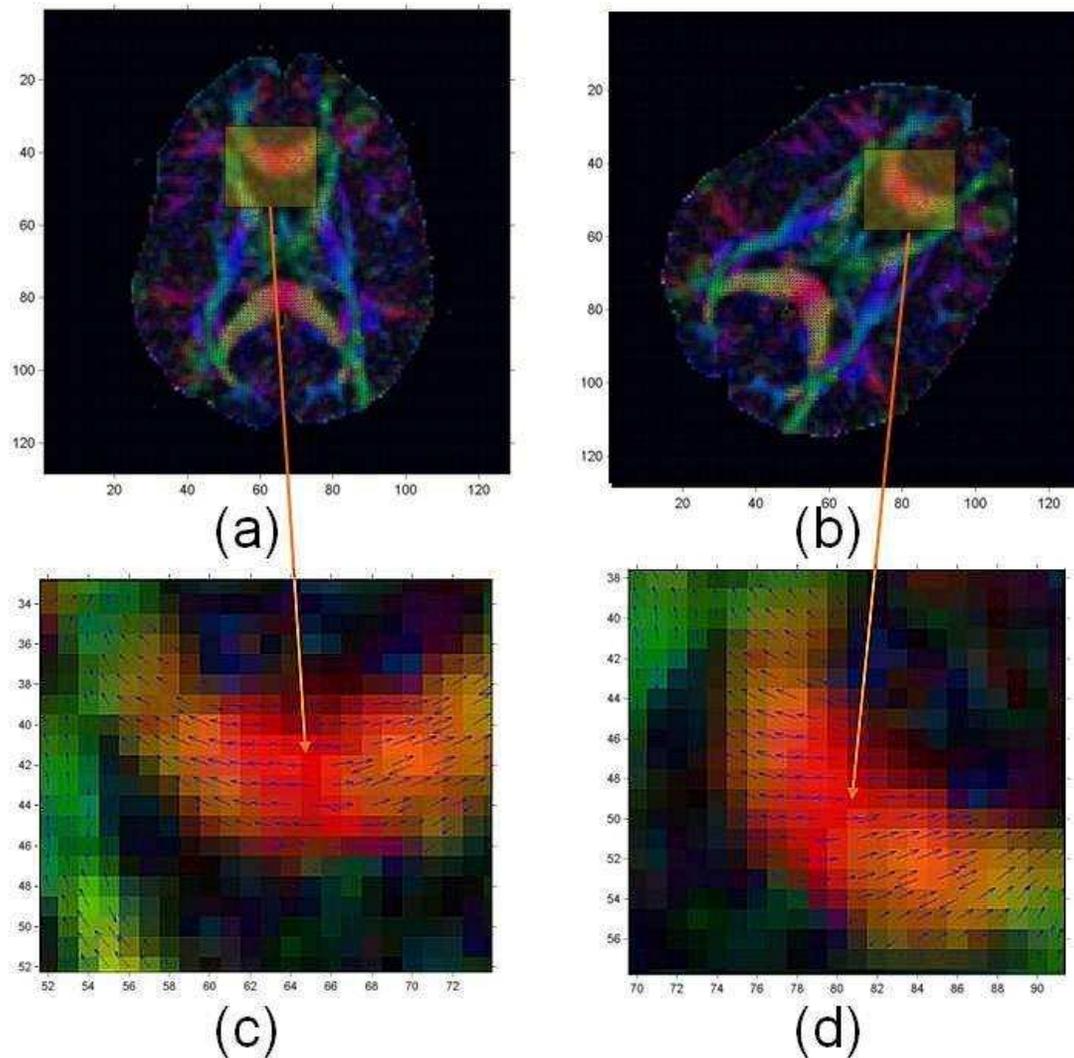


Figure 4.1: Original colour coded tensor image (a) and (c), and rotated image without reorientation (b) and (d)

Many registration approaches for DWI are also provided. Alexander and Gee [5] use the elastic matching algorithm to register DT-MRIs, but the tensor reorientation is not included in the transformation. Curran and Alexander [38] optimize an affine transformation to match diffusion images with no reorientation, FS reorientation and PPD, and they state that synthetic transformations are recovered more accurately using FS and PPD than using no reorientation. Cao et al. [29] develop a deformation registration algorithm for vector fields, which registers

DT images by matching their principal eigenvectors. Park et. al. [107] use a multiple-channel demons algorithm [55] for estimating deformation field in the spatial normalization to drive the registration, which includes a channel of T2-weighted MRI, a channel of FA, a channel of the difference of the first and second eigenvalues, two channels of FA and Trace, three channels of the eigenvalues of the tensor, and six channels of DT components. Registration using all the components of DTs gives the best results. Zhang et al. [146] provide a local affine registration algorithm to register DT images using FS reorientation. The optimal affine transformation is estimated by dividing the image into uniform regions. Gee and Alexander review the literature more completely in [50].

4.2 Intra-subject registration

As we have discussed in §2.6, intra-subject image misalignment can be caused by patients' motion and eddy current artifact. We review approaches to align the DWIs in the dataset in this section. The discussion is in terms of cost function, pre-processing and post-processing, reference image, model-based registration, and evaluation.

4.2.1 Cost function

Correlation coefficient (CC) gives a global linear relationship between two images, so it is not suited for multi-modality registration. Bastin et al. [91] indicate CC does not perform well with diffusion-weighted images acquired with b -values higher than $300s/mm^{-2}$. But Netsch and Muiswinkel [102] think if the neighborhood size is small enough, such as $3 \times 3 \times 3$ voxels, the assumption of a linear transfer function is valid. They use 3D affine registration with local correlation (LC) similarity. LC similarity is based on the CC.

Differences in image contrast caused by the diffusion gradients demand the cost function suitable for multi-modality registration, so MI is commonly used in diffusion image registration [141, 84].

Bammer and Auer [17] show that non-rigid registration, using MI as similarity function, corrects well for eddy-current induced distortion in diffusion-weighted single-shot EPI. Based on [17], Mangin et al. [88] append another step to reduce outlier-related artefacts. They replace objective function in least squares method (ϵ^2 , where ϵ is the residual on the estimated fit) with a similar but more robust function ($\epsilon^2/(\epsilon^2 + C^2)$, where C is a constant).

4.2.2 Pre-processing and post-processing

Masking out voxel values with low intensity values is commonly used as a pre-processing method in DW imaging. Andersson and Skare [8] apply a Gaussian filter to smooth the images before rendering the objective function.

Shams et al. [123] induced methods to initialise rotation and translation plus scale parameters, separately. Translation and scale parameters were determined by registering resized image. They use gradient intensity as measurement to determine the rotational parameters. Let \mathbf{K}_x , \mathbf{K}_y and \mathbf{K}_z be 3D differentiating kernels in x , y and z directions. The gradient images of 3D image \mathbf{I} are calculated using

$$\mathbf{G}_x = \mathbf{K}_x \otimes \mathbf{I}, \mathbf{G}_y = \mathbf{K}_y \otimes \mathbf{I}, \mathbf{G}_z = \mathbf{K}_z \otimes \mathbf{I}, \quad (4.1)$$

where \otimes denotes 3D convolution. The vector field is then expressed in spherical coordinate, where each gradient vector is represented with magnitude, zenith and azimuth angles. Shams et al. then transferred the information of magnitude into binary, and reduced spherical coordinate to a function of only two dependent variables, zenith and azimuth angles. Thus, the rotational misalignment problem was reduced in dimensions.

Rohde et al. [119] indicate that the distortion correction induces intensity errors. So after the registration, they apply a brightness correction function for intensity correction.

4.2.3 Reference image

Normally, registration task for motion correction within the data set selects a reference image from inside the series, and registers all the other images in turn to this fixed reference. Since the non-diffusion-weighted image does not suffer from the distortions induced by eddy current and has the highest SNR, it is usually chosen as the reference image for registration [59]. We refer this as the ‘‘Traditional method’’ in the later experiments of this thesis. For example, the popular *eddy_correct* program [43] in the FSL package [125] (fmrib.ox.ac.uk/fsl/) uses this approach.

In the same dataset, each diffusion-weighted image provides part of the unique information according to diffusion pulse direction, thus different diffusion measurement images do not have the same contrast to locate the same position. So the idea of minimising the sum of squared

differences between two images with different diffusion gradients weighted has problems.

Based on an adaptive intensity correction algorithm proposed by Guimond et al. [56], Ardekani and Sinha [10] use a local intensity-matching algorithm to make a T_2 -weighted image that matches the intensity of a diffusion-weighted image. Since the technique incorporates intensity outlier detection and performs iteratively on a subset of the data, its computation is expensive.

Landman et. al. [77] create a template set as reference images during registration. They declare that their template provide the unique reference image for each diffusion-weighted measurement. As the template is made from the registered diffusion-weighted data from non-diffusion image, there is no true improvement from the standard correction scheme. Because all the experiments are based on the comparison with their created template, the conclusion based on the experiment result is not convincing.

4.2.4 Model based registration in medical imaging

Hayton et al. [61] fit a pharmacokinetic model to correct for motion effects during the acquisition period, and demonstrate qualitatively that the contrast of MR images are enhanced significantly. Later, Buonaccorsi et al. [28] present a locally-controlled 3D translational registration process driven by tracer kinetic model to blood volume. Their technique involves fitting a parametric kinetic model to a time series of measurements in each voxel. They use a five-step iterative scheme in the registration process:

1. Fit the model to the original measurements;
2. Synthesize reference signal maps from the fitted model;
3. Register translation only to match each original time point volume to its corresponding reference volume;
4. Re-fit the model;
5. Repeat steps 2 to 4 until a minimum is found in the model fit errors.

Tracer kinetic model-driven registration method copes well with features that appear and disappear between images. Results show a significant improvement in model parameter estimates in the presence of motion corruption of the image sequence.

Melbourne et al. [92] perform registration repeatedly to an artificial time series of target images generated using the principal components of the current best-registered time-series data.

Because early principal components describe the major trends in the data, image profiles are expected to be relatively noise free. But the method relies on separating the effects of motion and contrast enhancement, which might fail if the registration is not successful. The authors show improved alignment across time series of contrast-enhanced MRI and later apply the same idea to aligning diffusion-weighted images. The approach is less suitable for the latter problem, because the orientational dependence in diffusion MRI may cause the algorithm to require larger numbers of principal components to capture the variation. However, in practice, early results appear promising.

4.2.5 Model based registration in diffusion imaging

Andersson and Skare [8] use a quadratic cost function to drive the optimisation. In other words, the algorithm minimizes the sum of the squared errors between the data and the linear model, by simultaneously registering all images and fitting the DT in each voxel. The residual error is directly minimized in two steps:

1. removing all the variance accounted by the diffusion tensor;
2. using the remaining variance to drive the registration.

This scheme can correct both eddy-current-induced distortion and subject motion, but only works when more than 6 measurements are available. They suggest that 12 to 20 gives the best results. Over all, the work provides a compelling theoretical approach to improve correction for distortions in DW-MRI. Theoretically, the method estimates eddy current-induced distortion and subject motion without need of any additional measurements. However, the method is complex and computationally demanding and anecdotal wisdom from several groups suggests that the algorithm is not very robust in practice. Most likely the key limitation is that initially poorly aligned measurements contribute strongly to the objective function that quantifies alignment. The methods we propose in later chapters avoid this problem by using outlier rejection to remove those contributions.

Previously in §2.1.3 Echo Planar Imaging and §2.6.2 Eddy current artifact, it has been mentioned that EPI causes distortion. Haselgrove and Moore [59] point out each diffusion measurement image is distorted individually. That is because, in order to get the different diffusion-weighted gradient directions, the strength and location of the residual fields change for each measurement. Jezzard et al. [70] state that the particular deformation model suited to diffusion-weighted EPIs is a slice-wise subset of an affine transform consisting of one shear, one scale and one translation parameter. Andersson and Skare [8] suggest models for how dis-

tortions vary with slice position and gradient direction. Spatial constraints model describes a linear or nonlinear relationship between slice number and distortion (such as shear) parameter. Another model is gradient constraints, which focuses on the relationship between diffusion gradient direction and distortion parameter. By implementing the two models, the dimensionality of the parameter space could be reduced.

4.2.6 Evaluation methods

The assessment of improvement of diffusion tensor image registration is mainly based on visual side-by-side comparison of images and parametric maps before and after registration [60, 119].

The most general quantitative evaluation method is the root mean square error (RMSE) between the images component. Image component can be the DWIs, diffusion model coefficients, and some model-derived scalar values such like FA maps. FA is a quantitative measure which marks tissue microstructure, thus is widely used to display appropriate statistical result. Kreher et. al. [75] calculate FA histograms before and after distortion correction.

Consistency testing is used to quantitatively evaluate registration results and compare different registration schemes [102]. The consistency error measures the largest displacement for all voxels x_i , and is defined as

$$\theta_{kl} = \max_{x_i} \|(T_{lk}T_{kl})x_i - x_i\|_2, \quad (4.2)$$

where T_{lk} is the transformation to register image X_k to image X_l ; and T_{kl} vice versa. But consistency testing can only be used when there is misalignment of two images; otherwise there are no transformations of T_{lk} and T_{kl} contributing to equation (4.2).

To overcome the limitation of consistency, Netsch and Muiswinkel [102] introduce cross consistency, which takes transformations from two algorithms (such as MI and LC) into account. If T and S are the registration results from two algorithms, the cross consistency error is defined as

$$\phi_{kl} = \max_{x_i} \{ \|(T_{lk}S_{kl})x_i - x_i\|_2, \|(S_{lk}T_{kl})x_i - x_i\|_2 \}. \quad (4.3)$$

But the authors point out that there is no truth about whether consistency or cross consistency can provide better results, and choice is made depending on the individual performances.

4.3 Summary

To study DTIs between datasets, inter-subject registration attempts to find proper similarity measures to drive registration and warp tensors to retain the same shape and orientation as before the transformation. Intra-subject registration aims to improve dataset quality after acquisition from scanners. Current literatures try to solve the problem by selecting proper cost functions, choosing good reference images, even developing some model driven approaches. MI cost function has been commonly accepted in DWI registration. However, the reference image is still mainly based on using non-DWI in the traditional method, and the contrast differences caused by the diffusion gradients have not been solved.

Chapter 5

Basic Model-based Registration

In this chapter, we propose a new class of model-based registration methods for motion correction in diffusion MRI. To improve the registration and remove the errors that the traditional correction method introduces, all the new model-based registration methods use different reference images to register diffusion-weighted images with different gradient directions, as illustrated in Figure 5.1. In this way, the registrations take into account the individual features

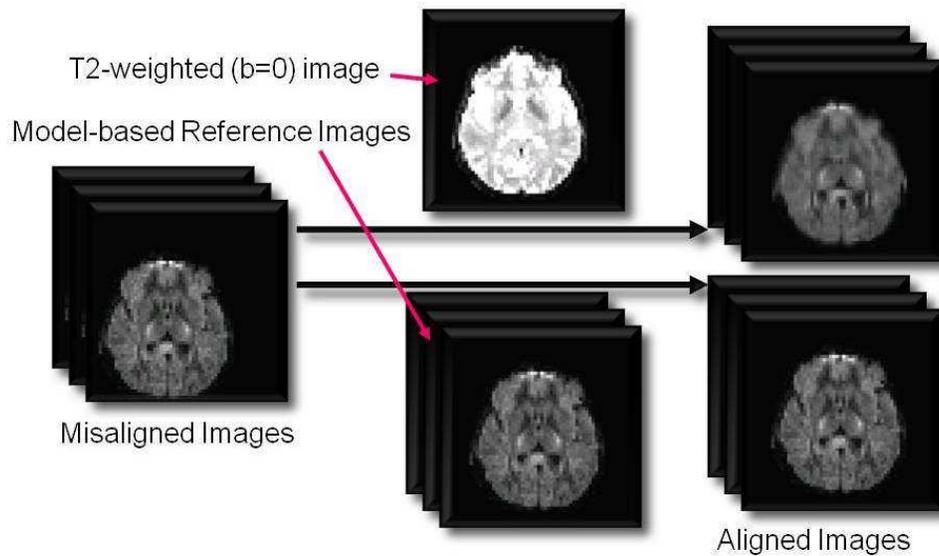


Figure 5.1: Comparison between Traditional Method (top) and Model Fit Correction (down)

of each measurement. Our method could correct eddy current artifact potentially, although we mainly focus on motion problem when designing. The basic idea is:

1. fit the diffusion model to the measurements;
2. synthesize reference data for each measurement from the fitted model;

3. register each measurement volume to the corresponding synthetic reference;
4. repeat previous steps to converge.

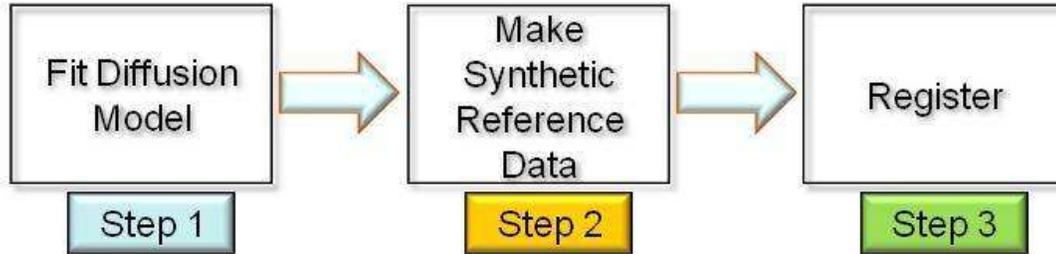


Figure 5.2: Model-Based Registration Methods Flow Chart

We use the simple DT model in step 1, although potentially other diffusion models could also be used. The only difference between the methods lies in the first step of fitting the diffusion tensor model to the measurements. In particular, the methods select different subsets of the measurements for fitting the model. In practice, we found the difference between the first outputs and the outputs from later cycles changes only very slightly. Even though minor fluctuations in the objective function occur during subsequent iterations, we consistently see no further improvement in alignment. Thus, we did not pursue the use of Step 4 in the end for all of our model-based registration methods in favour of seeking alternative algorithms that we describe in later chapters. Thus the various algorithms in this chapter all follow the three-step flow chart illustrated in Figure 5.2.

5.1 Primary model fitting method - FMAM

The most direct model-based registration is a direct adaptation of Buonaccorsi's method [28] to DT-MRI. We call this the FMAM (Fit the Model to All the Measurements) method.

Step 1: Fit the tensor to all scanner output measurement images

This is a diffusion tensor reconstruction step, which takes the scanner output images and fits the DT in each voxel as described in p.23 in §2.3.

Step 2: Making synthetic images

From the fitted tensor D from step 1, we generate target image volumes for each measurement by synthesizing the measurement from equation (2.7) using the fitted D in each voxel.

Step 3: Registration of the scanner output image data set

We register every measurement volume to the synthetic target image with the same \mathbf{q} .

The first two steps are the inverse transformations of each other, but the output measurements from **step 2** contain contributions from all the measurements. The normalized correlation (NC) or mutual information (MI) cost function is used for searching the transformation. We use FLIRT [69] to compute an affine registration here.

Compared with the standard correction scheme (choosing the T_2 -weighted image as the reference image for registration), there are some advantages for FMAM. It chooses different reference images for diffusion-weighted images with different gradient directions for registration as shown in Figure 5.1, so this registration considers the contrast differences of measurements. Since the method does not attempt to reject outliers, the model fitting procedure is influenced by measurements from misaligned images.

5.2 Models fitted from part of data set - FM

Step 1 of FMAM fits the diffusion tensor D to the whole set of diffusion weighted measurements. This means that all the measurements, including any that are poorly aligned, contribute to the fitted D . The corrupted measurements therefore affect the reference image made from D in **step 2**.

In order to avoid corrupt data disturbing the reference images, every sub-method of FM (**F**it the **M**odel) described in this section fits the tensor only to a subset of scanner output measurement images. All the sub-methods of FM still retain the three steps presented in Figure 5.2, but they use various strategies to select measurements for the tensor fitting in **step 1**. They are “FM.cons” “FM.a” “FM.b” and “FM.th”.

All FMs are trying to identify and reject the corrupted 3D DW volumes. In order to fit the DT model, a minimum of six DWIs is required. FM.cons fits the tensor to a fixed set of measurements, which are the first ones acquired during the scan. FM.a, FM.b and FM.th all use different automatic algorithms for selecting well aligned subsets of measurements for tensor fitting in **step 1**.

5.2.1 FM.cons

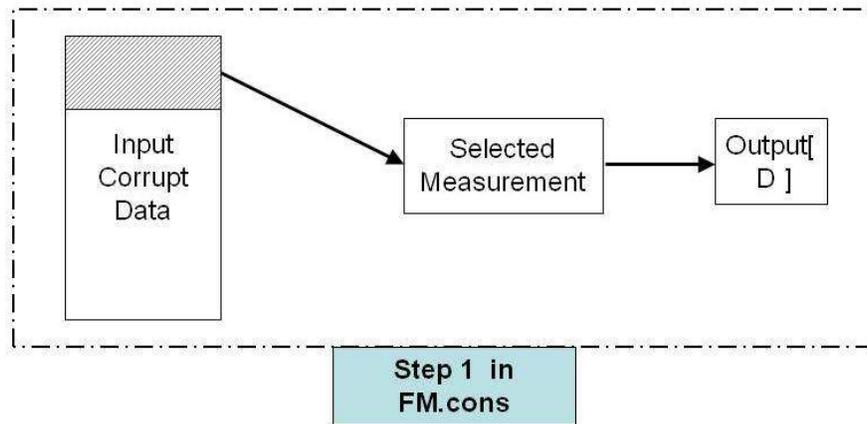


Figure 5.3: Flow Chart of **Step 1** in FM.cons

FM.cons only uses non-DWI(s) and the first few DWIs to do the tensor fitting, shown in Figure 5.3. For example, the first six DWIs taken from patients are used. The idea is that the patient is least likely to move at the start before they get tired. Some effects such as muscle relaxation or squashing of cushions on the scanner table may be more likely to affect the early measurements.

However, not all the patients can keep still, even during the early period of the scan. For example, baby patients make larger and more frequent movement than adults and these movements can occur at any time during the scanning. For people whose brains are damaged seriously, such as Parkinson's patients, the movements are similar. Thus, FM.cons may not always work, and we need more robust methods.

5.2.2 FM.a

In Step 1, FM.a identifies corrupted measurement, and generates synthetic reference image using part of dataset. The method is illustrated in Figure 5.4. FM.a fits DT model using the whole dataset, and generate reference images from the model. This first set of reference images is NOT used for registration, but only to identify corrupted measurements. It makes the identification by looking at the differences between the original images and the reference images which correspond to the same diffusion gradients. The measurements that have the smallest difference are selected, and used to generate reference for registration in Step 2.

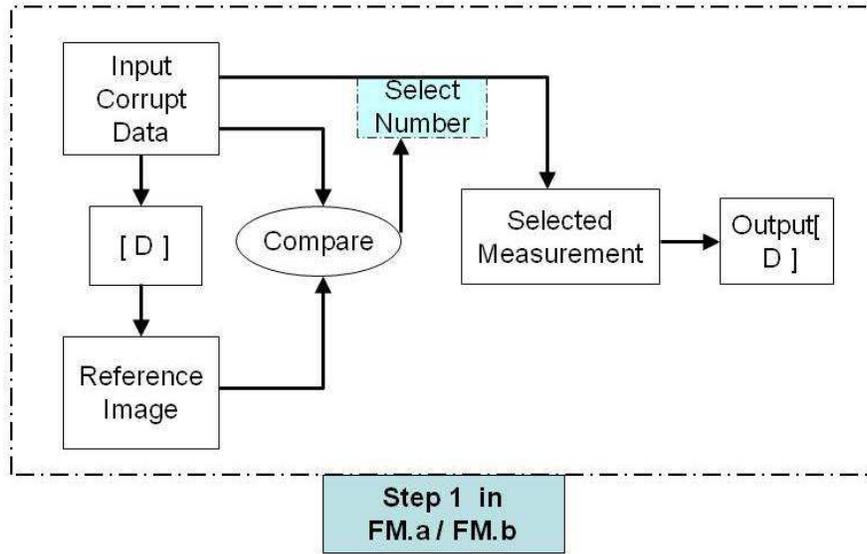


Figure 5.4: Flow Chart of **Step 1** in FM.a and FM.b

To compute the difference between the measurement volume $A(\mathbf{q}_m)$ and corresponding synthetic image $A_{ref}(\mathbf{q}_m)$ with the same gradient \mathbf{q}_m , we use:

$$S_m = \sum_{i=1}^V \sqrt{(A(\mathbf{q}_m)_i - A_{ref}(\mathbf{q}_m)_i)^2}, \quad (5.1)$$

where i indexes each voxel in the whole volume of V voxels.

5.2.3 FM.b

In a similar way to FM.a, FM.b selects measurements by computing the difference between the corrupt data set and the synthetic data, but every synthetic measurement image is built from a separate diffusion tensor, which is fit to all the measurements apart from the one being synthesized. For example, for making the reference image for the first measurement, we use the whole set of measurements **without** the first one to fit the tensor, then make the synthetic reference image from the fitted D matrix, as illustrated in Figure 5.4.

In FM.b, we measure the differences between the input corrupt data and the reference image data set in the same way as FM.a using equation (5.1).

FM.a is vulnerable to the effects of corrupted measurement volumes in the same way as FMAM. The motivation for FM.b is to exclude corrupted measurements from the dataset used to fit the tensor and thus emphasize their difference to the other measurement volumes.

5.2.4 FM.th

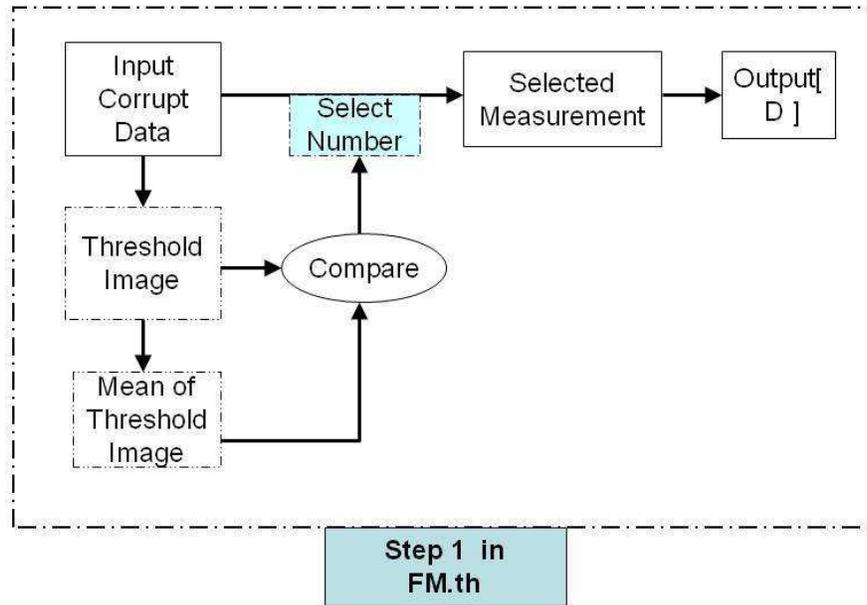


Figure 5.5: Flow Chart of **Step 1** in FM.th

FM.th method uses thresholded DWIs to identify corrupted measurements, and uses selected subset to fit diffusion model in step 1 in a similar way to FM.a and FM.b, as shown in Figure 5.5. The key idea is to identify corrupted measurements as those images whose foreground (brain) regions overlap the foreground regions of the other images least well. We identify foreground regions by a simple threshold on the image intensity. The histograms in Figure 5.6 shows the distribution of grey-scale intensities in one DWI. As there is a huge gray value gap between brain (higher intensity group) and the dark background (lower intensity group), threshold level TH can be chosen manually.

An alternative way is to make sure the same number of voxels remain after the threshold, which means threshold the first measurement by choosing a threshold level TH manually; to the second measurement, increasing the threshold from zero, until the number of foreground pixels is the same as the one in the first thresholded measurement; and keep doing the same work for the rest of the measurement images. This thresholding method is not used in the experiments of this thesis.

Once the threshold scheme is decided, we can implement FM.th scheme following the

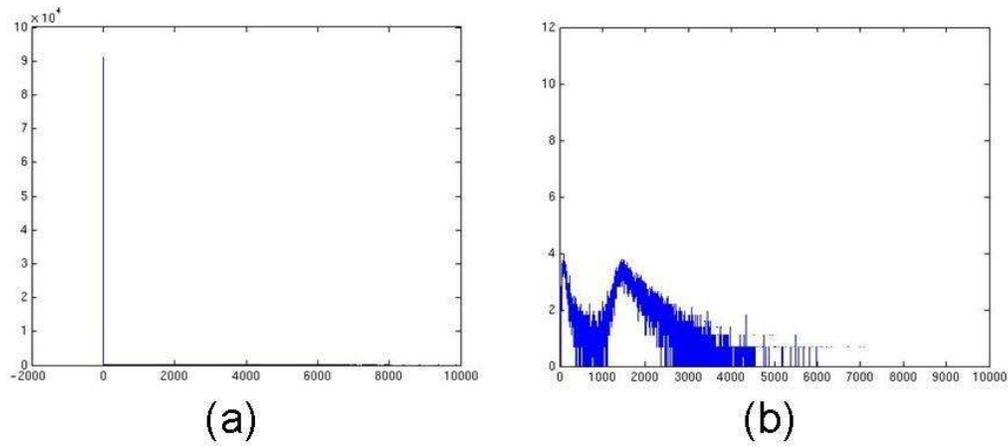


Figure 5.6: Image Histograms: (a) grey-scale value histogram; (b) log of grey-scale value histogram

steps as illustrated in Figure 5.5. FM.th selects the measurements to fit the tensor by computing the difference between the thresholded measurement image $A_{th}(\mathbf{q}_m)$ and the mean thresholded zero-weighted image $Mean(A(\mathbf{0})_{th})$, which can be denoted as

$$S_m = \sum_{i=1}^V (A_{th}(\mathbf{q}_m)_i - Mean(A(\mathbf{0})_{th})_i), \quad (5.2)$$

where i is the voxel index in volume V .

Thresholding images with different \mathbf{q} ignores the detailed microstructure information, so it makes the identification of corrupted measurements easier. However, FM.th relies on having sufficient signal in the DWIs to distinguish brain from background, so will be less effective for higher b-value acquisitions.

5.3 Data sets for testing

In this section, we first list all the datasets used in the experiments of this thesis, and then discuss the procedure of making synthetic datasets with corruption.

5.3.1 Data sets summary

One typical full DW dataset we are using has 66 measurements (6 $b = 0$ and 60 DWIs), and $128 \times 128 \times 60$ voxels for each measurement with the voxel size of $1.72 \times 1.72 \times 2.3\text{mm}^3$. It is acquired using a GE Signa 1.5-Tesla scanner with a standard quadrature head coil. DWIs were

obtained using a single-shot echo-planar acquisition with cardiac gating ($TR = 5 - 6 RR - 6 s$). The DWI images are at a fixed $|q|$ giving $b = 1050 s mm^{-2}$, with $\Delta = 0.04 s$ and $TE = 95 ms$. We call this dataset “Fcontrol” for further notice. Thank to Claudia Wheeler-Kingshott in Institute of Neurology (IoN) at UCL for providing the data.

Another dataset with small motion corruption, which we call “S32”, is used a lot in the experiments of this thesis. It is acquired on a Philips 3T Achieva scanner, using an 8-element SENSE head coil. A PGSE EPI sequence was implemented with $TE = 54 ms$, $TR = 6000 ms$. It has 64 measurements. There are $128 \times 128 \times 32$ voxels for each measurement, with a voxel size of $1.8 \times 1.8 \times 2.1 mm^3$. The dataset has 60 diffusion weighted images with $b = 1200 s/mm^2$, and $\Delta = 0.028 s$. Thank to Geoff Parker and Karl Embleton, University of Manchester, for providing the brain data.

“Fli” dataset has 68 measurements, with $96 \times 96 \times 60$ voxels for each measurement, including 61 DWIs with $b = 1000 s/mm^2$ and 7 with low b value. The voxel has isotropic dimensions of 2.3mm. It is acquired on a Siemens Trio 3T Scanner, using twice refocused diffusion encoding, with $\Delta = 0.035 s$, $TE = 90 ms$ and TR is 2s. Thank to Zoltan Nagy in Wellcome Trust Centre for NeuroImaging at UCL for providing the data.

“Olgacontrol” dataset is not motion corrupted during the scanning acquisition with cardiac gating. It has 68 measurements, including 6 $b = 0$ images and 62 diffusion weighted images with $b = 1600 s/mm^2$, with $\Delta = 0.04 s$, $TR = 2 - 3 s$. Each measurement has $128 \times 128 \times 60$ voxels with a size of $1.7 \times 1.7 \times 2.3 mm^3$. It is acquired on a GE 1.5T Signa Scanner. Thank to Olga Ciccarelli in Institute of Neurology (IoN) at UCL for providing the data.

In order to reduce the time required to test and compare different registration methods in the first stage, we construct a smaller dataset from the full “Fcontrol” dataset. We carefully selected this dataset, because the alignment among all the images appears good so it provides a good basis for constructing datasets with synthetic corruptions for testing. We introduce some artificial corruptions to the reduced-size dataset and create “Synthetic Dataset I” and “Synthetic Dataset II”. The size for each 3D measurement is $64 \times 64 \times 30$ voxels. The noise standard deviation σ is in the range 30 to 40. We simulate the situation that after a long time scan, a person moves unconsciously, so the last few DWIs are translation corrupted.

Synthetic Dataset I contains 2 non-DWIs and 12 DWIs, and the last 3 DWIs are translated by along different directions in xy -plane with the displacement vectors of

$$\mathbf{t}_1 = [-5 \quad -5 \quad 0]', \quad (5.3)$$

$$\mathbf{t}_2 = [5 \quad -5 \quad 0]', \quad (5.4)$$

and

$$\mathbf{t}_3 = [-5 \quad 5 \quad 0]'. \quad (5.5)$$

One slice of each measurement of the corrupted data set is shown in Figure 5.7. As the size of

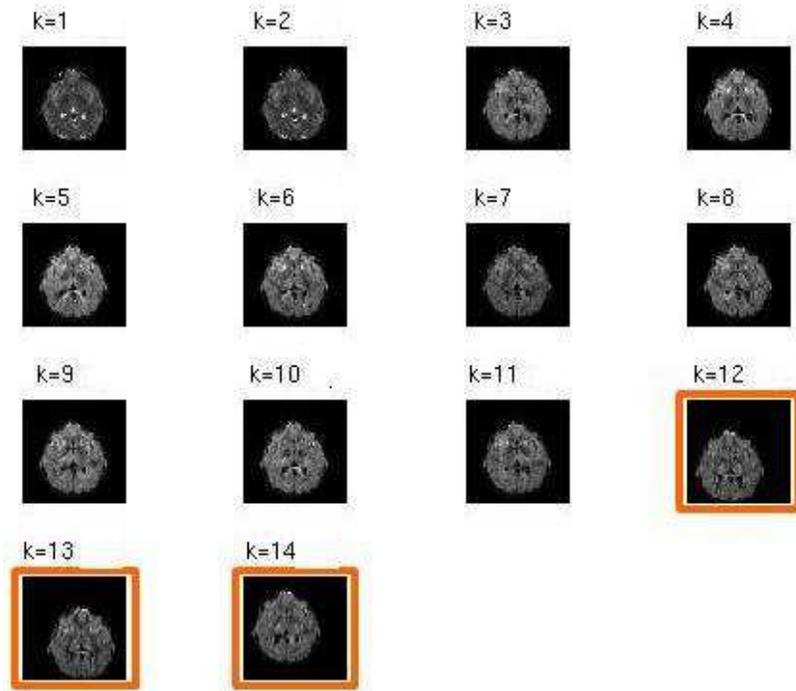


Figure 5.7: Synthetic Dataset I: last 3 of 14 measurements are translated by 8.5mm

each pixel is $1.7\text{mm} \times 1.7\text{mm} \times 2.3\text{mm}$, a 5-pixel displacement in xy -plane equals 8.5mm in the dataset. Synthetic Dataset II is almost the same as the Synthetic Dataset I, but with only last one of twelve DWIs corrupted by translating 1 voxel. Synthetic Dataset I is mainly used in the experiments of comparing different model-based registration methods, and Synthetic Dataset II is for outlier rejection methods.

The process for creating this synthetic dataset is illustrated in Figure 5.8. We discuss the procedure of resizing and corrupting in §5.3.2 and §5.3.3.

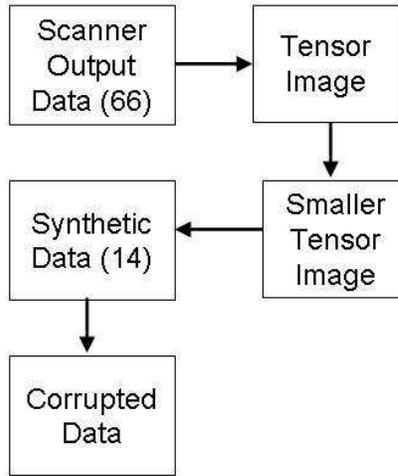


Figure 5.8: Data Simulation Flow Chart

5.3.2 Data resizing

We reduce the size of full DW dataset in two ways:

- reducing the number of voxels,
- reducing the number of measurements.

First, we use the linear least-squares algorithm to fit tensors, which is described in §2.3, and then the 3D images of the tensor elements are rescaled. We use the scaling factor of 0.5 on each dimension, and the reduced size for each 3D image is $64 \times 64 \times 30$ voxels. As exactly half number of pixels on each dimension are taken, nearest neighbour interpolation is used during rescale processing, so SNR is preserved. Even though the SNR of the individual DW images is preserved, the reduction in the total number of images causes a decrease in the SNR in derived scalar images, such as the FA.

To reduce the number of measurements, we choose a smaller set of \mathbf{q} values and use equation (2.7) to synthesize measurements from the reduced-size tensor field. The reduced-size data set has two zero-weighted and twelve measurements with non-zero \mathbf{q} . The twelve \mathbf{q} values have directions distributed evenly on the 3D unit sphere and $|\mathbf{q}|$ equal to the original dataset of “olgacontrol”. Comparing to the original, the SNRs of DW measurements and processed (eg FA) images in synthetic dataset are increased.

5.3.3 Data corruption

We introduce various translation and rotation artefacts to simulate corruptions in real datasets. Linear interpolation is used when transforming 3D images. Translation corruption is implemented simply by translating geometrical coordinate of some DWIs. Rotation is more complex. Because of the existence of diffusion gradient direction, the actual head rotation changes the DWI's contrast as well as spatial position of the image features.

To simulate a corrupted image with patient head rotation by R , we follow several steps:

Step1: Rotate diffusion gradient \mathbf{q}_i with R^{-1} .

Step2: Synthesize image $A(\mathbf{q}'_i)$ for $\mathbf{q}'_i = R^{-1}\mathbf{q}_i$, as illustrated in Figure 5.9(b).

Step3: Rotate image $A(\mathbf{q}'_i)$ by R .

Image $R(A(\mathbf{q}'_i))$ (as shown in Figure 5.9(c)) corresponds to the same gradient direction of original uncorrupted image $A(\mathbf{q}_i)$, but with the different image contrasts.

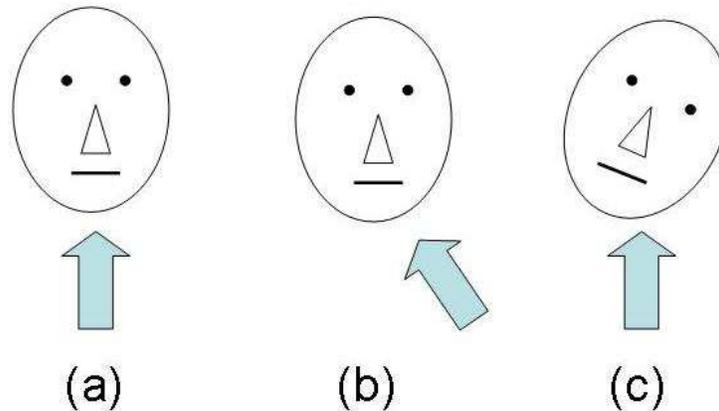


Figure 5.9: Rotation corruption simulation steps (a) head without rotation, (b) equivalent head contrast to the head with rotation (c) final corrupted image caused by head rotation

5.4 Experiments

In this section, we run experiments on “Synthetic Datasets I” and compare the Traditional Method, FMAM, FM.cons, FM.a, FM.b and FM.th.

5.4.1 Experiment

The result for matrices registering the image is denoted by the actual length unit, “mm”. Thus, ideally the transformation matrices used to register these three corrupted measurements should

be

$$T'_{12} = \begin{bmatrix} 1 & 0 & 0 & 8.5 \\ 0 & 1 & 0 & 8.5 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}, \quad (5.6)$$

$$T'_{13} = \begin{bmatrix} 1 & 0 & 0 & -8.5 \\ 0 & 1 & 0 & 8.5 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}, \quad (5.7)$$

and

$$T'_{14} = \begin{bmatrix} 1 & 0 & 0 & 8.5 \\ 0 & 1 & 0 & -8.5 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}. \quad (5.8)$$

The **Traditional Method**, which uses the zero-weighted image as the reference image for registration, should give an acceptable result. **FMAM** could make a set of reference images, but as the corrupted measurements contribute to fitted tensor D , the synthetic reference images could be badly affected. **FM.cons** fits the tensor D by using the first eight measurements, which contains two zero-weighted and six diffusion-weighted images. As all the corrupted data is outside the fixed set of measurements used to create the synthetic reference images, the result should be good. **FM.a**, **FM.b** and **FM.th** use different schemes to reject the corrupted measurements, so in experiment the last three measurements should be identified as the corrupted data. We use **normalised correlation (NC)** cost function, which is very commonly used in image registration. In Chapter 7, we will compare performances of different cost functions.

5.4.1.1 Traditional Method

The traditional method produces the following three transformation matrices to register the corrupted three measurements:

$$T_{12} = \begin{bmatrix} 1.00 & -0.03 & -0.02 & 9.67 \\ 0.03 & 1.01 & -0.01 & 7.27 \\ 0.01 & -0.00 & 1.03 & -0.21 \\ 0 & 0 & 0 & 1 \end{bmatrix}, \quad (5.9)$$

$$T_{13} = \begin{bmatrix} 1.00 & -0.03 & -0.02 & -7.10 \\ 0.027 & 1.02 & -0.01 & 6.44 \\ 0.013 & -0.00 & 1.03 & -0.33 \\ 0 & 0 & 0 & 1 \end{bmatrix}, \quad (5.10)$$

and

$$T_{14} = \begin{bmatrix} 1.00 & -0.04 & -0.01 & 10.59 \\ 0.03 & 1.01 & -0.02 & -9.79 \\ 0.01 & -0.00 & 1.03 & 0.25 \\ 0 & 0 & 0 & 1 \end{bmatrix}, \quad (5.11)$$

which are in the unit of mm. In the matrices $T_{12} \dots T_{14}$, elements $[1, 4]$ and $[2, 4]$ correspond to correction to the induced translation corruptions, so they should be around the absolute value of 8.5 in theory, but the transforms contain slightly higher values. The image result is shown in Figure 5.10. Comparing with dataset before registration (shown in Figure 5.7), we can see that the last three corrupted measurements are better aligned to the others than in the original corrupted dataset, but it is clear from the matrices above that realignment is not perfect.

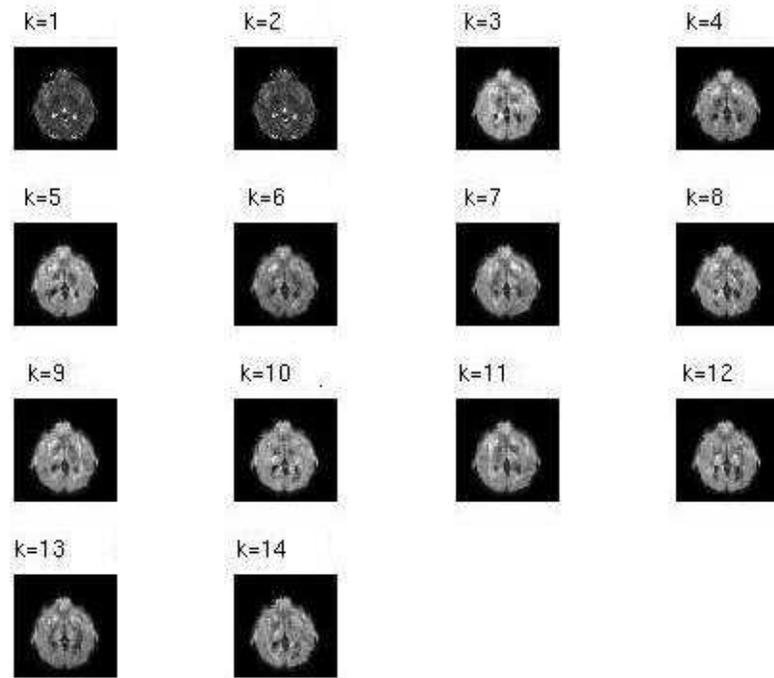


Figure 5.10: Output of Traditional Method in Experiment

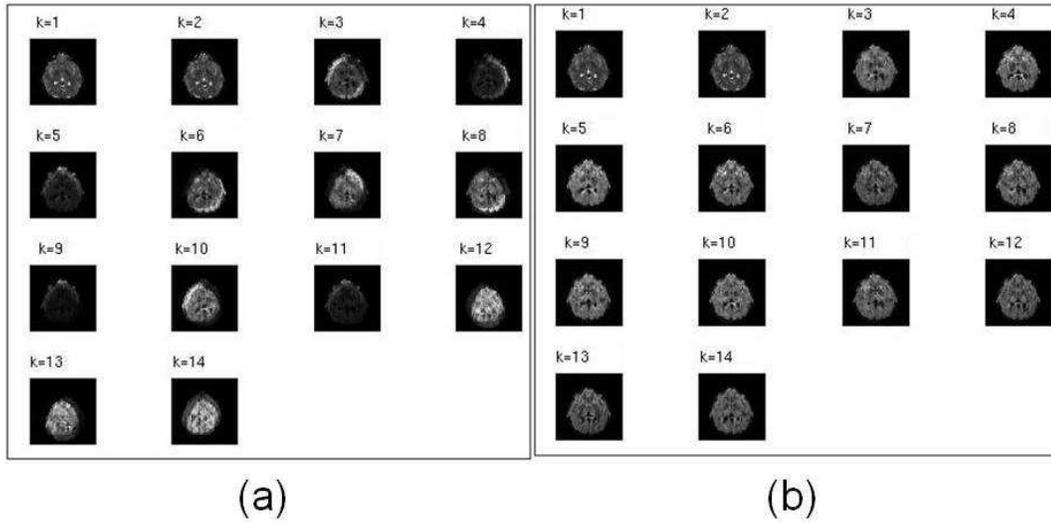


Figure 5.11: FMAM Reference Images (a) and FM.cons Reference Images (b)

5.4.1.2 FMAM

In the experiment with FMAM, as the corrupted measurements contribute to the tensor fitting, the reference images (Figure 5.11(a)) made from D show clear corruption. Transformation matrices obtained are:

$$T_{12} = \begin{bmatrix} 0.79 & -0.04 & 0.10 & 13.35 \\ 0.04 & 0.84 & 0.11 & 4.34 \\ -0.09 & -0.03 & 0.93 & 8.49 \\ 0 & 0 & 0 & 1 \end{bmatrix}, \quad (5.12)$$

$$T_{13} = \begin{bmatrix} 0.76 & 0.10 & -0.03 & 3.55 \\ -0.05 & 0.88 & 0.05 & 9.72 \\ 0.00 & -0.00 & 0.88 & 4.35 \\ 0 & 0 & 0 & 1 \end{bmatrix}, \quad (5.13)$$

and

$$T_{14} = \begin{bmatrix} 0.84 & 0.15 & 0.043 & 2.42 \\ -0.09 & 0.82 & 0.02 & 4.23 \\ -0.07 & 0.02 & 0.91 & 5.84 \\ 0 & 0 & 0 & 1 \end{bmatrix}. \quad (5.14)$$

The reference images produced on FMAM are so different from the corrupted source images that matrices in FMAM method (equation (5.12)–(5.14)) are not very close to the standard matrices T' (equation (5.6)–(5.8)).

5.4.1.3 FM.cons

For FM.cons, the last three corrupted measurements do not affect the tensor fitting in Step 1. Thus each reference image (Figure 5.11(b)) matches the source image perfectly giving the registration algorithm the best opportunity to find the correct transformation. Transformation matrices are:

$$T_{12} = \begin{bmatrix} 1.00 & 0.00 & 0.00 & 8.50 \\ 0.00 & 1.00 & -0.00 & 8.50 \\ -0.00 & 0.00 & 1.00 & -0.00 \\ 0 & 0 & 0 & 1 \end{bmatrix}, \quad (5.15)$$

$$T_{13} = \begin{bmatrix} 1.00 & -0.00 & -0.00 & -8.50 \\ 0.00 & 1.00 & -0.00 & 8.49 \\ 0.00 & 0.00 & 1.00 & -0.025 \\ 0 & 0 & 0 & 1 \end{bmatrix}, \quad (5.16)$$

and

$$T_{14} = \begin{bmatrix} 1.00 & -0.00 & 0.00 & 8.51 \\ 0.00 & 1.00 & -0.00 & -8.50 \\ 0.00 & 0.00 & 1.00 & -0.02 \\ 0 & 0 & 0 & 1 \end{bmatrix}. \quad (5.17)$$

Compared to matrices in traditional method (equation (5.9)–(5.11)), matrices in FM.cons method (equation (5.15)–(5.17)) are much closer to the standard matrices T' (equation (5.6)–(5.8)). This because the reference images made in FM.cons method are more similar to the corrupted images, than the non-diffusion-weighted images used in traditional method.

FM.cons used the first six DWIs, so the last three corrupted DWIs did not affect the tensor fitting. But if some of the first six DWIs were corrupted, FM.cons would fail in a similar way to FMAM. A further drawback of FM.cons that the simple experiment above does not reveal is that it never exploits potentially good information provided by the DWIs after the first six even if they are well aligned from the outset.

5.4.1.4 FM.a

FM.a selects corrupted measurements by comparing the whole set of corrupt measurement images with the synthetic reference (as shown in Figure 5.12(a)), which is fitted from the whole dataset. Figure 5.12(b) plots image differences S_m (using equation (5.1)) to measurement number, from which we cannot identify the last three corrupted measurements. FM.a fails because

the corrupted measurements contribute to the reference image making, which is the same as FMAM.

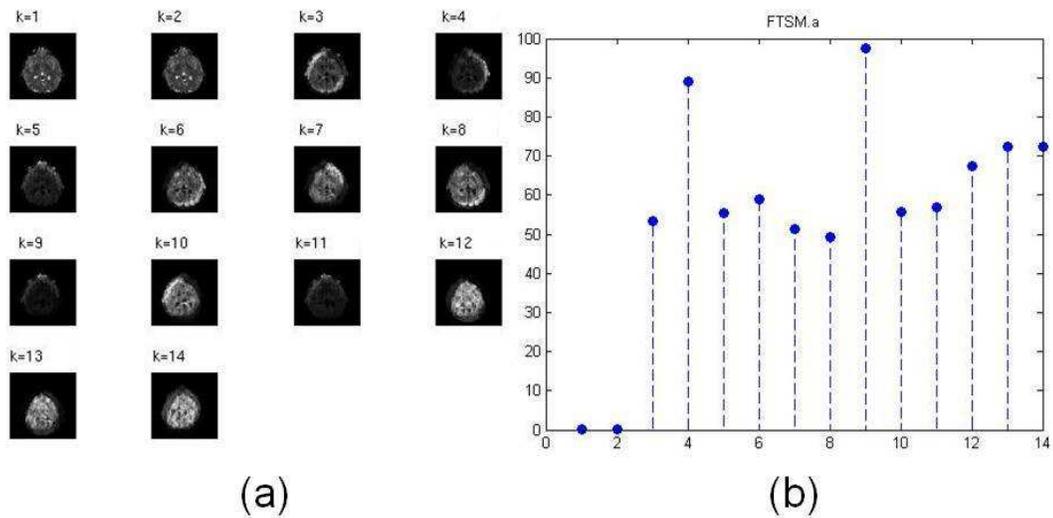


Figure 5.12: (a) Reference images used to identify the corrupted measurements in FM.a method; (b) Difference S_m between corrupted measurements and reference images in FM.a method

5.4.1.5 FM.b

FM.b method makes the reference images to identify the corrupted measurements in a slightly different way to FM.a, but as shown in Figure 5.13, FM.b still cannot identify the last three corrupted measurements. In FM.b, a 3D measurement won't contribute to the reference of its own, but its reference's synthesis cannot exclude the effect from other corrupted measurements. In other words, one corrupted measurement affects reference of all the others.

5.4.1.6 FM.th

FM.th identifies the corrupted measurements from S_m in equation (5.2), which is the difference between the thresholded measurements (Figure 5.14(a)) and the mean thresholded zero-weighted image (top right corner of Figure 5.14(b)). As shown by the plot in Figure 5.14(b), FM.th method clearly identifies that the last three measurements as the most corrupted. Thus, after deciding a proper threshold, FM.th can reject the last three corrupted measurements and fits tensor using all the other eleven least corrupted measurements. FM.th uses more DWIs than FM.cons for making the reference images, so the individual synthetic reference match its source image even more. Thus, the transformation matrices used to register the corrupted three

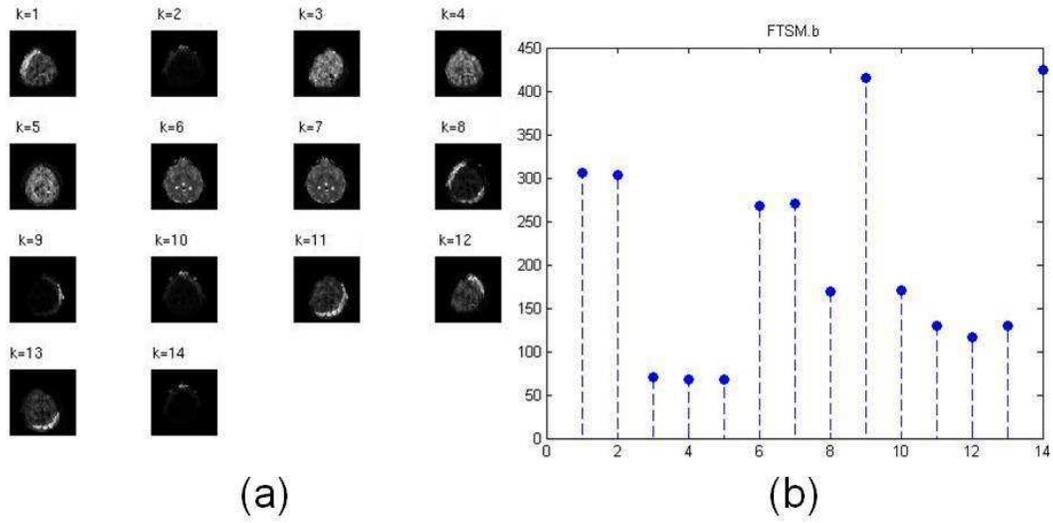


Figure 5.13: (a) Reference images used to identify the corrupted measurements in FM.b method; (b) Difference S_m between corrupted measurements and reference images in FM.b method

measurements

$$T_{12} = \begin{bmatrix} 1.00 & 0.00 & 0.00 & 8.50 \\ 0.00 & 1.00 & -0.00 & 8.50 \\ -0.00 & 0.00 & 1.00 & -0.02 \\ 0 & 0 & 0 & 1 \end{bmatrix}, \quad (5.18)$$

$$T_{13} = \begin{bmatrix} 1.00 & -0.00 & -0.00 & -8.50 \\ 0.00 & 1.00 & -0.00 & 8.49 \\ 0.00 & 0.00 & 1.00 & -0.03 \\ 0 & 0 & 0 & 1 \end{bmatrix}, \quad (5.19)$$

and

$$T_{14} = \begin{bmatrix} 1.00 & -0.00 & 0.00 & 8.51 \\ 0.00 & 1.00 & -0.00 & -8.50 \\ 0.00 & 0.00 & 1.00 & -0.02 \\ 0 & 0 & 0 & 1 \end{bmatrix} \quad (5.20)$$

are even closer to standard matrices T' (equation (5.6)–(5.8)) than matrices in FM.cons method.

5.4.2 Summary

In this Chapter, we discuss and develop a set of model-based registration methods to correct motion between acquisitions in diffusion MR images. FMAM is the primary method; FM.cons

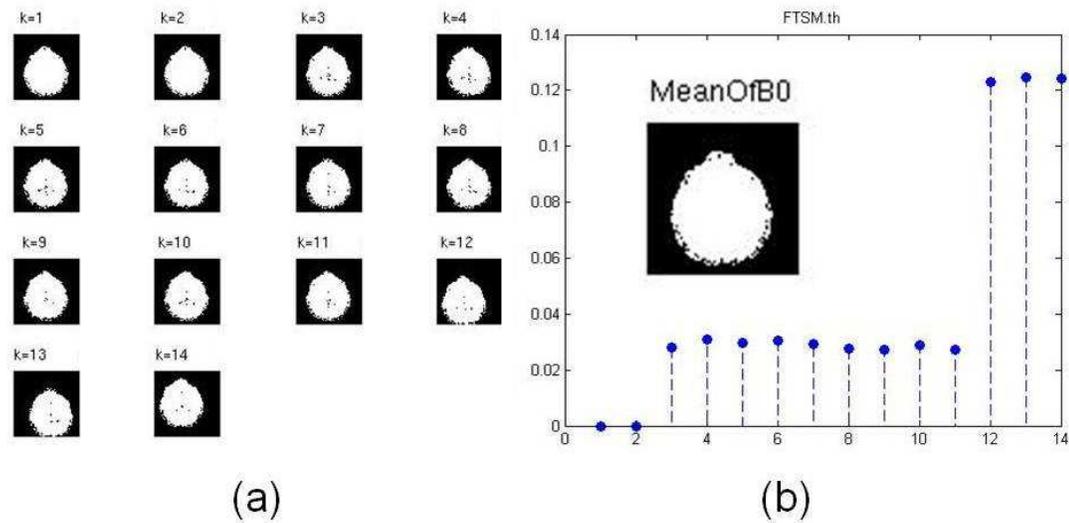


Figure 5.14: (a) Thresholded corrupted images; (b) Mean thresholded zero-weighted image, and the plot of differences S_m between thresholded corrupted images and mean thresholded zero-weighted image in FM.th method

FM.a FM.b and FM.th try to reject corrupted 3D volume measurements, and create synthetic reference by using part of the dataset. Experiments show that they either fail to work (e.g. FM.a and FM.b) or are not very robust (FM.cons). FM.th works best of all, but we still have not found a really efficient way to decide the threshold.

The key conclusion from this chapter is that the model-based approach works well if we can avoid the effects of corrupted on the references. Results from FM.cons and FM.th in this chapter show that in that case it outperforms the tradition approach. However, both FM.cons and FM.th have limitations and the following chapters explore ways to avoid corrupted measurements more robustly.

Chapter 6

Outlier Rejection Schemes

In Chapter 5, we discuss and develop several methods to reject corrupted 3D volume measurements, but experiments show they either fail to work (e.g. FM.a and FM.b) or are not very robust (FM.cons). The key problem with those methods is that they fail to identify misaligned images, which therefore contribute to the model fitting and thus the reference images. This observation motivates us to find better ways to exclude misaligned images from the model fitting. One way to proceed is to use robust estimation.

This chapter looks at two methods to identify outliers in the set of measurements in each voxel and thus fit the tensor more robustly. The first method is RESTORE, which comes from the diffusion MRI literature. The second is RANSAC, which is a computer vision technique. The application of RANSAC to DTI is a novel contribution of this thesis.

We first give an introduction to the two outlier rejection schemes; then, we propose some experimental hypotheses. In the experiment section, after introducing an evaluation method based on the principal direction of diffusion tensor, we discuss the affect of different parameters of RANSAC, and compare RANSAC with RESTORE based on different tolerance settings. Also, we compare the outlier rejection schemes with linear tensor fitting. At the end of this chapter, we draw a brief conclusion.

6.1 Two outlier rejection schemes

In this section, we introduce the two outlier rejection schemes, RESTORE and RANSAC, and also give a general introduction to outlier and the estimation of standard deviation σ .

6.1.1 Outlier

An outlier is an observation that lies outside the overall pattern of a distribution [96]. In the large sets of samples in a diffusion MRI acquisition, some measure points are further away from the position or the grey value it should be due to effects beyond simple measurement noise. This can be due to an incidental systematic error from the scanner, or motion of the subject or patient. Diffusion model fitting methods, such as the least-squares regression model [21], assume that the signal variability is only affected by thermal noise. Thus, outliers can cause major corruptions to features based on the fitted tensor.

6.1.2 RESTORE

Chang et al. [32] propose an approach to identify and exclude potential outliers in DWI, called RESTORE (Robust ESTimation of Tensors by Outlier Rejection). It uses an iteratively re-weighted least-square (LS) regression to identify potential outliers, which are then excluded. The final fit is performed with the remaining measurements.

Three diffusion tensor fitting approaches are related to RESTORE. They are linear least-squares fitting to the logarithmically transformed signal (linear LS) [21], nonlinear least-squares fitting to the untransformed signal (nonlinear LS) [71] and nonlinear LS fitting with robust Geman-McClure M-estimator (GMM) [51, 88]. All three methods minimise the value of

$$\chi^2 = \sum_i \omega_i \times (y_i - y(x_i))^2, \quad (6.1)$$

where y_i is the experimental value of the i th data point, the independent variable x_i is the \mathbf{b} -matrix for that data point, $y_i(x_i)$ is the corresponding fitted value, and ω_i is a weighting factor. The three tensor fitting approaches use $\frac{(S(\mathbf{b}))^2}{\sigma^2}$, $\frac{1}{\sigma^2}$ and $\frac{1}{r_i^2 + C^2}$ respectively for ω_i , where $S(\mathbf{b})$ is the signal intensity corresponding to \mathbf{b} -matrix \mathbf{b} , r_i is the residual between y_i and $y(x_i)$, and C is the scale factor. The way of estimating C is explained in [32] in detail.

The RESTORE algorithm includes several steps:

1. Initialise parameters using linear LS fitting;
2. Compute the DT using nonlinear LS method with constant weight, $\omega_i = \frac{1}{\sigma^2}$;
3. If the residuals of all data points lie within a given interval T , which means no outliers, the goodness-of-fit criterion is satisfied and the diffusion tensor from 2 is accepted;

Otherwise, an iterative re-weighting process using the GMM weighting function $\frac{1}{r_i^2 + C^2}$ is initiated. The reweighting process continues until it satisfies a convergence criterion. When the iteration is finished, points lying outside T are regarded as outliers and excluded. The final DT is recomputed using the nonlinear LS method.

The confidence interval, T , is set to three times of the estimated noise level σ .

From the steps above, we can see that no matter whether there are outliers rejected or not, the final DT from RESTORE is always computed using the nonlinear LS method.

6.1.3 Estimation of sigma

The RESTORE algorithm requires an estimate σ , the noise standard deviation. We assume the image measurement M is the magnitude of a real signal S plus complex Gaussian noise η [62], so that $M = |S + \eta|$. In background $S = 0$, so $E(M^2) = 2\sigma^2$, where E denotes expectation over an ROI. Thus, an estimate for the value of σ is $\sqrt{E(M^2)}/\sqrt{2}$ from an ROI entirely in background. Alternatively, Henkelman [62] just used $\sigma = 1.5267 \times$ (standard deviation of the background noise) to estimate the signal variance σ^2 .

In fact, the RESTORE algorithm contains a hidden parameter since the confidence interval T for outlier rejection is 3σ . For experiment, using different values of σ in the algorithm is equivalent to varying the width of T . Briefly speaking, bigger T allows the program to involve more measurements in the DT fitting; small T leads to more rejections.

6.1.4 RANSAC

The RANSAC (Random Sample Consensus) algorithm [46] is an algorithm for robust fitting of models in the presence of many data outliers. It has been successfully used for a variety of model fitting problems in computer vision [48]. In this section, we adapt it for fitting diffusion models to diffusion MRI data; specifically fitting the diffusion tensor. A basic assumption is that the data consists of “inliers”, which is a subset of data whose distribution can be explained by the model, and “outliers” which are data that do not fit the model.

Assume that we have N DW measurements and want to fit a n parameter diffusion model. For example, $n = 6$ for the DT model. The RANSAC algorithm uses random sampling to

search for a subset of the measurements that provides a fitted model that provides a close estimation to a high percentage of the other measurements. The algorithm proceeds as follows:

1. Pick n DW measurements from N at random;
2. Fit the model;
3. For all remaining $(N - n)$ measurements, compute error from model;
4. Divide the DW measurements into inliers and outliers according to a confidence interval T on the residual error;
5. Count fraction of inliers;
6. If fraction is greater than or equal to a fixed threshold f , then accept fit;
Otherwise, go to step 1 and repeat the whole procedure.

There are a few adjustable parameters in the RANSAC algorithm, which affect its model fitting performance. They include the size of the sample subset n in step 1, the confidence interval T for determining whether a data point fits the model, and fraction f indicating the number of inliers required to accept the model. The final fitted model from RANSAC is from a DW set containing both the sample set that provided the successful model and all other inlying measurements. For coding, we also need to decide the maximum iterations k allowed in the algorithm. After k attempts, if a good model has still not been found, RANSAC will simply use all the measurements to fit the model without any outlier rejection and output a warning.

Different diffusion models can be used in RANSAC, which can be either linear or non-linear. As the purpose is to identify outlier in a relative short time, we use linear LS for DT fit here. Using the DT model in RANSAC algorithm, for instance, the size of guessed model n cannot be less than 6.

6.2 Experiments

In this section, first we introduce an evaluation method for the model fitting routines based on the principal direction (PD) of DT. Second we discuss the affect of different parameters of RANSAC. Third, we compare RANSAC with RESTORE based on different tolerance settings. Finally, we compare the outlier rejection schemes with linear tensor fitting.

6.2.1 Orientation evaluation with ground truth

To evaluate tensor direction after registration, we use the metric

$$PD = \sum (\sqrt{f_1 f_2} \cos^{-1} |\mathbf{v}_1 \cdot \mathbf{v}_2|), \quad (6.2)$$

where \mathbf{v}_i is principal direction of the DT in image i , f_i is the ratio of the two largest eigenvalues of each DT and the sum is over all brain voxels. One of the two images contains ground truth directions and the other is the image to be evaluated. The absolute value of $\mathbf{v}_1 \cdot \mathbf{v}_2$ makes the inverse cosine give an angle between zero and $\pi/2$ and $\sqrt{f_1 f_2}$ weight down the contribution of more isotropic DTs, since their eigenvectors are less well defined.

For evaluating experiments on synthetic datasets containing various corruptions, we have the ground truth DWIs, which is the dataset before synthetic corruption. To compare the dataset after registration with its ground truth, root mean square error (RMSE) of DWIs, or RMSE of FAs from fitted DTs, is not be enough to assess all the tensor's information. In particular, such measures are insensitive to errors in principal directions. Equation (6.2) captures the differences between the principal directions and provides a numerical measure to evaluate a DWI dataset after registration by comparison with its ground truth. Similarly, Alexander et al.[2] use

$$PD = \frac{\sum (\sqrt{f_1 f_2} \cos^{-1} |\mathbf{v}_1 \cdot \mathbf{v}_2|)}{\sum \sqrt{f_1 f_2}} \quad (6.3)$$

to assess their reorientation strategies, with an extra normalising term $\sum \sqrt{f_1 f_2}$. Comparing with Equation (6.2), the value of PD in Equation (6.3) makes the comparison between datasets possible. In this thesis, all discussions on PD are to compare performance of different methods to one dataset (or datasets corrupted by the same original one), so we use Equation (6.2) for experiments.

6.2.2 Experiments with RANSAC parameters

In this section, we discuss the effect of parameters k , f and n in the RANSAC algorithm. We fix 2σ for the confidence interval T in this section, which should include 95% of signals for most measurements that have high enough intensity for the Rician distribution to approximate the Gaussian. Further discussion of T is in §6.2.3.

To test the affects of different parameter values, we run experiments on a Synthetic Dataset II. The size for each 3D measurement is $64 \times 64 \times 30$ voxels, and the skull-stripped brain region

covers about 24000(20%) voxels. The noise standard deviation σ is in the range 30 to 40.

To determine the quality of the tensor fitting using RANSAC, one quantity of interest is the number of voxels identified as outliers. For the same tensor fit, a good DW measurement should have few outliers, but a corrupted measurement should have many more. We also use RMSE of FAs and PD method mentioned in §6.2.1, to compare the fitted DT with its ground truth tensor, which is fitted from dataset before synthetic corruption.

6.2.2.1 Maximum number of iterations (k)

Unlike f and n , k is not a standard parameter of the RANSAC algorithm, since it is for pure coding purpose. However a good setting for k is important, so we run experiments to study how it affects performance before the discussion of f and n .

For this dataset, the ideal fraction f is $13/14 = 93\%$. To represent more general cases of f , we should fix a value here that is not too far away from the ideal setting 93% but still with some error space. We use $f = 85\%$ and change the iteration time k . If RANSAC algorithm runs well, the number of voxels identified as outlier and rejected in the 12th DWI should be many more than in the other 11 DWIs.

As shown in Table 6.1, the 12th DWI has the most voxels rejected after RANSAC fitting. Also, when the maximum iterations k changes from 5 to 50 the outlier histogram does not change much. That means RANSAC model can find the good fit model within 5 attempts in general. Theoretically, when more corrupted measurements are involved into the dataset, RANSAC is more likely to take more iterations to stabilize the performance. Out of consideration of both time and quality consistency, we decide to keep $k = 10$ for the remaining experiments.

6.2.2.2 Fraction required to accept model (f)

After we fix the value of k , we would like to test how changes of f affect the result. Fraction f decides the minimum number of measurements required to accept that model. Low value of f increases the likelihood of accepting a poor match based on outliers, but speeds up computation and makes the algorithm robust to larger numbers of outliers; if f is too high, the algorithm will often reach the maximum number of iterations and not reject any outliers.

T	0.85				
k	3	5	10	20	50
n	6				
DWIs	Number of Outliers				
1	214	199	211	194	209
2	178	183	188	202	196
3	134	114	135	119	142
4	141	160	149	138	147
5	290	303	307	320	301
6	139	141	154	145	159
7	431	417	454	434	425
8	323	332	315	323	327
9	129	128	126	130	122
10	210	237	220	235	250
11	192	200	223	224	230
12	20420	22849	23669	23697	23679

Table 6.1: For various setting of k , outlier histograms produced by RANSAC. Measurement 12 is the corrupted image.

In the experiment shown in Table 6.2, f changes from 70% to 95%, which corresponds to the actual minimum number of DWs to accept a good model ranging from 10 to 14. When f is too low, 70% for example, the assignment of outliers is less accurate. Although the 12th DWI still has the largest outlier set, many good DWIs are regarded as corrupted data as well. When f is so high that all the DWIs have to be included, 95% in the experiment, RANSAC has to use the whole set to fit model, without any outlier rejection. The best results come from $f = 0.9$ which requires that all but one measurement fit the model within tolerance. In Table 6.2, RMSE of FAs and PD are minimum also for $f = 0.9$ giving a consistent estimate of the best setting. When the proper value of f is used, RANSAC can identify the significant outlier in the corrupted measurement, and give a better fitted model (seen from the smaller values of RMSE and PD in Table 6.2).

Accepted DWIs for Good Model	≥ 10	≥ 11	≥ 12	≥ 13	≥ 14
f	0.7	0.78	0.85	0.9	0.95
k	10				
n	6				
DWIs	Number of Outliers				
1	3108	2525	212	49	0
2	2984	2508	191	29	0
3	1041	550	135	11	0
4	1333	335	149	45	0
5	2302	903	331	87	0
6	1352	333	156	31	0
7	4140	3311	447	147	0
8	2271	896	347	110	0
9	1101	597	148	29	0
10	1427	498	210	54	0
11	1337	494	2207	46	0
12	19063	20600	23698	24206	0
RMSE of FAs	0.0204	0.0152	0.0013	0.001	0.0389
PD(1.0e+003 *)	2.8901	2.0871	0.2681	0.218	7.8775

Table 6.2: For various setting of f , outlier histograms produced by RANSAC, root mean square error (RMSE) of FAs and principal directions (PD) differences between fitted tensors using RANSAC and ground truth. Measurement 12 is the corrupted image.

6.2.2.3 Size of sample subset (n)

Using DT model in RANSAC algorithm, the size of sample subset n should not be less than 6. In the previous experiments, we fixed $n = 6$ and changed the other parameters. Now, we would like to see how n affects the RANSAC algorithm.

In the experiment shown in Table 6.3, we change the size of subset n from the minimum 6 to the maximum 12. When $n = 12$, all the twelve DW measurements are used to fit the model, and there is no measurement left to be rejected. As n increases from 6, the affect of including the corrupted measurement in the sample set will decrease because other good DWIs are more likely to cover up its bad influence. That means the outlier is less likely to be

f	0.9						
k	10						
n	6	7	8	9	10	11	12(max)
DWIs	Number of Outliers						
1	49	98	180	406	898	1107	0
2	29	77	150	354	831	1069	0
3	11	16	6	6	18	129	0
4	45	85	143	279	690	1030	0
5	87	283	606	885	1405	1460	0
6	31	84	175	342	691	1016	0
7	147	533	867	1322	1448	1629	0
8	110	316	636	972	1572	1462	0
9	29	9	7	3	19	136	0
10	54	88	141	401	1027	1033	0
11	46	79	115	341	974	1019	0
12	24206	21891	18692	14021	6891	2081	0
RMSE(FA)	0.001	0.0024	0.005	0.0115	0.0258	0.0398	0.0389
PD(1.0e+003 *)	0.218	0.5129	1.0324	2.3954	5.2365	7.5736	7.8775

Table 6.3: Outlier histograms, RMSE and PD for various setting of n in RANSAC. Measurement 12 is the corrupted image.

identified. The outlier histograms shown in Table 6.3 confirm this hypothesis and we conclude that n should remain as small as possible at 6. We also notice that the 3rd and 9th DWIs have many less voxels rejected. Plots of gradients shown in Figure 6.1 suggest the reason is that the DW gradient directions of those two DWIs are farther away from most the others, so the 3rd and 9th DWIs are more dominant in model fit, and their voxels are less likely to be identified as outliers. Overall, the outlier histogram, RMSE and PD all suggest $n = 6$ is best.

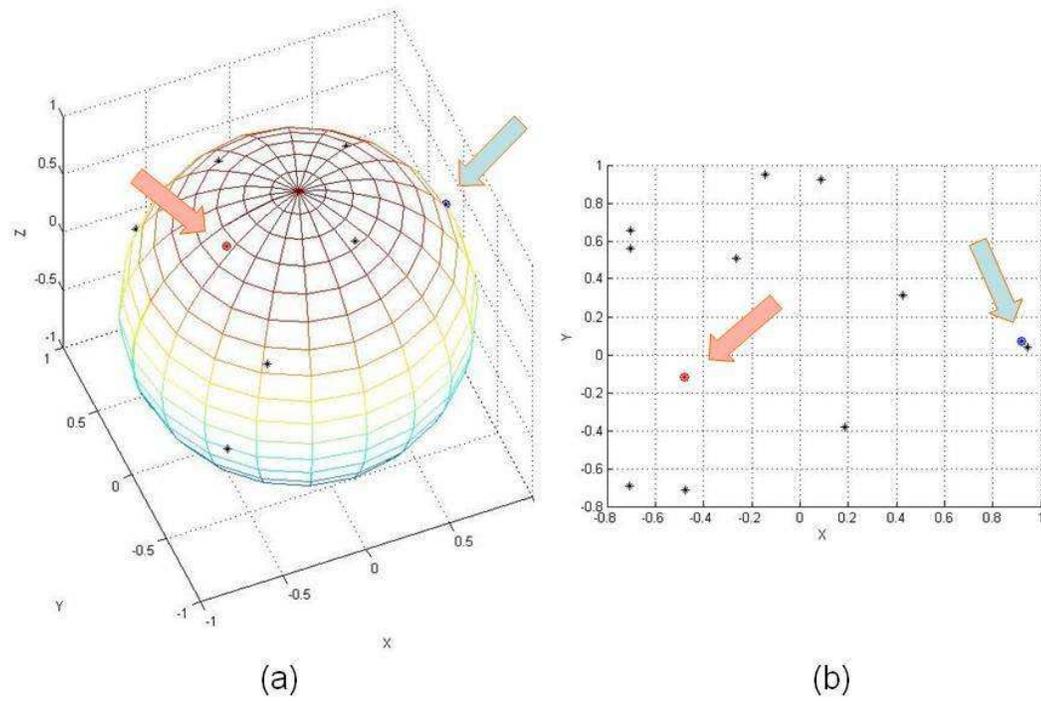


Figure 6.1: Plots of normalised diffusion-weighted gradients on xyz-space (a) and projected onto the xy-plane (b). What the red and blue arrows pointed are the gradients for the 3rd and 9th DWIs.

6.2.3 Confidence interval in RESTORE and RANSAC

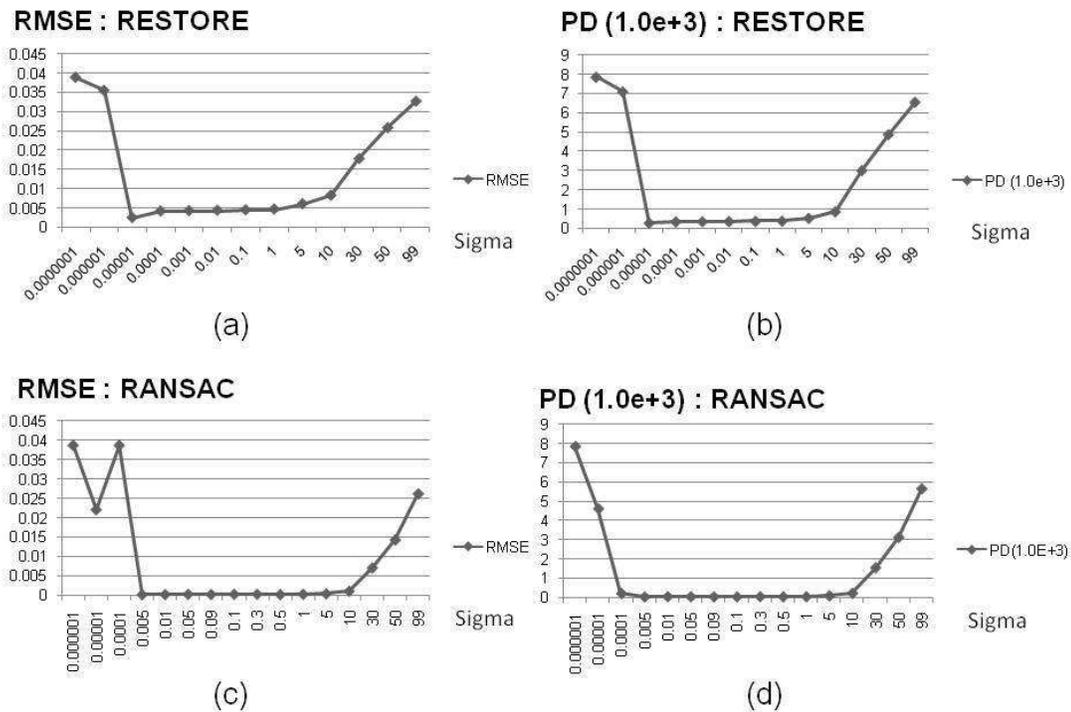


Figure 6.2: Plots of RMSE (FA) and PD against sigma (σ) for RESTORE (a)(b) and RANSAC (c)(d) methods

In §6.1.3 we described how to estimate noise standard deviation σ . In this section we investigate the effect of changing the confidence interval T , which we fixed to 2σ for RANSAC and 3σ for RESTORE by default. For experiment, using different values of σ in the algorithm is equivalent to varying the width of the confidence interval. That means, in theory high σ allows the program to involve more measurements in the DT fitting; low σ leads to more rejections.

In RESTORE, bigger σ makes larger residual errors within the confidence interval, so less DW measurements are regarded as outliers, and vice versa. The “Number of Outliers” shown in Table 6.4 and Figure 6.3(a) confirm this theoretical speculation. Figure 6.2(a) and (b) give the plots of RMSE(FA) and PD based on more detailed values of σ in $[10^{-7}, 10^2]$, and we can find that both have little change in $[0.0001, 5]$. That means in practice a good σ can be chosen within a quite big range, centred around $\sigma = 0.05$. As the estimated σ from the noisy background in this dataset is around 30 to 40, the experiment results also shows that using smaller value of σ could give better results than using the actual estimated value.

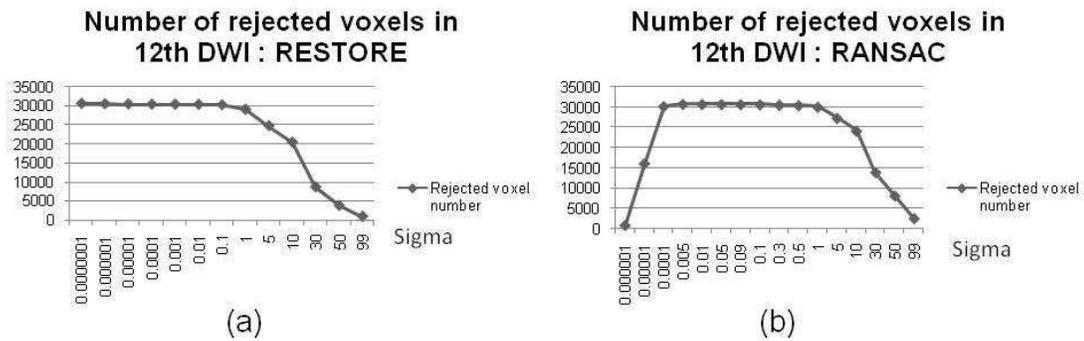


Figure 6.3: Plots of number of rejected voxels in 12th DWI against sigma (σ) for RESTORE (a) and RANSAC (b) methods

We test our RANSAC method with changing value of σ in $[10^{-6}, 10^2]$, and show comparative results in Figure 6.2(c) and (d), Figure 6.3(b) and Table 6.5. On one hand RANSAC is similar to RESTORE, and smaller σ leads more DW measurements to classified as outliers. In Figure 6.3(b), number of rejected voxels increases when σ decrease from 100 to 1. On the other hand, after k attempts, if a good model has still not been found, all the measurements will be used to fit the model without any outlier rejection. As shown in Figure 6.3(b), the number of rejected voxels stays high (around 30000) for σ in $[0.005, 1]$, but drops when σ is less than 0.005.

As shown in Figure 6.2(c) and (d), σ in $[0.0005, 5]$ gives the consistent best result. Adding check the number of outliers in Table 6.5, we would explain the reason is that the number of outliers between corrupted DWI (12th) and all the other DWIs are significant different, for σ in $[0.0005, 5]$. Similar with RESTORE, in RANSAC a good σ can be chosen within a quite big range, and using a smaller value of σ , around 0.05, gives better results than using the actual estimated value.

RESTORE								
σ	0.000001	0.00001	0.01	0.1	1	10	50	99
DWIs	Number of Outliers							
1	18363	1475	1021	1006	982	924	42	3
2	18315	1531	995	980	951	869	39	2
3	16675	799	566	538	471	341	5	0
4	19311	1316	632	630	628	590	21	0
5	22932	3934	648	645	576	231	17	0
6	19383	1519	670	670	670	642	15	0
7	23114	4077	1026	1007	966	464	25	3
8	22837	3497	614	607	513	154	9	0
9	16860	1047	582	543	447	288	4	0
10	20370	2024	554	531	427	112	0	0
11	20318	2265	585	568	474	202	5	0
12	30499	30439	30334	30225	29030	20396	3757	881
RMSE(FA)	0.0355	0.0025	0.0044	0.0045	0.0047	0.0083	0.0258	0.0327
PD(1.0e+003 *)	7.1144	0.29126	0.37529	0.3803	0.3916	0.8821	4.8903	6.5748

Table 6.4: Outlier histograms, RMSE and PD for various setting of σ in RESTORE. Measurement 12 is the corrupted image.

RANSAC								
f	0.9							
k	10							
n	6							
σ	0.000001	0.00001	0.01	0.1	1	10	50	99
DWIs	Number of Outliers							
1	0	0	0	1	6	49	97	75
2	0	0	0	0	5	29	99	90
3	0	0	0	0	0	11	74	82
4	0	0	0	0	4	45	102	98
5	0	0	0	2	9	87	226	128
6	0	0	0	0	8	31	100	117
7	0	0	0	3	13	147	410	244
8	0	0	0	3	9	110	249	186
9	0	0	0	0	6	29	86	93
10	0	0	0	0	5	54	167	113
11	0	0	0	1	2	46	124	107
12	954	16195	30803	30713	30133	24206	8215	2638
RMSE(FA)	0.0386	0.0386	0.0001185	0.0001777	0.0002	0.001	0.0142	0.0261
PD(1.0e+003 *)	7.8454	0.20052	0.0357503	0.0261	0.0324	0.218	3.1214	5.6465

Table 6.5: Outlier histograms, RMSE and PD for various setting of σ in RANSAC. Measurement 12 is the corrupted image.

6.2.4 Outlier rejection schemes and linear tensor fitting

In this section, we run experiments on ‘‘Synthetic Dataset I’’ and compare RESTORE and RANSAC to Linear DT fit.

First of all, we compare the number of outliers from RESTORE with RANSAC with $f = 90\%$ and $f = 75\%$. As illustrated in Figure 6.4, in all three methods the last three corrupted measurements are clearly distinguished from the others, and have the most rejected voxels. As we have discussed in §6.2.2.2, a good fraction f , which decides the minimum number of measurements required to accept the model, should be chosen according to the expected number of corrupted measurements in the dataset. This dataset contains 3 DWIs of the total 14 corrupted, so the suitable value of f should be smaller than $11/14(78\%)$ in order to allow 3 corrupted DW measurements to be rejected.

$f = 90\%$ makes RANSAC never find a good accepted model. From Figure 6.4(b) we can see that there are only about 5000(4%) voxels in each of the corrupted DWIs rejected, but the skull-stripped brain region covers about 24000(20%) voxels. Figure 6.4(c) illustrates the outlier histogram of RANSAC with $f = 75\%$ to fit DT. Like RESTORE, the last DWIs are clearly distinguished from the others, and their outlier numbers are even closer to the voxel number of the skull-stripped brain region, 24000, than RESTORE.

The FA maps in Figure 6.5 show that both RESTORE and RANSAC provide better tensor fitting than least-square. Comparing with FA from fitted DT using least-square linear method (Figure 6.5(b)), RESTORE improves DT fitting (Figure 6.5(c)). FA maps from RANSAC with $f = 90\%$ (Figure 6.5(d)) is just a little better than tensor directly fitted from least-square method, but very different to the ground truth (Figure 6.5(a)) and RESTORE. Comparing with RANSAC with $f = 90\%$, RANSAC with $f = 75\%$ fits DT much better and its FA is shown in Figure 6.5(e).

Table 6.6 provides numerical comparison, using RMSE of FAs and PD. RMSEs give the same expression as the FA maps: RANSAC ($f = 90\%$) does slightly better than Linear DT; RANSAC ($f = 75\%$) improves a lot from $f = 90\%$; RESTORE does slightly better and RANSAC ($f = 75\%$). But results in PD rows show that RANSAC ($f = 75\%$) preserves the DT principal directions better than RESTORE. All the calculation excludes background region.

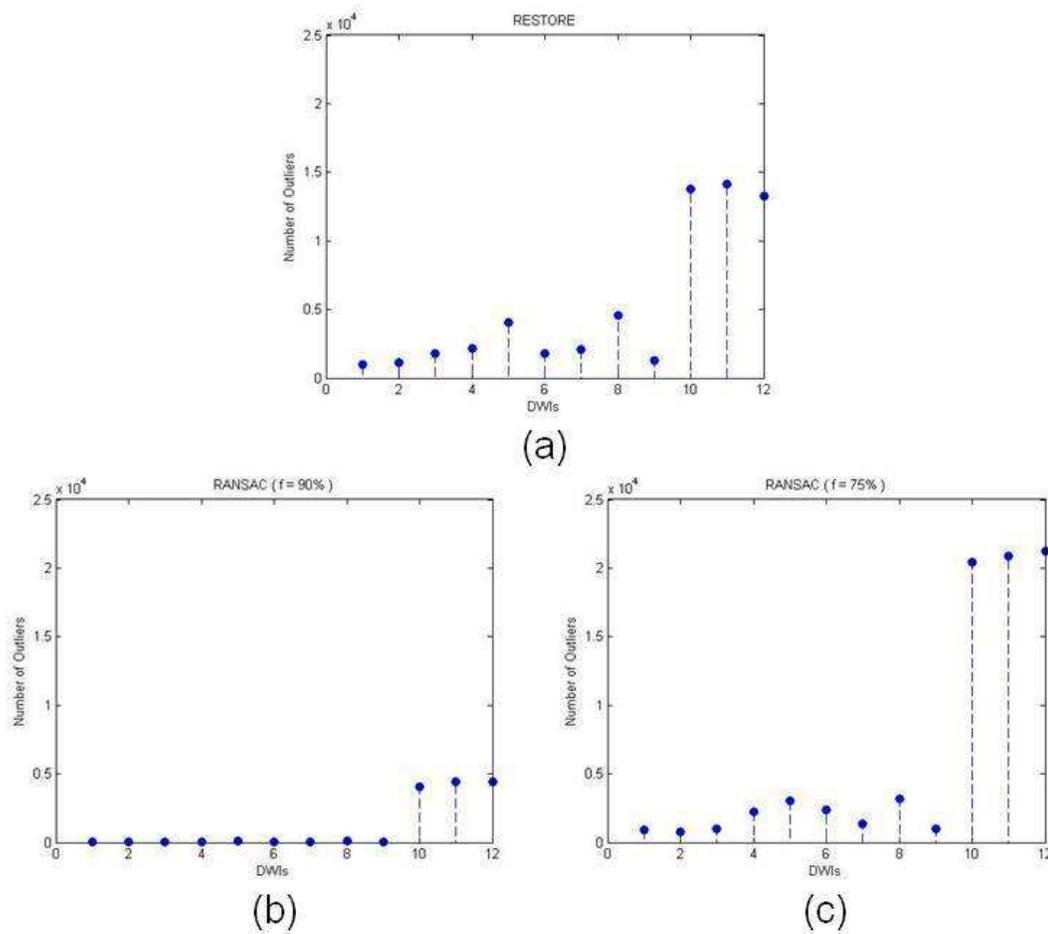


Figure 6.4: Plot of number of outliers against DWI measurements using RESTORE (a), RANSAC with $f = 90\%$ (b) and $f = 75\%$ (c)

	Linear DT Fit	RESTORE	RANSAC ($f = 90\%$)	RANSAC ($f = 75\%$)
RMSE of FAs	0.3546	0.0760	0.3363	0.1097
PD($1.0e+004$ *)	4.9675	2.6516	4.6853	1.4523

Table 6.6: Root mean square error (RMSE) of FAs and principal directions (PD) differences between fitted tensors directly using least-square linear DT fitting, RESTORE, RANSAC with $f = 90\%$ and $f = 75\%$ and ground truth.

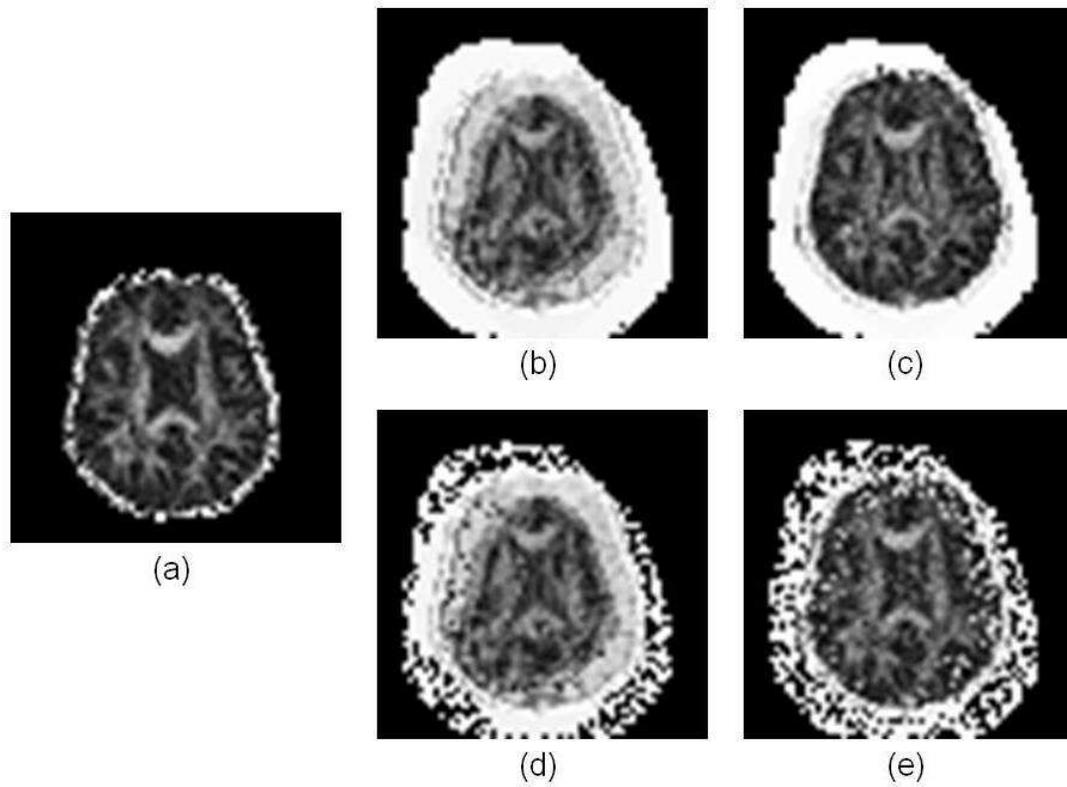


Figure 6.5: FA maps of ground truth tensor (a), tensor fitted directly using least-square linear tensor fitting (b), RESTORE (c), RANSAC with $f = 90\%$ (d) and $f = 75\%$ (e)

6.3 Conclusion

In this chapter, we first introduce two outlier rejection schemes, RESTORE and RANSAC, with a general introduction to outlier and the estimation of standard deviation σ . In experiment section, we first introduced an evaluation method based on principal direction of DT. Then in experiments of RANSAC, we started testing k (maximum number of iteration), and then tested f (fraction required to accept model) and n (size of sample subset), which are more important in RANSAC algorithm.

A clear limitation of the search for the best combination of algorithm parameters that we perform here is that we only vary each parameter individually. Thus, we cannot guarantee that we find the best overall combination. Preliminary experiments located the starting values which give reasonable performance. We do not expect to find other combinations with dramatically better performance. The main aim of the evaluation is to demonstrate the effects of each parameter rather than to locate the absolute optimal combination.

We consider only the simplest kind of transformation for these experiments, which is image translation. We could repeat the experiment using more complex transformations, such as rotations or shears, however, the purpose of experiments is to test outlier rejection for each method with different parameters' setting, which we do not expect the exact transformation to affect significantly.

We analysed the result from outlier histograms, RMSE of FAs and principal direction differences between fitted tensors using RANSAC and ground truth. All aspects of evaluation shows that the values of k and n can be fixed easily, but f needs to be decided according to the dataset individually. After that, we discussed confidence interval for both RESTORE and RANSAC. The confidence interval for both can be chosen within a big range. Although the experiment result also shows that smaller values confidence interval could give better performance, we still keep the conclusion conservative, since it is only based on one dataset. Finally, we compare the outlier rejection schemes with linear tensor fitting. Both RESTORE and RANSAC provide better tensor fitting than linear least-square. When f is chosen according to the expected number of corrupted DWIs in the dataset, RANSAC fits tensor very efficiently.

Chapter 7

Model-based Registration with Outlier Rejection

In Chapter 5, we provide a set of model-based registration methods. Experiments in Chapter 5 show the most important step in those procedures is fitting a good (DT) model, so that the synthetic reference images can be predicted well. Chapter 6 provides two outlier rejection schemes, RESTORE and RANSAC, for DT model's fitting. In this chapter, we will introduce and discuss the combination of the two outlier rejection schemes with the model-based registration method. The idea is that the robust fitting techniques reject measurements corrupted by motion or distortion while fitting the DT so reference images computed from the DT are uncorrupted. We call the combined method FMR (Fit the Model using RESTORE or RANSAC), which includes FMRestore and FMRansac.

In §7.1, we give a general introduction of FMRestore and FMRansac, with experiments similar to those with FMAM and FM in Chapter 5. In §7.2, we compare FMRs with methods introduced in Chapter 5. Then we propose two novel contributions in §7.3, which have been published in [14]: The first one is a new orientation correction technique introduced in §7.3.1, which is used after diffusion MRI registration to update diffusion gradients. In §7.3.2, we provide an evaluation method for Diffusion MRI registration and use it to compare the algorithms using human brain data.

7.1 FMRs

In this section, after introducing FMR methods briefly, we run experiments for FMRestore and FMRansac and analyse performance in terms of *transformation matrices* used to register the corrupted DWIs, in a similar way to Chapter 5.

7.1.1 FMR methods

FMR methods contain FMRestore and FMRansac. The DT model to create synthesized reference images in FMR is fitted by using RESTORE or RANSAC outlier-rejection method. Those rejections occur at the **voxel level**, ie a separate set of outliers is identified in each image voxel. The detail of RESTORE and RANSAC algorithms have been explained in Chapter 6.

7.1.2 Experiments

In this section, we run experiments on ‘‘Synthetic Dataset I’’ to compare FMRestore and FM-Ransac. As in §5.4, normalised correlation (NC) cost function is also used for registration in the experiments. We will compare performances of NC and mutual information (MI) cost functions in §7.2.

7.1.2.1 FMRestore

FMRestore identifies the corrupted measurements by RESTORE outlier-rejection method [32], and RESTORE should identify and reject outliers, and predict good reference images for FMR. FA maps from synthetic reference images used in FMRestore are shown in Figure 6.5(c). RESTORE improves DT fitting which improves the quality of the reference images in FMRestore for registration.

In Chapter 5, we compare three transformation matrices, which are used to register the corrupted three measurements, to the correct matrices T' (equation (5.6)–(5.8)) used to simulate the corruption. Here, we perform a similar comparison. The three transformation matrices that FMRestore computes are

$$T_{12} = \begin{bmatrix} 1.05 & 0.02 & -0.02 & 4.74 \\ 0.02 & 1.03 & -0.03 & 5.70 \\ -0.01 & -0.00 & 1.00 & 0.94 \\ 0 & 0 & 0 & 1 \end{bmatrix}, \quad (7.1)$$

$$T_{13} = \begin{bmatrix} 1.01 & -0.01 & 0.01 & -9.08 \\ -0.01 & 1.00 & -0.00 & 8.51 \\ 0.01 & -0.00 & 1.00 & -0.16 \\ 0 & 0 & 0 & 1 \end{bmatrix}, \quad (7.2)$$

and

$$T_{14} = \begin{bmatrix} 1.02 & -0.06 & -0.01 & 9.49 \\ 0.00 & 1.03 & 0.01 & -9.57 \\ -0.01 & 0.02 & 1.00 & -0.09 \\ 0 & 0 & 0 & 1 \end{bmatrix}. \quad (7.3)$$

Compared to the transformations in FM.cons and FM.th methods, T_{12} – T_{14} in FMRestore stay a little further away from the ideal matrices T' (equation (5.6)–(5.8)) which have absolute value of 8.5 for the two non-zero translation elements. This is because RESTORE rejects the corrupted measurements separately for voxel; while, FM.cons and FM.th reject the whole volume of identified corrupted measurements. Although FM.cons and FM.th recover the transformation matrix better in this particular example, in general we expect FMRestore to be more robust, because FM.cons and FM.th are unlikely to reject voxels from images with only minor corruption. However a potential advantage of FM.cons and FM.th is that they ignore all of corrupted images rather than only selected voxels.

7.1.2.2 FMRansac

FMRansac identifies the corrupted measurements by using RANSAC outlier-rejection method in DWI, which has been introduced in Chapter 6. Previously, we use $f = 90\%$ as default. Transformation matrices used to register the corrupted three measurements to the bad fitted references are

$$T_{12} = \begin{bmatrix} 0.81 & -0.07 & 0.08 & 14.52 \\ 0.06 & 0.85 & 0.09 & 3.40 \\ -0.06 & -0.01 & 0.94 & 5.91 \\ 0 & 0 & 0 & 1 \end{bmatrix}, \quad (7.4)$$

$$T_{12} = \begin{bmatrix} 0.78 & 0.08 & -0.03 & 2.87 \\ -0.03 & 0.88 & 0.05 & 8.03 \\ 0.00 & -0.00 & 0.89 & 4.05 \\ 0 & 0 & 0 & 1 \end{bmatrix}, \quad (7.5)$$

$$T_{12} = \begin{bmatrix} 0.85 & 0.11 & -0.01 & 5.89 \\ -0.04 & 0.84 & 0.05 & 0.08 \\ -0.01 & 0.01 & 0.91 & 3.73 \\ 0 & 0 & 0 & 1 \end{bmatrix}. \quad (7.6)$$

The absolute values of translation elements ($[1, 4]$ and $[2, 4]$) in Equations (7.4)–(7.6) are not close to 8.5, since the reference images are generated from badly fitted DT, whose FA is shown in Figure 6.5(d). The problem is that with $f = 90\%$, only one of the 14 measurements in each voxel can be rejected as an outlier, which leaves two of the three corrupted measurements.

Then we use RANSAC with $f = 75\%$ to fit DT. Figure 6.4(c) illustrates the outlier histogram. Comparing with RANSAC with $f = 90\%$, RANSAC with $f = 75\%$ fits DT much better because it can reject all three outliers. Better DT generates better reference images, and also leads better registration. Transformation matrices are

$$T_{12} = \begin{bmatrix} 1.00 & 0.00 & -0.00 & 8.38 \\ 0.00 & 1.00 & -0.00 & 8.08 \\ -0.00 & 0.00 & 0.99 & 0.42 \\ 0 & 0 & 0 & 1 \end{bmatrix}, \quad (7.7)$$

$$T_{12} = \begin{bmatrix} 0.99 & -0.00 & -0.00 & -8.12 \\ -0.00 & 0.99 & 0.00 & 8.12 \\ -0.00 & 0.00 & 0.99 & 0.43 \\ 0 & 0 & 0 & 1 \end{bmatrix}, \quad (7.8)$$

$$T_{12} = \begin{bmatrix} 0.99 & -0.00 & -0.00 & 8.65 \\ 0.00 & 0.99 & 0.00 & -8.48 \\ -0.00 & 0.00 & 0.99 & 0.31 \\ 0 & 0 & 0 & 1 \end{bmatrix}. \quad (7.9)$$

Absolute values around 8.5 dominate the translation elements in all three transformations.

7.1.3 Conclusion

We incorporate RESTORE and RANSAC outlier rejection methods into our model-based registration method, and develop FMRestore and FMRansac. From the number of outliers, reference images, as shown in §6.2.4, and recovered transformation matrices in Equations (7.7)–(7.9) used to register the corrupted DWIs, both FMR methods work reasonably, since misaligned images are often excluded for making synthetic references. The matrices recovered using FMRestore differ significantly from the true transformations but do improve alignment. FMRansac recovers the transformation matrices well and realigns the data successfully, but relies on an

appropriate choice of f for good performance.

7.2 Comparison of FMR with other methods

From Chapter 5 to §7.1, we use normalised correlation (NC) cost function to discuss types of model-based registration methods and provide only qualitative performance evaluation by inspection of the transformation matrices. This section provides a more comprehensive evaluation and comparison of the different methods. In particular, we will compare FMR with other methods, using NC and MI cost functions. Evaluation is based on transformation matrix (§7.2.1), experiments on synthetic datasets with various translated levels (§7.2.2) and full datasets (§7.2.3). Furthermore in §7.2.4, we evaluate different methods using landmark coordinate. Meanwhile, we will also continue the discussion of confidence interval and fraction f required to accept model, from RESTORE and RANSAC prior to FMRestore and FMRansac, based on large datasets.

7.2.1 Evaluation of transformation matrices

To measure the quality of transformation matrices T_{12} – T_{14} used to register the corrupted measurements from different methods, we use their root mean square error (RMSE) to the ideal matrices T' (equation (5.6)–(5.8)). In §6.2.3 we discussed the effect of confidence intervals on RESTORE and RANSAC, but the discussion is based on the “Synthetic Dataset II” which contains one measurement corrupted by a one-voxel-translation, and the conclusion was that a good confidence interval T (or saying σ for experiment) can be chosen within a big range, and using smaller values around 0.05, could give better results than using the actual estimated value. In this section, we compare the effect of a small confidence interval ($\sigma = 0.05$) to the default ($\sigma = 35$), for FMRestore and FMRansac methods, based on the “Synthetic Dataset I” used in Chapter 5 and §7.1, which has three measurements corrupted by a five-voxel-translation.

In experiments using NC cost function, Table 7.1 shows the RMSE, using formula $\frac{1}{N} \sum_{i=1}^N \sqrt{(T_i - T'_i)^2}$, between transformations used to register corrupted measurements (T) and standard matrices (T'), where N is the number of transformation matrix elements. When the RMSE is greater than 1, the reference image is *visually* badly corrupted. The synthetic reference dataset used in FMAM is generated without rejecting outliers (the FA is shown in Figure 6.5(b)), thus the badly corrupted reference leads to bad registration. Both FM.cons and FM.th methods improve the accuracy from traditional method, and work best of all, as none

	Registration Methods							
T v.s. T'	Tradition	FMAM	FM.cons	FM.th	FMRestore ($\sigma = 0.05$)	FMRestore ($\sigma = 35$)	FMRansac ($\sigma = 0.05$)	FMRansac ($\sigma = 35$)
T_{12}	0.1723	1.1475	0.0018	0.0017	0.2846	0.4806	0.0923	0.0619
T_{13}	0.2464	1.1454	0.0023	0.0029	0.2753	0.0504	0.1583	0.0768
T_{14}	0.2362	1.5922	0.0021	0.0021	0.3491	0.1440	0.1046	0.0321
Mean	0.2183	1.2950	0.0021	0.0022	0.3030	0.2250	0.1184	0.0569

Table 7.1: RMSE between transformations used to register corrupted measurements (T) and standard matrices (T') using NC cost function

	Registration Methods							
T v.s. T'	Tradition	FMAM	FM.cons	FM.th	FMRestore ($\sigma = 0.05$)	FMRestore ($\sigma = 35$)	FMRansac ($\sigma = 0.05$)	FMRansac ($\sigma = 35$)
T_{12}	0.0147	8.6607	0.0031	0.0007	0.3028	0.0952	0.0530	0.0545
T_{13}	0.0168	10.4680	0.0064	0.0047	10.6000	0.0237	0.0536	0.0848
T_{14}	0.0280	10.3860	0.0041	0.0031	0.2004	0.0190	0.0808	0.0865
Mean	0.0198	9.8385	0.0045	0.0028	3.7009	0.0460	0.0625	0.0753

Table 7.2: RMSE between transformations used to register corrupted measurements (T) and standard matrices (T') using MI cost function

of the corrupted 3D DWIs contributes to generating reference dataset. For FMRestore and FMRansac, small value of σ does not provide better result, unlike the experiment in §6.2.3. The result suggests that smaller value of σ might increase accuracy to datasets with little corruption, but make RESTORE and RANSAC less robust to the occurrence of larger numbers of corrupted measurements.

We also repeat all the experiments using MI cost function, and the result table is shown in Table 7.2. Comparing with NC (Table 7.1), MI improves performance in general. But there are some special cases, the FMAM column and “ T_{13} ” in FMRestore ($\sigma = 0.05$). That means *NC works better than MI for some corrupted reference images*, such as the 13th reference used in FMRestore ($\sigma = 0.05$) shown in Figure 7.1.

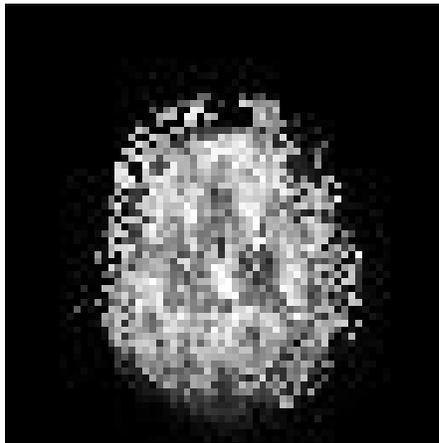


Figure 7.1: Synthetic reference image used in FMRestore with $\sigma = 0.05$ from dataset with 3 of 12 DWIs translated by 5 pixels

7.2.2 Experiments with translated synthetic datasets

From the previous experiments based on “Synthetic Dataset I” and “Synthetic Dataset II”, we get some conclusions, and they become assertions that we test in this section’s experiments:

1. In RANSAC and FMRansac, a good fraction f , which decides the minimum number of measurements required to accept the model, is best chosen according to the expected number of corrupted measurements in the dataset. If we cannot decide the precise f , lower f gives safer performance than higher f . (Assertion 1);
2. Smaller confidence interval increases accuracy for datasets with little corruption, but makes RESTORE and RANSAC less robust (Assertion 2);
3. Both FMRansac and FMRestore work better than FMAM, also better than traditional method for datasets with small corruptions (Assertion 3);
4. MI improves performance in general from NC; but NC works better than MI for some corrupted reference images (Assertion 4).

In this section, we run experiments on datasets with different levels of corruption, like “Synthetic Dataset I” and “Synthetic Dataset II”. We make a series of corrupted datasets with translations from 1.7mm to 8.5mm affecting between 1 and 5 of the 14 measurement images. The result tables include **three groups**: 1) Tables of RMSE between registered and ground truth datasets using different methods (Table 7.3–7.9); 2) Tables of RMSE of FAs between registered and ground truth datasets using different methods (Table 7.10–7.16); 3) Tables of PD differences between registered and ground truth datasets using different methods (Table 7.17–7.23).

7.2.2.1 Assertion 1

Assertion 1 is related to what we have done in §6.2.2.2, 6.2.4 and 7.1.2.2, but using datasets with different levels of corruption. Tables related to Assertion 1 are: Table 7.3 and 7.4 from Group 1, Table 7.10 and 7.11 from Group 2, and Table 7.17 and 7.18 from Group 3.

First, we can mainly focus on two tables from Group 1, Table 7.3 and 7.4. To test Assertion 1, we use the default confidence interval $T_0 = 2\sigma$ with $\sigma = 35$ and compare fractions $f = 90\%$ and $f = 75\%$. In Table 7.3, datasets with 1 of 12 DWIs corrupted (first row) have the minimum RMSE from ground truth. (We will discuss the differences between NC and MI cost functions in Assertion 4.) For RANSAC, 90% is the best setting for datasets only with 1/12 DWIs corrupted, so it fits datasets in the first row better than the rest. For the three rows corresponding to datasets with 2/12, 3/12 and 5/12 corrupted DWIs, RANSAC with $f = 90\%$ can never find 11/12 good DWIs to fit model, so it reached the maximum number of iterations and did not reject any outliers.

In Table 7.4, RANSAC works with $f = 75\%$. 75% is the best setting for datasets with 3 of 12 DWIs corrupted, but it can also find 9 DWIs to fit a good model for datasets with 1/12 and 2/12 corrupted DWIs. So it works well in the first three rows. But for the datasets with 5/12 corrupted DWIs, it failed just like $f = 90\%$ setting.

From Table 7.10 and 7.11 in Group 2 and Table 7.17 and 7.18 in Group 3 we can have similar conclusion. FMRansac ($f = 75\%$) works better on corrupted datasets in general; FMRansac ($f = 90\%$) gives the most accurate registration to datasets with 1 DWI corrupted, works badly for most of the others. That means, in practice, if we cannot decide the precise f , lower f can give safer performance than higher f .

7.2.2.2 Assertion 2

What we are doing in Assertion 2 is related to what we have done in §6.2.3 and 7.2.1. To test Assertion 2, we compare FMR methods with default confidence interval T_0 to smaller value $T' = \frac{1}{700}T_0$. T' can be regarded as 2σ with $\sigma = 0.05$. Tables related to Assertion 2 are: Table 7.4 – 7.7 from Group 1, Table 7.11 – 7.14 from Group 2, and Table 7.18 – 7.21 from Group 3.

To analyse FMRansac, we compare Table 7.4 which has been used in Assertion 1 and Table 7.5 from Group 1. Comparing the first two rows from two tables, which correspond to the

datasets with 1/12 and 2/12 DWIs corrupted, small confidence interval T' ($\sigma = 0.05$) provides lower RMSEs from ground truth. But for the 3rd and 4th row corresponding to 3/12 and 5/12 corruption, the default T_0 gives better results. In the tables of FMRestore (Table 7.7 and 7.6), we can have the same findings.

Together with the tables from Group 2 and 3, we can say that in both FMRansac and FMRestore, small confidence interval T' ($\sigma = 0.05$) works better for datasets with small corruptions; but default confidence interval T_0 ($\sigma = 35$) may be more robust for higher levels of corruption.

7.2.2.3 Assertion 3

Comparing with results from FMAM (Table 7.8) from Group 1, both FMRansac and FMRestore work better than FMAM. Tables of Traditional method (Table 7.9) shows that FMRs are better than traditional method for datasets with small corruptions, but traditional method is more robust for dataset with extreme corruption. These corruptions, however, are not likely to happen in practice and if they do, we could perform a preliminary step to correct them approximately. The same conclusion can be driven from Group 2 and 3 tables as well.

7.2.2.4 Assertion 4

In this Assertion, we are trying to compare NC and MI cost functions. From Table 7.3 to 7.23, **Bold** means MI is worse than NC. MI improves performance in general from NC; but NC works better than MI in highly corrupted datasets. It is same conclusion from what we have in §7.2.1. One badly corrupted reference is shown in Figure 7.1.

For FMRansac ($f = 90\%$), Table 7.3, 7.10 and 7.17, MI is worse than NC for datasets with one measurement corrupted. We find that NC performs better usually when both NC and MI in those three tables have very small values. A possible explanation is that the values from NC have become too small to be improved by MI, or NC is more sensitive to small corruptions than MI.

	NC				MI			
	Corrupt Shift Scale							
N	1.7(mm)	3.4(mm)	5.1(mm)	8.5(mm)	1.7(mm)	3.4(mm)	5.1(mm)	8.5(mm)
1	1.2378	1.265	1.3085	1.2416	1.5519	1.5937	1.5574	1.5785
2	11.547	17.123	23.168	32.626	10.723	13.541	16.472	29.944
3	13.321	23.392	31.361	43.637	12.285	20.555	24.602	38.604
5	16.766	28.681	36.338	50.86	16.552	27.633	38.877	55.148

Table 7.3: RMSE between registered and ground truth datasets, from FMRansac with $f = 90\%$, $\sigma = 35$, using normalized correlation (NC) and mutual information (MI) cost functions. The first column N is the number of corrupted measurements. **Bold** means MI is bigger than NC.

	NC				MI			
	Corrupt Shift Scale							
N	1.7(mm)	3.4(mm)	5.1(mm)	8.5(mm)	1.7(mm)	3.4(mm)	5.1(mm)	8.5(mm)
1	2.4561	3.0868	3.1872	3.3534	2.0544	2.4276	2.432	2.5535
2	3.272	3.9462	4.3676	4.5449	2.9357	2.9756	2.8142	2.9588
3	5.2986	8.0756	9.6819	11.324	4.4939	6.2147	7.0604	8.3988
5	13.622	24.772	32.651	46.008	12.779	22.106	27.623	44.962

Table 7.4: RMSE between registered and ground truth datasets, from FMRansac with $f = 75\%$, $\sigma = 35$.

	NC				MI			
	Corrupt Shift Scale							
N	1.7(mm)	3.4(mm)	5.1(mm)	8.5(mm)	1.7(mm)	3.4(mm)	5.1(mm)	8.5(mm)
1	2.2905	2.7442	3.0925	3.1934	1.8673	2.0564	2.1088	2.2114
2	2.0376	2.8782	3.1714	3.5451	1.802	2.2133	2.1946	2.4218
3	6.4224	9.8724	12.078	14.33	5.4903	7.4632	8.7207	27.27
5	17.087	28.882	36.393	50.932	17.057	27.817	32.927	50.116

Table 7.5: RMSE between registered and ground truth datasets, from FMRansac with $f = 75\%$, $\sigma = 0.05$.

	NC				MI			
	Corrupt Shift Scale							
N	1.7(mm)	3.4(mm)	5.1(mm)	8.5(mm)	1.7(mm)	3.4(mm)	5.1(mm)	8.5(mm)
1	2.0127	2.7322	2.9443	3.0396	1.5588	1.7067	1.5079	1.6709
2	4.3212	5.5946	6.2631	5.9937	2.6485	3.2065	2.7444	2.4579
3	8.9849	15.158	19.31	22.645	8.2683	13.682	17.19	27.187
5	16.874	28.566	35.867	49.672	16.788	28.424	34.662	63.32

Table 7.6: RMSE between registered and ground truth datasets, from FMRestore with $\sigma = 0.05$.

	NC				MI			
	Corrupt Shift Scale							
N	1.7(mm)	3.4(mm)	5.1(mm)	8.5(mm)	1.7(mm)	3.4(mm)	5.1(mm)	8.5(mm)
1	3.3183	4.769	5.0455	5.2211	2.6342	2.5994	2.5729	2.0761
2	5.6495	7.3323	7.9281	7.7599	4.4915	3.8021	3.2364	2.7185
3	6.8621	10.641	12.006	12.342	5.161	7.5776	8.593	8.8649
5	11.394	36.639	21.537	26.873	9.4318	35.573	17.811	39.081

Table 7.7: RMSE between registered and ground truth datasets, from FMRestore with $\sigma = 35$.

	NC				MI			
	Corrupt Shift Scale							
N	1.7(mm)	3.4(mm)	5.1(mm)	8.5(mm)	1.7(mm)	3.4(mm)	5.1(mm)	8.5(mm)
1	8.8856	14.884	20.002	27.765	7.867	11.158	11.183	9.7057
2	11.989	21.461	27.92	38.453	11.563	19.193	21.545	28.475
3	14.948	24.856	32.73	45.459	14.338	22.894	26.601	45.283
5	17.144	28.964	36.455	51.158	17.289	28.497	39.445	52.016

Table 7.8: RMSE between registered and ground truth datasets, from FMAM.

	NC				MI			
	Corrupt Shift Scale							
N	1.7(mm)	3.4(mm)	5.1(mm)	8.5(mm)	1.7(mm)	3.4(mm)	5.1(mm)	8.5(mm)
1	28.295	28.284	28.293	28.312	11.992	11.558	11.577	11.486
2	28.308	28.302	28.291	28.301	11.899	11.499	11.495	11.745
3	28.317	28.299	28.302	28.31	11.994	11.835	11.871	11.162
5	28.327	28.3	28.299	28.338	11.216	11.231	11.073	10.764

Table 7.9: RMSE between registered and ground truth datasets, from traditional method.

	NC				MI			
	Corrupt Shift Scale							
N	1.7(mm)	3.4(mm)	5.1(mm)	8.5(mm)	1.7(mm)	3.4(mm)	5.1(mm)	8.5(mm)
1	0.0039529	0.0043158	0.0044115	0.0041761	0.0054571	0.0053866	0.0055919	0.005289
2	0.064399	0.091452	0.12941	0.24017	0.061031	0.076726	0.098917	0.66843
3	0.070364	0.12887	0.19142	0.33357	0.065792	0.12238	0.16391	0.69194
5	0.088558	0.15875	0.22518	0.36335	0.087635	0.15397	0.62516	0.69429

Table 7.10: RMSE of FAs otherwise, as Table 7.3, from FMRansac with $f = 90\%$, $\sigma = 35$.

	NC				MI			
	Corrupt Shift Scale							
N	1.7(mm)	3.4(mm)	5.1(mm)	8.5(mm)	1.7(mm)	3.4(mm)	5.1(mm)	8.5(mm)
1	0.013646	0.017899	0.018261	0.018918	0.0087866	0.011593	0.011786	0.012455
2	0.019473	0.021832	0.024158	0.024963	0.015487	0.014406	0.013515	0.014009
3	0.028265	0.042187	0.049195	0.055898	0.023233	0.032503	0.03744	0.044047
5	0.069162	0.12874	0.18179	0.31239	0.066079	0.12197	0.16659	0.67793

Table 7.11: RMSE of FAs from FMRansac with $f = 75\%$, $\sigma = 35$.

	NC				MI			
	Corrupt Shift Scale							
N	1.7(mm)	3.4(mm)	5.1(mm)	8.5(mm)	1.7(mm)	3.4(mm)	5.1(mm)	8.5(mm)
1	0.012292	0.014922	0.017334	0.017494	0.0077081	0.0097238	0.0096178	0.010734
2	0.010187	0.014143	0.016277	0.018223	0.0067543	0.0093724	0.0099096	0.011039
3	0.034365	0.051821	0.061668	0.07075	0.0277	0.039464	0.044426	0.047965
5	0.090606	0.16032	0.22665	0.3645	0.090344	0.15644	0.21304	0.67854

Table 7.12: RMSE of FAs from FMRansac with $f = 75\%$, $\sigma = 0.05$.

	NC				MI			
	Corrupt Shift Scale							
N	1.7(mm)	3.4(mm)	5.1(mm)	8.5(mm)	1.7(mm)	3.4(mm)	5.1(mm)	8.5(mm)
1	0.010056	0.014091	0.015406	0.016052	0.0053588	0.006511	0.005166	0.0063807
2	0.024904	0.030628	0.033869	0.032814	0.014195	0.015596	0.013315	0.010758
3	0.047968	0.081084	0.1042	0.12162	0.044628	0.073201	0.093889	0.62406
5	0.088882	0.15985	0.2241	0.36986	0.089272	0.16254	0.21768	0.59127

Table 7.13: RMSE of FAs from FMRestore with $\sigma = 0.05$.

	NC				MI			
	Corrupt Shift Scale							
N	1.7(mm)	3.4(mm)	5.1(mm)	8.5(mm)	1.7(mm)	3.4(mm)	5.1(mm)	8.5(mm)
1	0.018422	0.026326	0.027559	0.028267	0.01305	0.012568	0.01293	0.0088022
2	0.033625	0.039081	0.041944	0.041067	0.02618	0.020381	0.016995	0.012721
3	0.037508	0.056902	0.063021	0.064658	0.028194	0.042984	0.048804	0.054293
5	0.056696	0.41423	0.10696	0.14899	0.048818	0.67893	0.10236	0.6791

Table 7.14: RMSE of FAs from FMRestore with $\sigma = 35$.

	NC				MI			
	Corrupt Shift Scale							
N	1.7(mm)	3.4(mm)	5.1(mm)	8.5(mm)	1.7(mm)	3.4(mm)	5.1(mm)	8.5(mm)
1	0.047876	0.080042	0.11723	0.19328	0.043133	0.067083	0.082788	0.13631
2	0.066718	0.11443	0.16819	0.27834	0.064772	0.10896	0.14284	0.62887
3	0.078533	0.13882	0.20685	0.34622	0.077901	0.13524	0.18633	0.69377
5	0.090798	0.16132	0.22671	0.36568	0.090964	0.16157	0.62489	0.68176

Table 7.15: RMSE of FAs from FMAM.

	NC				MI			
	Corrupt Shift Scale							
N	1.7(mm)	3.4(mm)	5.1(mm)	8.5(mm)	1.7(mm)	3.4(mm)	5.1(mm)	8.5(mm)
1	0.11413	0.1141	0.11413	0.11416	0.059584	0.057915	0.058092	0.057606
2	0.11415	0.11412	0.11412	0.11413	0.060103	0.058547	0.058608	0.059701
3	0.11417	0.11414	0.11416	0.11414	0.060855	0.060502	0.060383	0.057291
5	0.11422	0.11412	0.11419	0.11432	0.058344	0.05833	0.056551	0.056689

Table 7.16: RMSE of FAs from traditional method.

	NC				MI			
	Corrupt Shift Scale							
N	1.7(mm)	3.4(mm)	5.1(mm)	8.5(mm)	1.7(mm)	3.4(mm)	5.1(mm)	8.5(mm)
1	388.49	450.09	490.17	438.65	528.65	499.07	523.57	490.44
2	11785	16299	21888	37908	11065	13731	17202	75456
3	12779	21839	29716	48769	11787	20716	25537	128720
5	15581	24867	31836	46815	15389	24457	109130	111840

Table 7.17: PDs otherwise, as Table 7.3, from FMRansac with $f = 90\%$, $\sigma = 35$.

	NC				MI			
	Corrupt Shift Scale							
N	1.7(mm)	3.4(mm)	5.1(mm)	8.5(mm)	1.7(mm)	3.4(mm)	5.1(mm)	8.5(mm)
1	2058.7	2807.7	2857.2	2990.5	1203	1748.2	1808.4	1851.3
2	2886.4	3201.4	3576.7	3645.2	2324.9	1998.8	1815.5	1785
3	4586.8	7340.1	8625	9982.6	3651	5495.2	6431	7595
5	12363	21511	27667	40969	11816	20193	25321	122600

Table 7.18: PDs from FMRansac with $f = 75\%$, $\sigma = 35$.

	NC				MI			
	Corrupt Shift Scale							
N	1.7(mm)	3.4(mm)	5.1(mm)	8.5(mm)	1.7(mm)	3.4(mm)	5.1(mm)	8.5(mm)
1	1934.3	2448.3	3020.2	3009.4	853.68	1227	1266.6	1464.1
2	1286.7	1802.4	2100.8	2367.8	658.12	1080.4	1154.6	1345.1
3	5519.7	9029.5	11043	12904	4215	6469.4	7415.5	8161.4
5	15892	25026	32125	48023	15888	24843	30793	125020

Table 7.19: PDs from FMRansac with $f = 75\%$, $\sigma = 0.05$.

	NC				MI			
	Corrupt Shift Scale							
N	1.7(mm)	3.4(mm)	5.1(mm)	8.5(mm)	1.7(mm)	3.4(mm)	5.1(mm)	8.5(mm)
1	1546.6	2330.3	2608.3	2641.2	495.94	721.53	498.26	652.03
2	4397.7	5339.3	6031.8	5646.5	2059.5	2162.9	1645.5	1307
3	8482.9	14896	18732	21466	7750	13638	17272	104780
5	15658	25368	31904	47182	15753	25261	31170	43076

Table 7.20: PDs from FMRestore with $\sigma = 0.05$.

	NC				MI			
	Corrupt Shift Scale							
N	1.7(mm)	3.4(mm)	5.1(mm)	8.5(mm)	1.7(mm)	3.4(mm)	5.1(mm)	8.5(mm)
1	3119.6	4896.8	5151.5	5270.8	1978.1	1863.7	1743.3	1003.3
2	6180	7892.9	8413.8	8249.2	4636.4	3245.6	2495.2	1603.1
3	6835.3	11106	12178	12294	4930.9	8559.1	9743.3	10841
5	10786	69784	19121	23078	9298.9	126700	17756	131110

Table 7.21: PDs from FMRestore with $\sigma = 35$.

	NC				MI			
	Corrupt Shift Scale							
<i>N</i>	1.7(mm)	3.4(mm)	5.1(mm)	8.5(mm)	1.7(mm)	3.4(mm)	5.1(mm)	8.5(mm)
1	8227.5	13701	19420	30855	7535.9	11647	13168	20047
2	12322	20027	27058	42458	11757	19009	22785	104890
3	14334	23392	31464	50740	14114	22581	28025	140730
5	15963	25058	32043	47176	15987	25336	109820	124990

Table 7.22: PDs from FMAM.

	NC				MI			
	Corrupt Shift Scale							
<i>N</i>	1.7(mm)	3.4(mm)	5.1(mm)	8.5(mm)	1.7(mm)	3.4(mm)	5.1(mm)	8.5(mm)
1	19425	19420	19422	19432	6884.5	6689.8	6719.4	6762.4
2	19433	19427	19420	19434	6881.9	6709.4	6750	6967.2
3	19435	19423	19424	19430	6794.1	7091.3	6962.2	6646.6
5	19431	19429	19415	19459	6424.4	6651	6649.3	6170.7

Table 7.23: PDs from traditional method.

7.2.3 Experiments on full datasets

In this section, we run experiments on full datasets, “S32”, “Fcontrol” and “Fli”, and compare the results based on visual judgement of FA maps before providing more quantitative assessment in later sections. “S32” and “Fli” have small motion corruption, and “Fcontrol” is well aligned before registration. We first run all the methods (the traditional method, and our FMAM and FMR) using both NC and MI cost functions on the “S32” dataset, and then run and compare some methods using the “Fcontrol” and “Fli” datasets, which provide some additional information.

The FA maps of aligned datasets using different methods are shown in Figure 7.2 and Figure 7.3. Figure 7.2 provides axial view of middle slices of original dataset and aligned ones using different methods. The image from traditional method using NC (Row A) is very different from all the others, which can be explained by referring to the sagittal view in Figure 7.3. Figure 7.3 shows FA maps overlaid with a hand segmented outline of the corpus callosum from the non-diffusion-weighted image before registration. The traditional method using T2 image (Top left in Figure 7.3) as target for registration and the NC cost function leads to misalignment and the difference we observe in Figure 7.2.

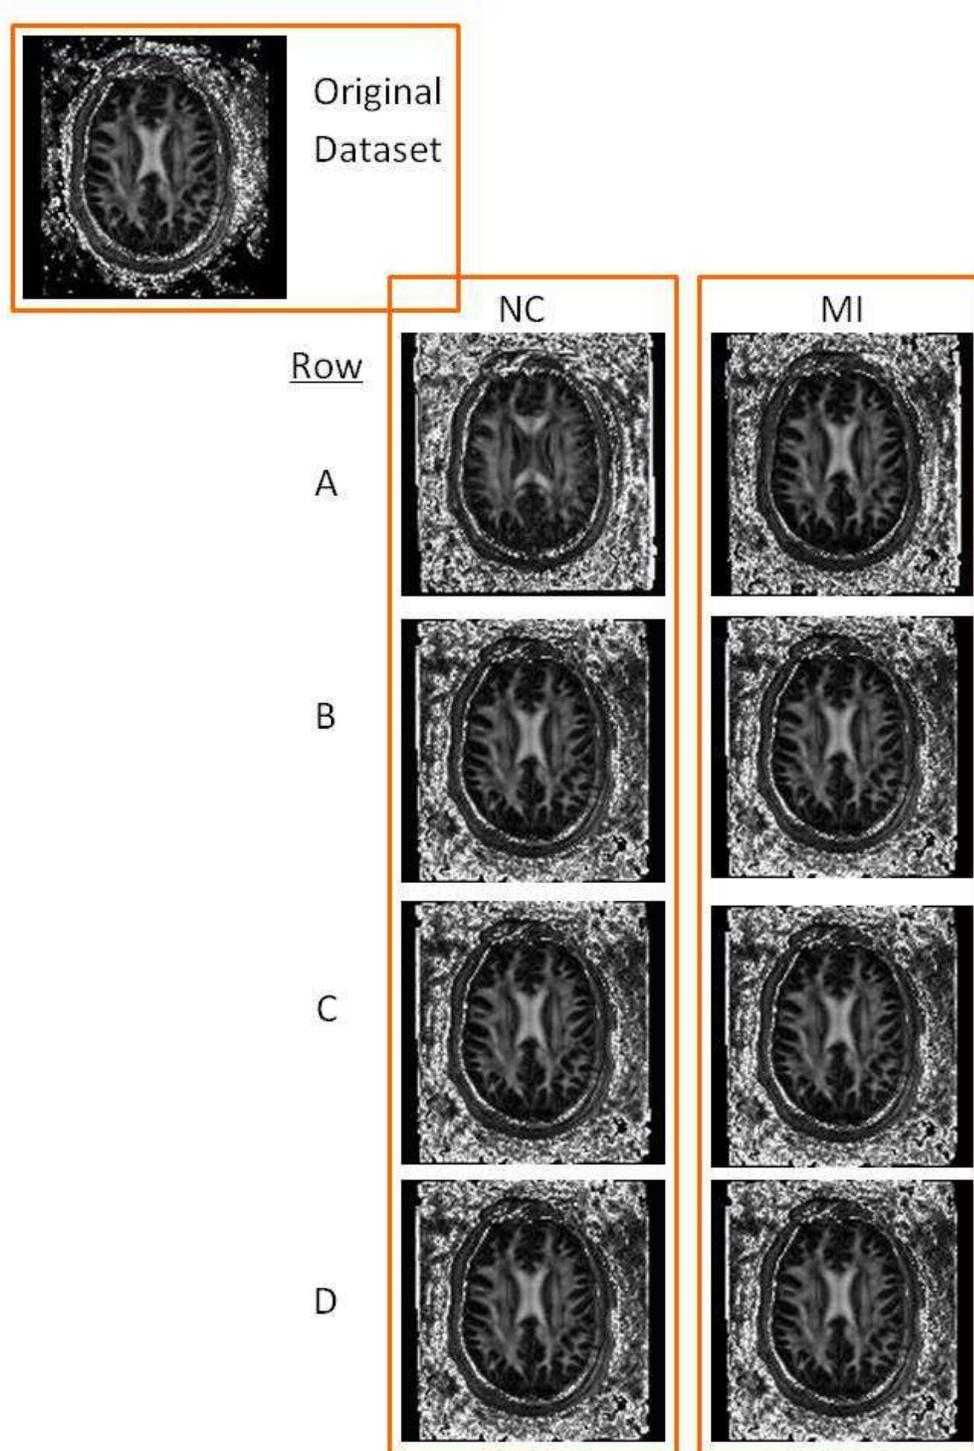


Figure 7.2: FAs from registered datasets using traditional method (Row A), FMAM (Row B), FMRestore (Row C) and FMRansac (Row D)

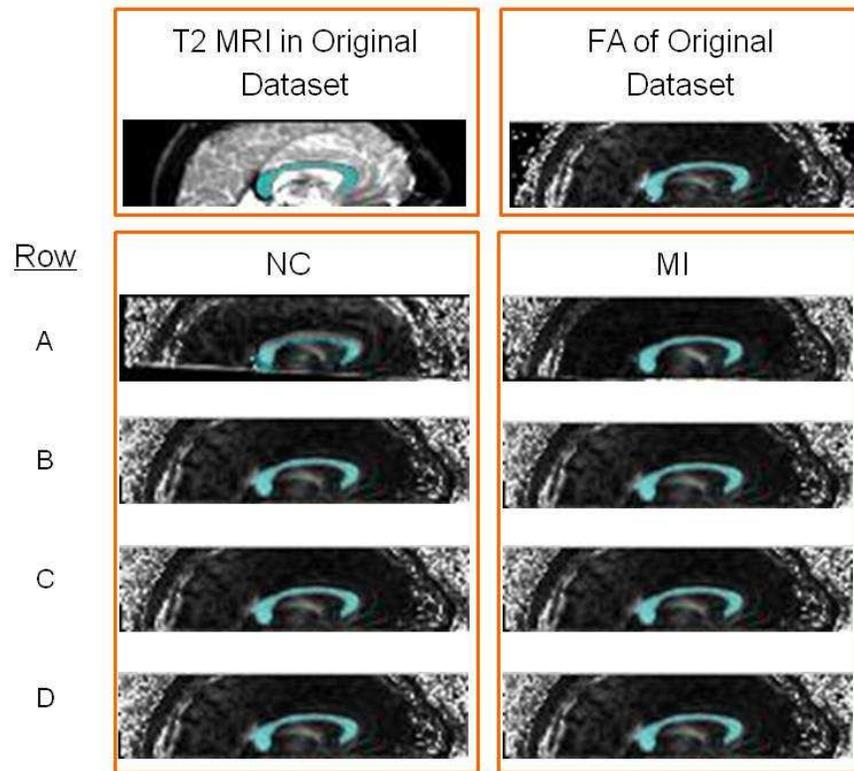


Figure 7.3: FA maps overlaid with a hand segmented outline of the corpus callosum from the non-diffusion-weighted image

We run Traditional Method, FMAM and FMRestore methods using the NC cost function on “Fcontrol” dataset, and the results are shown in Figure 7.4. Figure 7.4(a) illustrates the misalignment of a hand-defined corpus callosum outline with the FA map after the traditional method which used the NC cost function to align each diffusion-weighted image to the non-diffusion-weighted image. This is an artifact of the traditional approach that we observe frequently and arises from poor alignment of the DW images with the non-diffusion-weighted image. Figure 7.4(b) and (c) show that the problem disappears after correction with FMAM and FMRestore.

We compare traditional method and FMRestore methods using the MI cost function on the “Fli” dataset. We also overlay FA maps with a hand segmented outline of the corpus callosum from the non-diffusion-weighted image. The experiment result is shown in Figure 7.5. Still, FMRestore performs better than Traditional Method. As shown in Figure 7.5(b), not only the misalignment of corpus callosum disappears, but also the brain structure gets clearer (such as the part above corpus callosum).

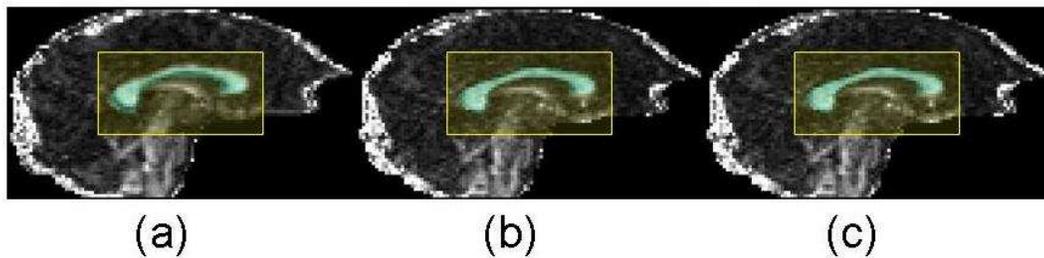


Figure 7.4: FA maps overlaid with a hand segmented outline of the corpus callosum from the non-diffusion-weighted image, after (a) standard motion correction and (b) FMAM. correction (c) FMRestore correction using NC cost function

The experiments illustrate a common problem with the traditional method. Figure 7.2, Figure 7.3, Figure 7.4 and Figure 7.5 show FA maps computed after correction using traditional method and our new methods. Overlaid on each map is a hand segmentation of the corpus callosum from the non-diffusion-weighted image. In Figure 7.3(Row A, NC Column), Figure 7.4(a) and Figure 7.5(a) the boundary of the hand segmentation does not align crisply with the edge of the apparent corpus callosum region in the FA image. The traditional method aligns boundaries of anisotropic regions poorly causing bleeding of high FA outside the true anatomic region. FMRs avoids this problem.

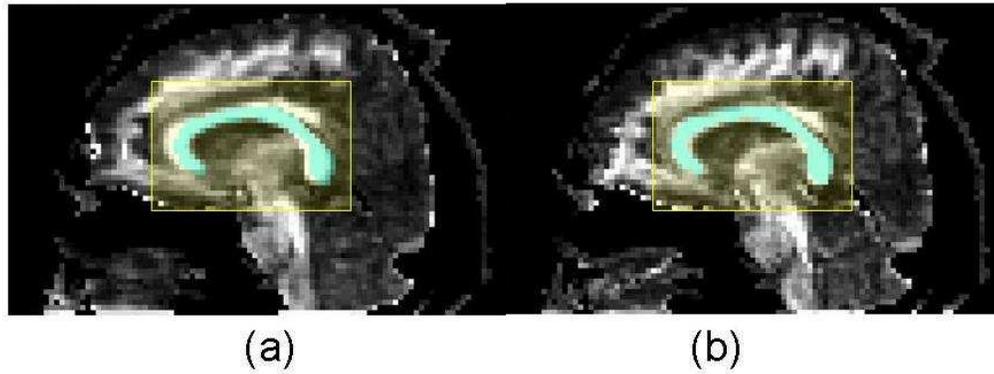


Figure 7.5: FA maps overlaid with a hand segmented outline of the corpus callosum from the non-diffusion-weighted image, after (a) the traditional method and (b) FMRestore correction using MI cost function

7.2.4 Landmark evaluation method and experiment

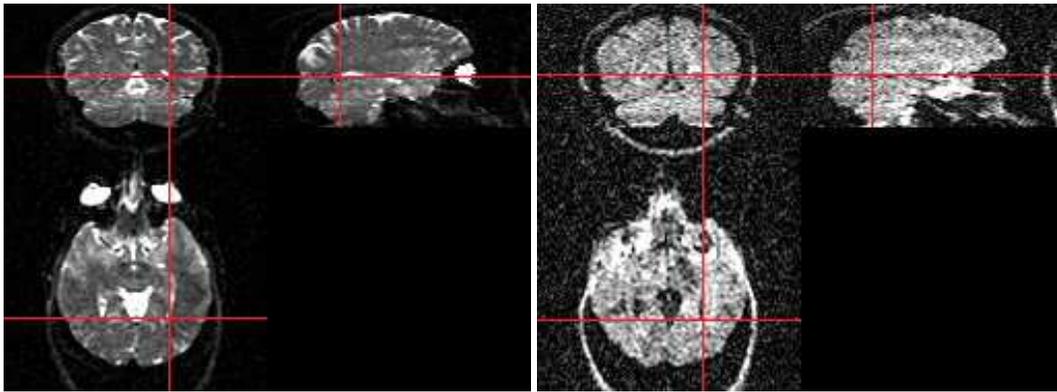
In this section, we discuss a method using the covariance of the landmarks coordinates to evaluate the alignment of all the measurements of one data set.

The landmarks are carefully chosen anatomical features that we can locate reliably in 3D (as illustrated in Figure 7.6), such as the saddle point at the top of corpus callosum or right, posterior, inferior corner of the right ventricle. For example, imagine we have a dataset having $2 b = 0$ and 9 diffusion weighted images. We label 5 corresponding landmarks on each image. After registrations using different methods, we compute transformed locations of all 55 landmarks. For each group of 11 transformed points we compute the trace of the covariance of the set of points. Finally we compute the mean covariance trace (MCT) over the five landmark locations.

7.2.4.1 Experiment

In this experiment section, we compare the traditional method, FMAM, FMRestore and FM-Ransac on the “Olgacontrol” dataset. It has 6 $b = 0$ images and 62 DWIs. Mutual information cost function is used for all the methods. We pick 2 $b = 0$ and 9 DWIs, and label 5 corresponding landmarks on each image. If FMR methods provide the best performance, they should have the lowest MCT values of all.

The MCT result using different methods is shown in Table Figure 7.24. The original data



(a) Landmark on $b = 0$ image (b) Corresponding landmark on diffusion-weighted image

Figure 7.6: Landmarks of the same anatomical location in different measurement images

Method	MCT (mm^2)
Original Data	1.4109
Traditional method	2.9837
FMAM	2.0417
FMRestore	2.0191
FMRansac	2.0686

Table 7.24: Mean of covariance matrix Trace (MCT)

set is high quality and contains very little motion corruption particularly at the centre of the brain where the landmarks are. The value of MCT of the original dataset is thus very small. All the registrations lead to slightly higher values, although our model-based methods, FMAM, FMRestore and FMRansac give a better result than the traditional method.

The result may indicate that, for this high quality dataset, realignment is actively detrimental, but lower MCT with no alignment may also reflect bias in landmark positioning from the user tending toward similar image locations. Manual landmarking is very time consuming, and the accuracy is difficult to maintain and guarantee. Thus we abandon the approach and seek a more efficient evaluation method, since labelling landmarks on a large number of datasets is impractical and we cannot draw firm conclusion from the metric anyway. We propose a more automatic evaluation metric in the next section.

7.3 Orientation correction

This section contains two novel contributions of this thesis, which have been published in [14]. First in §7.3.1, we propose an new orientation correction technique, which is used after diffusion MRI registration to update diffusion gradients. A similar idea has also now been explored by Rohde et. al. [119], Maniega et. al. [89] and Leemans and Jones [78, 79]. All authors find minor improvements in datasets and results of post processing, such as tractography, from updating gradient directions. Here we test the influence on our realignment procedure. The second contribution in §7.3.2, we provide an evaluation method for Diffusion MRI registration, which does not require manual labelling or ground truth.

7.3.1 Orientation correction

The registration methods outlined in the previous chapters, including the traditional method, do not account for the effect of rotation on the DWIs. Rotational head motion causes additional contrast changes because of the change in diffusion gradient direction with respect to the head. Figure 7.7 illustrates the effect. The arrows indicate gradient directions. (a) Head without rotation. Suppose the mouth is a fibre. The signal is high because the gradient is perpendicular to the fibre so little diffusion occurs in the gradient direction. (b) Head with rotation. In the rotated head, the signal is lower in the mouth fibre, which is more aligned with the gradient so more diffusion occurs in the gradient direction. (c) The unrotated head (after registration) retains the signal from its rotated position. The effective gradient direction for the corrected image is rotated.

Having computed a correcting affine transformation for each acquisition, we update the effective diffusion-weighting gradient direction to account for the head rotation at imaging time. From the affine transformation T_i obtained from the registration, we use the standard polar decomposition [87] to extract the rotation R_i for each image:

$$R_i = (T_i T_i^T)^{-1/2} T_i, \quad (7.10)$$

and use it to correct each gradient direction $\mathbf{q}_i \rightarrow R_i(\mathbf{q}_i)$.

7.3.2 Orientation evaluation without ground truth

In this section, we provide a “subgroup” evaluation method for Diffusion MRI registration. When the original dataset contains only small motion corruption, it is hard to tell the differences and improvement by visualising FA maps, as shown in §7.2.3. When there is no ground

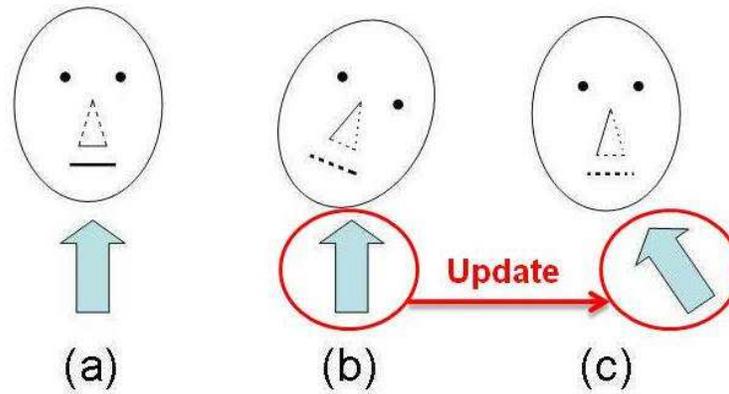


Figure 7.7: Illustration of how rotation affects the effective gradient direction. The arrows indicate gradient directions. (a) Head without rotation. (b) Head with rotation. (c) The unrotated head (after registration) retains the signal from its rotated position.

truth for our registration dataset, to provide numerical evaluation to assess the performance of different methods, we use the method introduced in this section.

We evaluate the variance of statistics derived from the diffusion tensor fitted to subsets of all the measurements, for example, in each voxel. The better aligned all the DWIs, the more consistent statistics like FA, $\text{Tr}(\mathbf{D})$ and principle diffusion direction should be among different subsets. We use Cook et al.’s method in [36] to choose subsets so that each contains measurements with well separated and evenly distributed gradient directions, which should produce fitted DTs with similar statistics. If we divide the whole DWI datasets into 4 groups, to evaluate the variance of the principal directions $\mathbf{e}_1, \dots, \mathbf{e}_4$ in each voxel, we use the largest eigenvalue λ_1 of the dyadic tensor $1/4 \sum_{i=1}^4 \mathbf{e}_i \mathbf{e}_i^T$. When all four \mathbf{e}_i are aligned, $\lambda_1 = 1$. As they become less aligned, λ_1 decreases to a minimum of zero when they are maximally separated.

7.3.2.1 Experiment

We run experiment on “S32” dataset. To do subgroup evaluation of the aligned datasets using different methods, we divide the 60 DW measurements into 4 groups of 15, illustrated in Figure 7.8 (Diffusion gradient vectors from the whole dataset and four subsets, mapped on unit spheres) and Figure 7.9 (FA maps of DTs fitted from the whole dataset and four subsets).

Figure 7.10 illustrate the λ_1 image from original dataset. High value of λ_1 s locate in the area with high FAs, as the local fibres have consistent orientation. We can see the numerical

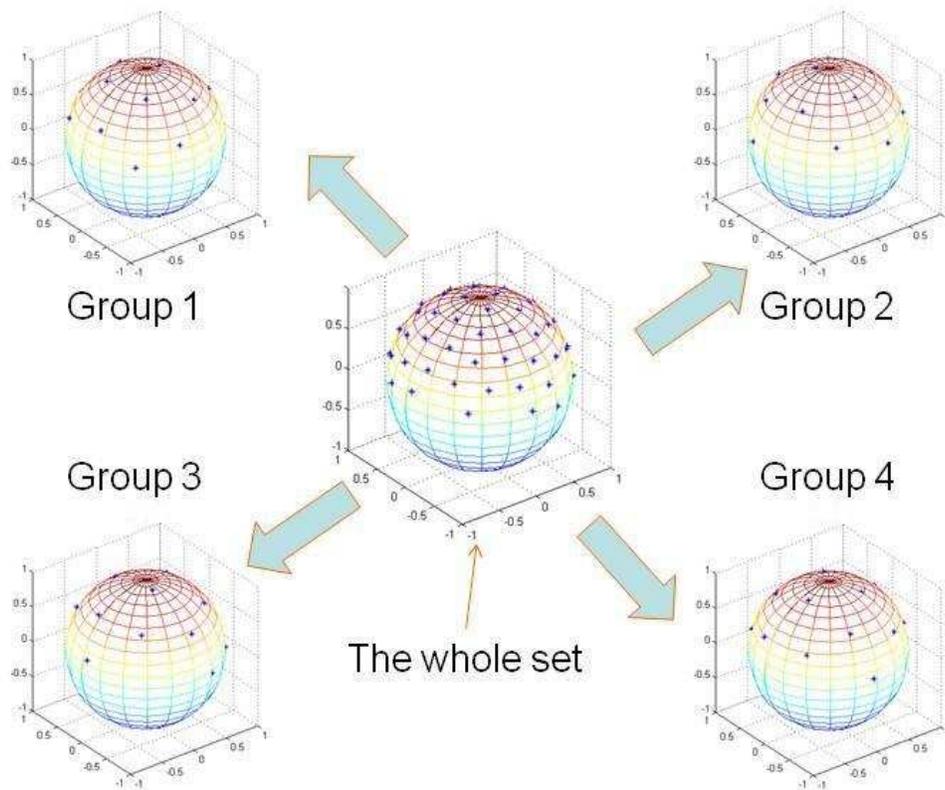


Figure 7.8: Diffusion gradient vectors from the whole dataset and four subsets, mapped on unit spheres.

results in Table 7.25, which shows the average mean and standard deviation (STD) of λ_1 over four regions of the image after alignment of the DWIs using various algorithms. One of the regions is the whole of the skull-stripped brain, which we outlined by thresholding the background. The other regions are shown in Figure 7.11. Two of the regions (cc01, cc02) contain only coherent white matter with a single orientation in the corpus callosum. The third contains white matter with different orientations (cingulum and corpus callosum) and partial volume voxels between those structures.

The first five rows of Table 7.25 compare four methods, traditional, FMAM, FMRestore and FMRansac, with no alignment; the remaining four rows compare mean λ_1 after gradient-direction correction. Similarly, Table 7.27, Table 7.28, Table 7.29 and Table 7.30 shows mean and STD (values inside bracket), through the voxels in the selected volumes/regions, of STD or mean of four groups in FA and Trace.

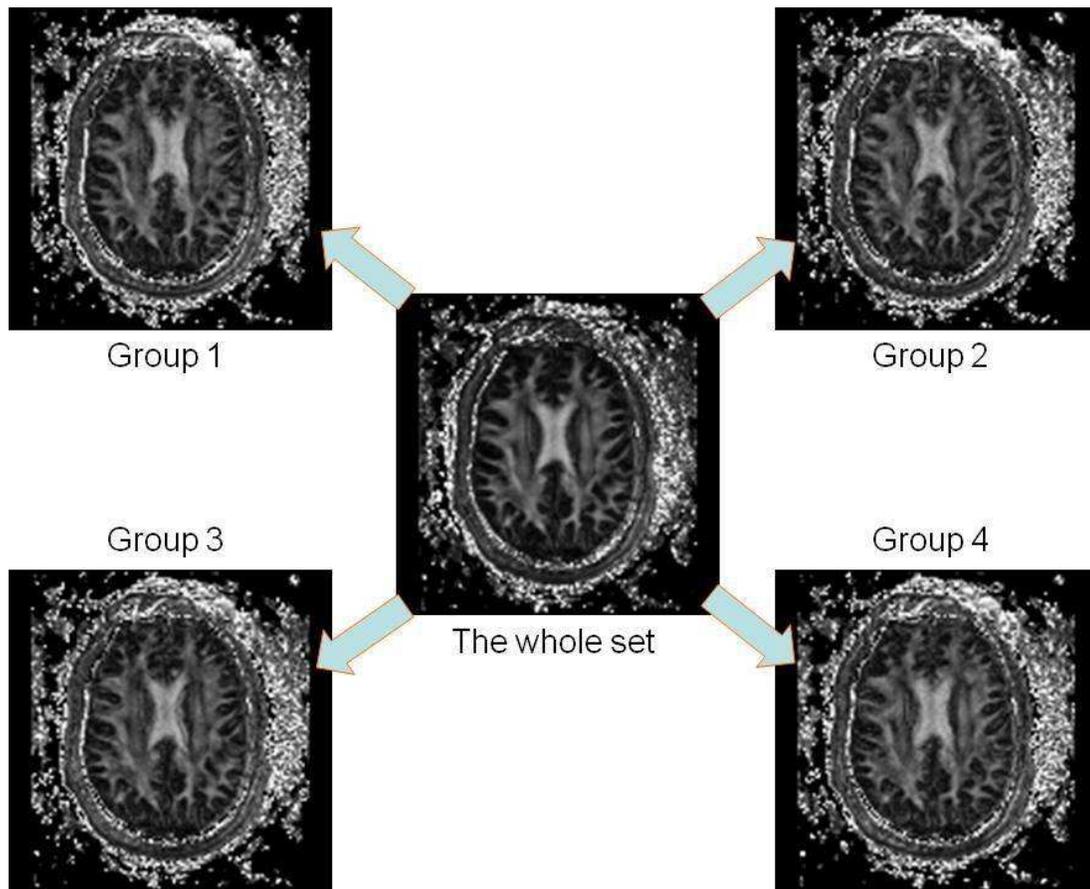


Figure 7.9: FA maps of DTs fitted from the whole dataset and four subsets

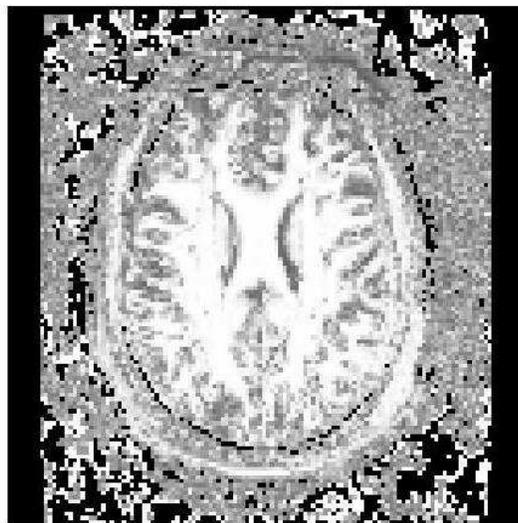


Figure 7.10: λ_1 image from original dataset.

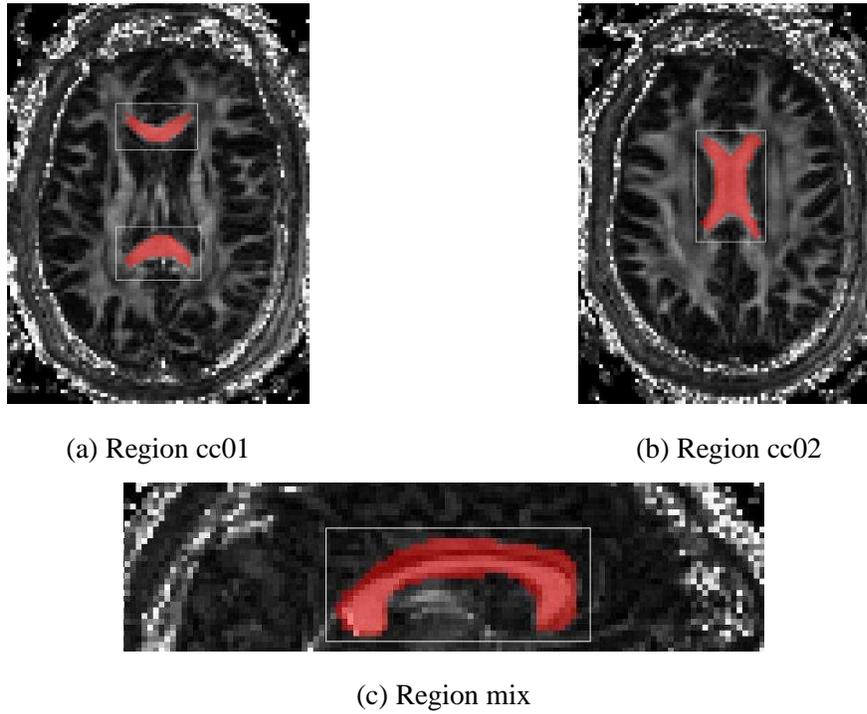


Figure 7.11: Three specified regions (coloured in red) for comparing measurement's alignment.

	Brain	cc01	cc02	mix
Original data	0.8180 (0.1694)	0.9961 (0.0037)	0.9835 (0.0489)	0.9065 (0.1256)
Traditional	0.9206 (0.1136)	0.9989 (0.0010)	0.9959 (0.0282)	0.9676 (0.0615)
FMAM	0.9358 (0.1008)	0.9989 (0.0008)	0.9989 (0.0008)	0.9811 (0.0449)
FMRestore	0.9335 (0.1031)	0.9989 (0.0008)	0.9988 (0.0008)	0.9804 (0.0456)
FMRansac	0.9355 (0.1002)	0.9989 (0.0008)	0.9989 (0.0008)	0.9810 (0.0454)
Correction with Gradient Updated				
Traditional	0.9280 (0.1064)	0.9988 (0.0009)	0.9961 (0.0293)	0.9623 (0.0752)
FMAM	0.9364 (0.0999)	0.9988 (0.0008)	0.9987 (0.0008)	0.9819 (0.0450)
FMRestore	0.9350 (0.1012)	0.9987 (0.0009)	0.9987 (0.0008)	0.9814 (0.0447)
FMRansac	0.9361 (0.1002)	0.9988 (0.0008)	0.9987 (0.0008)	0.9818 (0.0454)

Table 7.25: Mean & (STD) of λ_1 's of dyadic tensors from four groups in four regions (the larger the better and 1 is perfect)

	Brain	cc01	cc02	mix
Original data	2.5390 (1.6326)	5.8575 (0.7867)	5.1666 (1.1789)	3.5996 (1.8506)
Traditional	3.8585 (1.8957)	7.0685 (0.7320)	6.8649 (1.1097)	5.0970 (1.9795)
FMAM	4.0488 (1.8006)	7.0849 (0.8223)	7.0076 (0.7220)	5.4897 (1.6617)
FMRRestore	4.0145 (1.8104)	7.0517 (0.8090)	6.9630 (0.7203)	5.4556 (1.6573)
FMRansac	4.0536 (1.8019)	6.9685 (0.7744)	6.8570 (0.7068)	5.5041 (1.6096)
Correction with Gradient Updated				
Traditional	3.9401 (1.8578)	6.9647 (0.7304)	6.7664 (1.0230)	5.0157 (1.9832)
FMAM	4.0587 (1.8004)	6.9695 (0.7758)	6.8548 (0.7046)	5.5103 (1.6095)
FMRRestore	4.0368 (1.8065)	6.9347 (0.7654)	6.8152 (0.6999)	5.4728 (1.6071)
FMRansac	4.0536 (1.8019)	6.9685 (0.7744)	6.8570 (0.7068)	5.5041 (1.6096)

Table 7.26: Mean & (STD) of $\gamma = -\log(1 - \lambda_1)$'s of dyadic tensors from four groups in four regions (the larger the better)

	Brain	cc01	cc02	mix
Original data	0.0506 (0.0309)	0.0364 (0.0192)	0.0444 (0.0287)	0.0633 (0.0396)
Traditional	0.0216 (0.0118)	0.0204 (0.0100)	0.0209 (0.0096)	0.0263 (0.0158)
FMAM	0.0210 (0.0117)	0.0207 (0.0093)	0.0204 (0.0093)	0.0262 (0.0168)
FMRRestore	0.0213 (0.0118)	0.0210 (0.0094)	0.0206 (0.0095)	0.0273 (0.0192)
FMRansac	0.0211 (0.0117)	0.0208 (0.0093)	0.0203 (0.0093)	0.0261 (0.0165)
Correction with Gradient Updated				
Traditional	0.0214 (0.0117)	0.0198 (0.0094)	0.0218 (0.0096)	0.0251 (0.0153)
FMAM	0.0209 (0.0116)	0.0200 (0.0090)	0.0228 (0.0097)	0.0256 (0.0166)
FMRRestore	0.0212 (0.0119)	0.0204 (0.0092)	0.0231 (0.0098)	0.0262 (0.0165)
FMRansac	0.0210 (0.0117)	0.0201 (0.0090)	0.0228 (0.0097)	0.0255 (0.0163)

Table 7.27: Mean & (STD) of STD of FAs from four groups through the voxels in the selected volumes/regions

	Brain	cc01	cc02	mix
Original data	0.2825 (0.1473)	0.7067 (0.1026)	0.6955 (0.0827)	0.5248 (0.2311)
Traditional	0.2164 (0.1530)	0.6812 (0.1182)	0.6201 (0.1367)	0.4134 (0.2553)
FMAM	0.2264 (0.1526)	0.6939 (0.1035)	0.6568 (0.0936)	0.4715 (0.2465)
FMRestore	0.2250 (0.1524)	0.6912 (0.1029)	0.6552 (0.0960)	0.4740 (0.2456)
FMRansac	0.2263 (0.1526)	0.6941 (0.1035)	0.6570 (0.0935)	0.4716 (0.2466)
Correction with Gradient Updated				
Traditional	0.2172 (0.1520)	0.6803 (0.1188)	0.6238 (0.1316)	0.4089 (0.2579)
FMAM	0.2266 (0.1527)	0.6947 (0.1031)	0.6574 (0.0930)	0.4713 (0.2467)
FMRestore	0.2258 (0.1529)	0.6951 (0.1024)	0.6597 (0.0912)	0.4714 (0.2471)
FMRansac	0.2265 (0.1527)	0.6947 (0.1031)	0.6575 (0.0929)	0.4713 (0.2468)

Table 7.28: Mean & (STD) of FAs from four groups

$(m^2/s)*10^{-11}$	Brain	cc01	cc02	mix
Original data	12.07 (13.24)	8.67 (4.03)	12.33 (6.06)	16.08 (11.77)
Traditional	4.17 (3.85)	4.58 (2.28)	5.41 (2.49)	5.36 (3.81)
FMAM	4.41 (4.79)	4.41 (2.09)	4.70 (2.06)	5.42 (3.69)
FMRestore	4.42 (4.74)	4.51 (2.12)	4.80 (2.08)	5.57 (3.94)
FMRansac	4.45 (4.83)	4.42 (2.10)	4.70 (2.07)	5.46 (3.67)
Correction with Gradient Updated				
Traditional	4.46 (4.05)	5.17 (2.35)	6.33 (2.80)	6.05 (3.97)
FMAM	4.52 (4.80)	4.87 (2.34)	5.17 (2.28)	5.72 (3.69)
FMRestore	4.54 (4.78)	4.96 (2.37)	5.32 (2.30)	5.90 (3.93)
FMRansac	4.55 (4.85)	4.88 (2.33)	5.17 (2.28)	5.76 (3.67)

Table 7.29: Mean & (STD) of STD of Traces

$(m^2/s)*10^{-9}$	Brain	cc01	cc02	mix
Original data	3.35 (1.69)	2.61 (0.38)	2.72 (0.35)	3.23 (1.38)
Traditional	3.21 (1.42)	2.70 (0.59)	3.13 (0.93)	3.20 (1.17)
FMAM	3.29 (1.56)	2.63 (0.42)	2.83 (0.50)	3.21 (1.26)
FMRestore	3.29 (1.56)	2.63 (0.41)	2.82 (0.49)	3.22 (1.26)
FMRansac	3.29 (1.56)	2.63 (0.42)	2.83 (0.50)	3.21 (1.26)
Correction with Gradient Updated				
Traditional	3.26 (1.44)	2.73 (0.60)	3.18 (0.95)	3.24 (1.18)
FMAM	3.28 (1.55)	2.62 (0.41)	2.81 (0.50)	3.20 (1.25)
FMRestore	3.29 (1.56)	2.62 (0.41)	2.81 (0.49)	3.21 (1.26)
FMRansac	3.28 (1.55)	2.62 (0.41)	2.81 (0.50)	3.20 (1.25)

Table 7.30: Mean & (STD) of Traces

The better aligned all the DWIs, the more consistent FA, Tr(D) and principal diffusion direction should be among different subsets. In Table 7.25, λ_1 s in the first row are the smallest. That means all four methods, with and without gradient updating, improve the alignment of the four sub-tensors on average, which indicates that the signal contains less noise or corruption, since the principal directions within voxels are more consistent. However, the differences between the mean λ_1 are quite small compared to the STD of λ_1 across voxels. Although the relatively large STDs suggest the differences in mean are not significant, we consistently observe those trends, which is evidence for improved alignment from FMAM and FMR.

One problem with λ_1 as a comparative measure of directional coherence is that it is bounded by the range zero to one. Differences in the numerical value of λ_1 become small for collections of very similar directions, where the numerical value of λ_1 is close to 1. This may mask significant differences in populations of λ_1 given in Table 7.25. To address this potential problem, we rescale by computing $\gamma = -\log(1 - \lambda_1)$, which provides an alternative directional coherence measure for which the numerical value approaches infinity as the directions become perfectly aligned. Standard deviations of γ are more meaningful, as its distribution is less skewed for collections of well aligned directions with λ_1 close to 1. Table 7.26 shows γ . In fact, here we do not see any more significant differences using γ than using λ_1 directly.

According to the average λ_1 , the FMAM and FMRs methods give better performance than the traditional method, and FMRansac does slightly better than FMRestore. We can draw similar conclusions from tables of mean of FA and Trace in table 7.28 and 7.30. We observe an increase in mean FA and reduction in mean Trace using the new alignment procedures, which also suggests improved alignment reducing small partial-volume effects that artificially reduce FA and increase Trace. However, once again the STD of mean FA and Trace are high compared to the differences, so firm conclusions are difficult to draw. We also expected to observe reductions in the average STD of both FA and Trace, in table 7.27 and 7.29, which would reflect more consistency of these values within voxels. However, we do not see consistent reduction with FMAM and FMR compared to the traditional method, although all methods show improvement compared to the original data.

The extra step of correcting the diffusion gradients generally improves alignments slightly further on average. The rotation component from equation (7.10) includes contribution from the shear caused by eddy-current distortion, which does not affect the gradient orientation. We

may see further improvement if we can separate the rotational corruption from motion and the eddy-current distortion and use only the former to correct gradient orientations. Other work, for example by Leemans and Jones [78, 79] suggest that although these improvements are minor, they may have significant effects on postprocessing operations such as tractography.

We note that in this example, FMRs offer no advantage over FMAM. However, it introduces no significant disadvantage and, in cases with larger distortions or motion problems, it is often more robust so we advocate its use.

7.4 Conclusion

For good performance in RANSAC and FMRansac, a good fraction f , which decides the minimum number of measurements required to accept the model, should be chosen according to the expected number corrupted measurements of dataset. A smaller confidence interval might increase accuracy to datasets with little corruption, but may make RESTORE and RANSAC less robust. MI improves performance in general from NC; but NC works better than MI for some corrupted reference images. So the best choice remains unclear, but both perform fairly well.

Compared to the traditional method, FMAM and FMRs use different reference images for diffusion-weighted images with different gradient directions for registration, in order to improve alignment. FMRs detect outliers and use better aligned subset to predict reference images, so give more robust performance. Our experiments provide evidence for this improvement, although they do not achieve statistical significance so more work is required to verify the assertion.

The extra step of correcting the diffusion gradients generally improves performance metrics further. Although improvements are slight, the procedure is simple to perform and experiments with tractography [78] suggest the differences can be significant.

Chapter 8

Conclusion

In this chapter, we first give a summary of this thesis, and then point out several potential areas for further work.

8.1 Discussion

For fitting the diffusion tensor, voxels in different diffusion-weighted images must correspond to the same anatomical location. Thus, all the measurement images need to be well aligned. The long scan time introduces patient movement. Moreover, EPI induces displacement and distortion in diffusion-weighted MR imaging.

The traditional correction methods use the same reference image to register all the other diffusion-weighted images with different gradient directions. Although they correct both eddy-current-induced distortion and subject motion, the contrast differences still cause misregistration. Our methods predict separate reference images for each diffusion gradient and thus avoid the mismatching caused by the intensity differences between component images.

All our new methods are based on a three-step procedure to register DWI datasets. They use different reference images for diffusion-weighted images with different gradient directions for registration. So the registrations take into account the contrast differences of measurements. The main difference between methods lies in the first step, fitting the diffusion tensor to the measurements.

The model-based approach works well if we can avoid the effects of corrupted images on the references. The RESTORE and RANSAC identify outliers in the set of measurements in each voxel and thus fit the tensor more robustly. The application of RANSAC to DTI is a novel

contribution of this thesis.

FMAM fits the diffusion tensor D to the whole set of diffusion weighted measurements, and most the other FM methods (FM.cons, FM.a, FM.b, FM.th) and FMRs (FMRestore and FMRansac) fit the tensor only to a selected subset of measurement images. FM.a, FM.b, FM.th, FMRestore and FMRansac all use different automatic algorithms for selecting this subset. Results show that for FMAM, FM.a and FM.b, the corrupted measurements affect tensor fitting too strongly to obtain accurate reference images; FM.th can identify the most corrupted measurements but requires careful tuning for specific datasets; FMRestore and FMRansac generalize better, because RESTORE and RANSAC reject outliers using more robust and established statistical techniques.

We run experiments for different correction methods on both synthetic reduced-size and full-size diffusion MRI datasets. To evaluate performances of different methods, we not only judge the result from qualitative image inspection, such like DWIs and FAs, but also develop quantitative methods, which includes using outlier histogram, transformations used in registration, RMSE of DWIs, RMSE of FAs, principal direction differences to ground truth images, covariance of landmark coordinates, and subgroup evaluation, etc.

Qualitative results consistently suggest that the FMAM and FMRs improve on the traditional alignment procedure. MI cost function accounts for non-linear contrast differences, and improves traditional method using NC cost function; it also improves model-based registration methods in general, but NC works better for some corrupted reference images. The best choice is thus still open to debate, but both choices are effective and we find MI works better more often. Quantitative results provide further evidence for better performance of the model-based techniques, but statistically significant differences proved hard to establish.

Body rotation leads the mismatching of diffusion pulse and diffusion weighted image in registration. Orientation correction step can improve alignments from all methods further.

FMR methods were built into Camino, which is a free, open-source, object-oriented software package for analysis and reconstruction of Diffusion MRI data. Computation time for a typical dataset such as those used in Chapter 7 on a modern desktop is about 120mins, which is about the same as eddycorrect in FSL. Man page and tutorial document are also provided and

published on line (www.camino.org.uk), which are included in Appendix A and B.

8.2 Further Outlook

In this section, we list some possible aspects of further work, which include update reference image model, reject 3D measurement using outlier histogram and some alternative choices for what we have done.

8.2.1 Update reference image model

The gradient update correction can be based on two levels: The orientation correction can be implemented after registration, which is what we have done in §7.3.1; the other possible scheme is using the gradient updated image as the reference, and register source images for the second time. This may improve results, particularly for large rotations that cause significant changes in contrast as well as alignment.

8.2.2 Reject 3D measurement using outlier histogram

FM.cons, FM.a, FM.b, FM.th use some less robust scheme to reject the whole set of DW measurement; FMRs use robust schemes to reject outliers in voxel level. We can combine two approaches and use the outlier histogram to decide which measurements (3D volume) can be rejected totally during the tensor fitting in synthetic reference image making.

8.2.3 Some alternative choices

In this thesis, we only use DT model to generate reference, but other diffusion model is possible as well. As the purpose of fitting diffusion model in step 1 is to generate a reference DW images, comparing with linear tensor model, more complicated non-linear diffusion model seem not necessary. However, for diffusion MRI acquisition protocols that include much higher b -values, the limitations of the DT model become more apparent and model-based registration based on the DT is likely to break down. Alignment of such datasets is a significant problem and the traditional approach of aligning to $b = 0$ images does not work at all because the contrast is so different. More sophisticated model-based approaches potentially offer a good solution.

In registration we have done, only rigid and affine transformations were used. The method

extends to non-linear registration, which may be beneficial for correcting susceptibility distortion. Non-linear registration can introduce artifacts, but we expect the model-based approach to be more robust than matching to $b = 0$ where local contrast changes may cause significant errors with non-linear registration.

So far, we have used only 3D transformations to mainly correct the subject motion. Results may improve by using more flexible transformations such as 3D rigid body combined with slice by slice 2D affine transformations [70] to correct for both patient subject motion and eddy-current-induced distortions. Andersson and Skare [8] suggest models for how distortions vary with slice number and diffusion gradient direction. This shows us a way to develop our method on the slice-by-slice registration.

Appendix A

Tutorial: Motion Correction for Diffusion Weighted Images

This tutorial gives a general introduction of using function `malign` in Camino package to align the diffusion-weighted images within a single acquisition.

SYNOPSIS

`malign` [options]

A.1 Preparation before running `malign`

A.1.1 General data files and information

Before running `malign`, we need to have input and scheme files ready, and use

`-inputfile` <Input voxel-order file>

`-schemefile` <Scheme file name>

to specify in command options. If the input file is in scanner-order, we can use

`-scanner -inputfile` <Input scanner-order file>.

And also we need to make some other information of input image data clear and specify a few other options

`-datadims X Y Z` <Number of voxels in each dimension>

`-voxeldims x y z` <Voxel size in mm>

`-sigma` <Standard deviation of noise>

A.1.2 Sigma

The sigma is approximate noise standard deviation. A suggested value is $\sqrt{E(S^2)/2}$, where S is the signal in background and E denotes expectation over an ROI.

A camino program, `datastats`, can work it out for you as well.

```
datastats -schemefile S.scheme -bgmask S32_M100.Bshort -inputfile S32.Bfloat
```

where S32_M100.Bshort is a mask file for ROI (we will discuss more about generating mask in Section 3), then the screen output would be:

```
=====
Foreground voxel count: 142584
Component E(S) E(S^2) Var(S) Std(S)
1 2.882503E02 9.834995E04 1.526169E04 1.235382E02
2 3.010773E02 1.070645E05 1.641700E04 1.281288E02
3 2.876059E02 9.854322E04 1.582610E04 1.258018E02
4 2.564062E02 7.868863E04 1.294450E04 1.137739E02
5 7.893223E01 7.438512E03 1.208215E03 3.475939E01
6 7.813621E01 7.298962E03 1.193695E03 3.454989E01
7 8.115909E01 7.870362E03 1.283565E03 3.582688E01
8 7.870379E01 7.465566E03 1.271279E03 3.565500E01
9 8.104558E01 7.776110E03 1.207724E03 3.475232E01
10 8.053585E01 7.688916E03 1.202894E03 3.468276E01
11 8.019136E01 7.728569E03 1.297915E03 3.602659E01
12 7.751786E01 7.154714E03 1.145696E03 3.384814E01
13 8.196583E01 8.091765E03 1.373367E03 3.705897E01
14 7.725234E01 7.106027E03 1.138104E03 3.373580E01
15 8.016002E01 7.625693E03 1.200064E03 3.464194E01
=====
```

So we can choose 7.5 as the value of $E(S^2)$. Then σ , $\sqrt{E(S^2)/2}$, would be 1.9.

A.1.3 Make slice to check input volume

If the image file is in voxel-order, we need to transfer it to scanner-order first:

```
voxel2scanner -voxels $((128*128*32)) -inputdatatype float -outputdatatype float -
components 64 -inputfile S32.Bfloat > S32.scan.Bfloat
```

Then we can use camino function shredder to extract one slice from each of the 32 3D components, and build a 3D image shown in Fig1:

To make the command easy to read, we can use variables representing data information:

```
COMPONENT=64
DATADIM_X=128
DATADIM_Y=128
DATADIM_Z=32
```

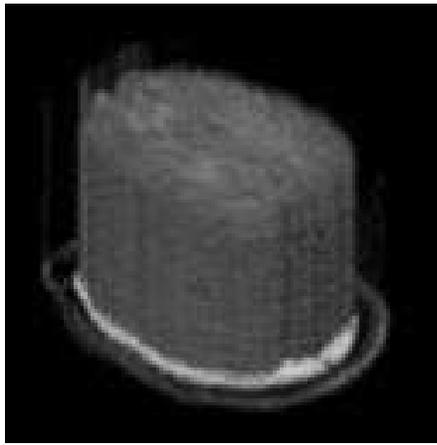


Figure A.1: An image volume containing 32 slices from 64 diffusion components.

Since our dataset used in the example is float, so

```
TYPESIZE=4
```

If we would like to extract the middle slice along z direction, we can set

```
OFFSET=$(( $DATADIM_X * $DATADIM_Y * (($DATADIM_Z / 2)) * $TYPESIZE ))
```

Then, calling shredder

```
shredder $OFFSET $($DATADIM_X * $DATADIM_Y * $TYPESIZE) $($DATADIM_X * $DATADIM_Y * (($DATADIM_Z - 1)) * $TYPESIZE) < S32.scan.Bfloat > S32..SLICECHECK.img
```

Make a header file (*.hdr) to make *.img file readable by many visualisation softwares:

```
VOXELDIM_X=1.88
```

```
VOXELDIM_Y=1.88
```

```
VOXELDIM_Z=2.0
```

```
analyzeheader -voxeldims $VOXELDIM_X $VOXELDIM_Y $VOXELDIM_Z -datadims $DATADIM_X $DATADIM_Y $COMPONENT -datatype float > S32..SLICECHECK.hdr
```

Using some visualisation tools, such like MRIcro(Fig.2) and camino, we can check the slice motion of input data set.

A.2 Run mbalign

A.2.1 Run mbalign simply

An example of using mbalign in the simplest way is:

```
mbalign -datatype float -scheme S.scheme -datadims 128 128 60 -voxeldims 1.88 1.88 2.0 -sigma 1.9 -inputfile S32.Bfloat
```

then the screen output would be:

```
=====
```

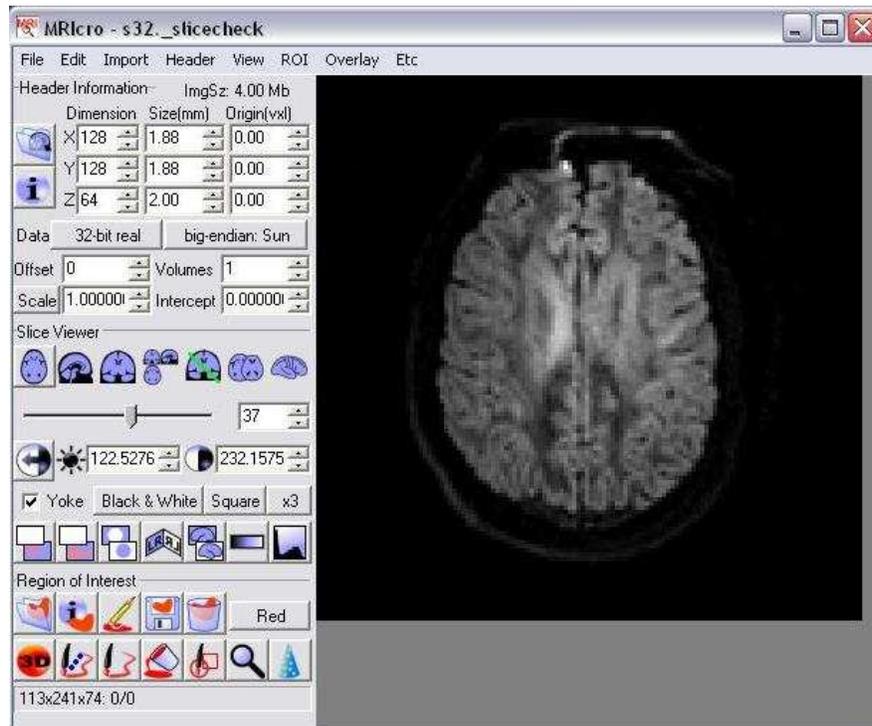


Figure A.2: Using MRIcro to see slice volume

:-) I have everything I need!

WARNING: No -bgmask and -bgthresh input, using zero background threshold. Performance may improve with better threshold.

WARNING: No -components input. Using 64 components from schemefile. If wrong, restart with -components option.

64 components inside inputfile.

No temp directory specified. Trying to use /tmp/S32_17.03.08_180229.

Successfully created /tmp/S32_17.03.08_180229.

Linux system detected.

No output file name specified. Output file will be: /tmp/S32.out.Bfloat

Disk space temporarily used during calculation is about 2264924160. Make sure space is available!

=====

The program will create a temporary directory used for calculation process. It can be either specified by -tmpdir or automatically create according to the file name and current time.

We do not have to use -outfile to specify the output file name, and the program can

generate it according to the input file name, which is like the example shown above.

To run mbalign, computer need to have registration software flirt([fmrib.ox.ac.uk/fsl/flirt/](http://www.fmrib.ox.ac.uk/fsl/flirt/)) installed, which is part of FSL library(<http://www.fmrib.ox.ac.uk/fsl/>). We also need to specify the flirt direction by using -fsldir, but more easily, once we have camino and FSL installed, default value of variable DIR_FSL can be set which can be easily found inside mbalign.

During the running of mbalign, there could be a lot of warning messages showing inside terminal window. Normally, just do not worry about it too much. Once the registration has got done, we may see the message shown below.

```
=====
...
Registering ori.ScannerOrder.Bfloat.ck by using ref.ScannerOrder.Bfloat.ck
Registering ori.ScannerOrder.Bfloat.cl by using ref.ScannerOrder.Bfloat.cl
Transferring output to Big-endian format...
Making img and hdr files for slice checking...
Transfer output file to voxel-order...
Removing junk files...
Scheme file not updated.
Aligned data set output to /tmp/S32.out.Bfloat.
Program finished at
Mon Mar 24 23:36:24 GMT 2008
=====
```

After program finishes, the temporary directory would be deleted (Removing junk files...), but we can still use -keepjunk to make it kept, which could be useful to help us to analyse the final output. Certainly, if the program is interrupted, this temporary fold will remain as junk files in computer.

If we did not specify -outputfile, the program can generate it according to the input file name (Aligned data set output to /tmp/S32.out.Bfloat).

We will discuss scheme file updating in Section 4.Update Gradient.

A.2.2 Run mbalign in an advanced way

Some other options might need to use:

`-flirtsearchcost <Search cost function used in flirt>`

Default cost function is `mutualinfo` (Mutual Information). Other options are `corratio`, `normcorr`, `normmi`, `leastsq`.

`-flirttransform <Transformation used in flirt>`

Default transformation is `affine`. The other option is `rigid`.

`-searchrange <angle>`

Default is 90, which means search range is between -90 and 90 in all x, y and z directions.

`-eddy`

Specifies registration for eddy-current induced distortion.

`-datatype <Data type for input and output files>`

Default is `float`.

`-scanout <output scanner-order file>`

Adds an extra output file in scanner-order. This won't stop default voxel-order output.

`-omat <File name>`

Output transform matrix in ascii format.

`-slicecheck <File name>`

Output a pair of `<File name>.img` and `<File name>.hrd` files. Default is no calculation.

When all the options are decided, we can run `mbalign` in an advanced way, such like

```
mbalign -datatype float -schemefile S.scheme -datadims 128 128 60 -voxeldims 1.88
1.88 2.0 -sigma 1.9 -fsldir /cs/research/medim/common0/green/common/fsl/fslRH9/ -inputfile
S32.Bfloat -slicecheck S32.fmr.slice.check -outputfile S32.fmr.Bfloat -omat S32.fmr.mat.txt
-scanout S32.fmr.scanout.Bfloat -keepjunk -tmpdir tmp.fmr
```

A.2.3 Improve performance

There are a few options can be used to improve the performance of `mbalign`.

`-sigma`

High *sigma* allows the program to involve more measurement in the DT fitting, and low *sigma* leads rejections during the DT fitting. Based on this theory, we can change the value of *sigma* to improve the reference making.

`-bgmask <Mask file>`

Use a mask file can improve the quality of the reference images used in `mbalign` registration.

And the data type of mask file should be "short".

Camino function mask can help to create a background mask from a voxel-ordered DW data file by thresholding the average $b=0$ measurement.

```
mask -inputfile S32.Bfloat -inputdatatype float -schemefile S.scheme -bgthresh 100 -
outputdatatype short > S32_M100.Bshort
```

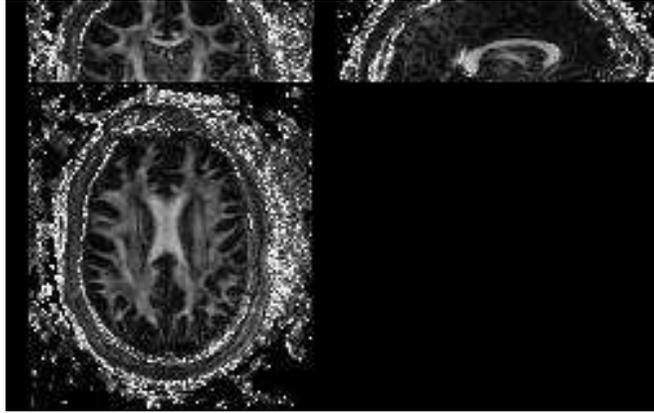


Figure A.3: View projection of FA map generated by camino

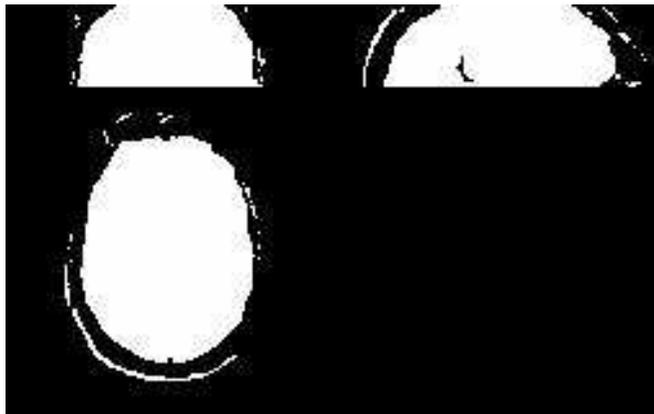


Figure A.4: View projection of mask file generated by camino

We can also use matlab to make a mask file.

-bgthresh <Background threshold>

Decide the value of threshold, and improve fitting the diffusion tensors.

-searchrange <angle>

Sometimes, the pitch of histogram will cause the failure of registration. Simply narrow down the angle search range can cover this problem for most of the time.

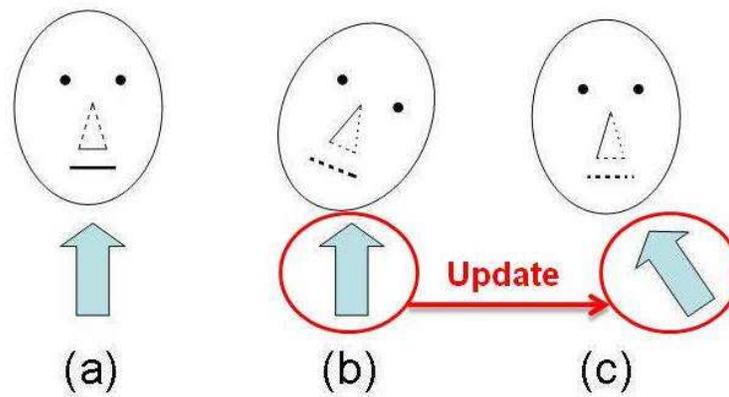


Figure A.5: Illustration of how rotation affects the effective gradient direction. The arrows indicate gradient directions. (a) Head without rotation. (b) Head with rotation. (c) The unrotated head (after registration) retains the signal from its rotated position.

A.2.4 Update Gradient

Updating diffusion gradients after registration is not an essential procedure, but can improve the registration result.

Fig.5 explains the reason why diffusion gradients need to be updated after registration.

We can use our matlab function UpdateGradient.m to update diffusion gradient for registered data set, and the usage method is explained inside the matlab file. Since the transformations used in the registration contribute to the gradient updating, we **MUST** save the transform matrix when running mballign, using -omat.

A.2.5 More hints

Quite a few default options can be change in the source code of mballign. Make the default values to the ones most frequently used can make the everyday use of mballign much simpler. Inside mballign source code, we can find and change the default options from the following part.

```
#####
##### Change default variables to match system #####
#####
# Hint: To make your input arguments simple, set
# default input which you most often to use.
# FSL directory
```

```
DIR_FSL=/cs/research/medim/common0/green/common/fsl/fslRH9
# LIM_ROTATE is default for -searchrange
LIM_ROTATE=90
# Available cost functions are:
# mutualinfo corratio,normcorr,normmi,leastsq.
SEARCH_COST=mutualinfo
#Degree of freedom
# 12 for affine; 6 for rigid.
DOF=12
```

A.3 Acknowledgement

We would like to thank Geoff Parker and Karl Embleton, University of Manchester, for providing the brain data.

A.4 Reference

- Y. Bai and D. C. Alexander, "Model-Based Registration to Correct for Motion between Acquisitions in Diffusion MR Imaging", *The Fifth IEEE International Symposium on Biomedical Imaging (ISBI 2008)*, May 2008.

Appendix B

Man Page of mbalign

mbalign(1)

B.1 NAME

mbalign - Aligns the diffusion-weighted images within a single acquisition.

B.2 SYNOPSIS

```
mbalign -inputfile <input voxel-order file> -schemefile <filename> -datadims <X Y Z> -  
voxeldims <x y z> -sigma <noise standard deviation>
```

B.3 DESCRIPTION

Reads DWI in voxel-order data from the standard input, fits the DT model using RESTORE, synthesizes reference images from the model, and aligns each measurement of input to its synthetic reference. Finally, it gives the output as voxel-order big-endian data file.

This program requires an installation of FSL, as it uses FLIRT for alignment. By default the program uses a 3D affine registration to align each image, although you can specify 3D rotations. Occasionally FLIRT fails when trying to compute 3D affine registrations. When it does, the program attempts to align the diffusion weighted image volume to the reference volume using a 3D rigid transformation instead. If that fails too, the output file contains a copy of the input image volume for that measurement.

B.4 OPTIONS

(1) Required options:

-inputfile < *input voxel-order file* > The input file must have BIG-ENDIAN ordering.

-schemefile < *Scheme file name* >

-datadims < *X Y Z* > Specifies the number of voxels in each dimension.

- voxeldims** < *x y z* > Specifies the voxel sizes in each dimension, in millimetres.
- sigma** < *noise standard deviation* > The approximate noise standard deviation, sigma. A suggested value is $\sqrt{E(M^2)/2}$, where M is the signal in background and E denotes expectation over an ROI. A camino program called `datastats` works it out for you. See `datastats(1)`, `modelfit(1)`.

(2) Optional options:

- outputfile** < *output voxel-order file* > Default outfile is derived from input file name, but user can specify the file name including the directory.
- datatype** < *data type for input and output files* > Specifies the same data type for the input and output file. The data type can be any of the following strings: “char”, “short”, “int”, “long”, “float” or “double”. By default, the input type is “float”.
- bgmask** < *mask file* > The data type of mask file should be “short”. See `modelfit(1)`.
- bgthresh** < *background threshold* > See `modelfit(1)`.
- tmpdir** < *temp directory for calculation* > The program creates quite a lot of temporary files, which are stored in this directory. Deault name is derived from input file name, current date time, and is a subdirectory of “/tmp”. But if “/tmp” is small, specifying an alternative location is necessary. User needs a separate directory for each data set when running “mbalgin” multiple processes concurrently. By default, the program removes the directory containing all the temporary files, but user can tell it to keep it all by adding `-keepjunk`.
- fsldir** < *FSL directory* > Specifies the location of FSL installation, which is used to do the registration. Such as “../common/fsl/fslSolaris”. FSL must be installed to run `mbalign`.
- flirtsearchcost** < *search cost function used in flirt* > Default cost function is “mutualinfo” (Mutual Information). Other options are “corratio”, “normcorr”, “normmi” and “leastsq”.
- flirttransform** < *Transformation used in flirt* > Default transformation is “affine”. The other option is “rigid”.
- omat** < *file name* > Output transform matrix in ascii format.
- keepjunk** Tells the program to keep temporary files in the directory specified by “-tmpdir”. Default behaviour is to remove all temporary files before program finish.

-slicecheck < *file name* > Specifies the root name of an analyze file that you can check the alignment in. After registration, the file contains the corresponding slice of each DWI for comparison and to check that nothing weird happened and that the alignment is good. If you omit -slicecheck, it won't output this pair of files.

-eddy Specifies registration for eddy-current induced distortion.

-scanner Regards input file is in scanner-order. Command line example: -scanner -inputfile <file name>

-scanout < *output scanner-order file* > Adds an extra output file in scanner-order. This won't stop default voxel-order output.

-searchrange < *angle* > Default is 90, which means search range is between -90 and 90 in all x, y and z directions.

B.5 EXAMPLES

Do the registration on data set A.Bfloat, and store the output in A_AlignAffine.Bfloat:

```
mbalign -inputfile A.Bfloat -schemefile A.scheme -datadims 128 128 60 -voxeldims 2 2
2.7 -bgthresh 200 -sigma 50 -fsldir /cs/research/medim/common0/green/common/fsl/fslRH9/
-slicecheck /tmp/Ser02SliceCheckRigid -outputfile A_AlignAffine.Bfloat
```

For the example above, if we do not specify -outputfile, the output will be /tmp/A.out.Bfloat.

B.6 AUTHORS

BAI Yu <camino@cs.ucl.ac.uk>

B.7 SEE ALSO

modelfit(1), dat asynth(1), analyzeheader(1), datastats(1).

Bibliography

- [1] Freeborough P A, Woods R P, and Fox N C. Accurate registration of serial 3DMRbrain images and its application to visualizing change in neurodegenerative disorders. *Journal of Computer Assisted Tomography*, 20(6):1012–1022, 1996.
- [2] D. C. Alexander, C. Pierpaoli, P. J. Basser, and J. C. Gee. Spatial transformation of diffusion tensor magnetic resonance images. *IEEE Transactions On Medical Imaging*, 20(11):1131–1139, November 2001.
- [3] Daniel C Alexander. *Visualization and Processing of Tensor Fields*, chapter An Introduction to Computational Diffusion MRI: the Diffusion Tensor and Beyond, pages 83–106. Springer Berlin Heidelberg, June 2007.
- [4] Daniel C. Alexander and Gareth J. Barker. Optimal imaging parameters for fiber-orientation estimation in diffusion MRI. *NeuroImage*, 27:357–367, 2005.
- [5] Daniel C. Alexander and James C. Gee. Elastic matching of diffusion tensor MRIs. *Computer Vision and Image Understanding*, 77:233–250, 2000.
- [6] D.C. Alexander, G.J. Barker, and S.R. Arridge. Detection and modeling of non-gaussian apparent diffusion coefficient profiles in human brain data. *Magnetic Resonance in Medicine*, 48:331–340, 2002.
- [7] Yali Amit. A nonlinear variational problem for image matching. *SIAM Journal on Scientific Computing*, 15(1):207–224, 1994.
- [8] Jesper L. R. Andersson and Stefan Skare. A model-based method for retrospective correction of geometric distortions in diffusion-weighted EPI. *NeuroImage*, 16:177–199, 2002.
- [9] N. Arad, N. Dyan, D. Reisfield, and Y. Yeshurun. Image warping by radial basis functions: Application to facial expressions. *Computer Vision, Graphics and Image Processing: Graphical Models and Image Processing*, 56(2):161–172, 1994.

- [10] Siamak Ardekani and Usha Sinha. Geometric distortion correction of high-resolution 3 T diffusion tensor brain images. *Magnetic Resonance in Medicine*, 54:1163–1171, 2005.
- [11] John Ashburner and Karl J. Friston. Nonlinear spatial normalization using basis functions. *Human Brain Mapping*, 7(4):254–266, 1999.
- [12] Y. Assaf, R. Z. Freidlin, G. K. Rohde, and P. J. Basser. New modelling and experimental framework to characterize hindered and restricted water diffusion in brain white matter. *Magnetic Resonance in Medicine*, 52:965–978, 2004.
- [13] Mordecai Avriel. *Nonlinear Programming: Analysis and Methods*. Number ISBN 0-486-43227-0. Dover, 2003.
- [14] Yu Bai and D.C. Alexander. Model-based registration to correct for motion between acquisitions in diffusion MR imaging. *The Fifth IEEE International Symposium on Biomedical Imaging*, pages 947–950, May 2008.
- [15] H. S. Baird. *Model-based image matching using location*. MIT Press, 1984.
- [16] R. Bajscy and S. Kovacic. Multiresolution elastic matching. *Computer Vision, Graphics, and Image Processing*, 46:1–21, 1989.
- [17] R. Bammer and M. Auer. Correction of eddy-current induced image warping in diffusion-weighted single-shot EPI using constrained nonrigid mutual information image registration. In *International Society for Magnetic Resonance in Medicine (ISMRM)*, page 508, Glasgow, U.K., 2001. Proc. Int. Soc.
- [18] Roland Bammer, Michael Augustin, Siegrid Strasser-Fuchs, Thomas Seifert, Peter Kapeller, Rudolf Stollberger, Franz Ebner, Hans-Peter Hartung, and Franz Fazekas. Magnetic resonance diffusion tensor imaging for characterizing diffuse and focal white matter abnormalities in multiple sclerosis. *Magnetic Resonance in Medicine*, 44(4):583–591, October 2000.
- [19] E. I. Barnea and H. F. Silverman. A class of algorithms for fast digital image registration. *IEEE Transactions on Computers*, C-21:179–186, 1972.
- [20] P. J. Basser, J. Mattiello, and D. Le Bihan. MR diffusion tensor spectroscopy and imaging. *Biophysical Journal*, 66:259–67, 1994.

- [21] P. J. Basser, J. Mattiello, and D. LeBihan. Estimation of the effective self-diffusion tensor from the NMR spin echo. *Journal of Magnetic Resonance Imaging*, 103(3):247–254, Mar 1994.
- [22] PJ Basser, S Pajevic, C Pierpaoli, J Duda, and A Aldroubi. In vitro fiber tractography using DT-MRI data. *Magnetic Resonance in Medicine*, 44:625–632, 2000.
- [23] T. E. J. Behrens, M. W. Woolrich, M. Jenkinson, H. Johansen-Berg, R. G. Nunes, S. Clare, P. M. Matthews, J. M. Brady, and S. M. Smith. Characterization and propagation of uncertainty in diffusion-weighted MR imaging. *Magnetic Resonance in Medicine*, 50:1077–1088, 2003.
- [24] J. A. Bieszk and I. Fram. Automatic elastic image registration. In *Proceedings of Computers in Cardiology*, 1987.
- [25] D Le Bihan, E Breton, D Lallemand, P Grenier, E Cabanis, and M Laval-Jeantet. MR imaging of intravoxel incoherent motions: application to diffusion and perfusion in neurologic disorders. *Radiology*, 161(2):401–7, November 1986.
- [26] A. Bjorck. *Numerical methods for least squares problems*. SIAM, 1996.
- [27] L.G. Brown. A survey of image registration techniques. *ACM Computing Surveys*, 24(4):325–376, 1992.
- [28] Giovanni A Buonaccorsi, Caleb Roberts, Sue Cheung, Yvonne Watson, Karen Davies, Alan Jackson, Gordon C Jayson, and Geoff JM Parker. Comparing tracer kinetic model-driven registration to time series mean image registration for dynamic contrast enhanced MRI. In *MIUA*, pages 139–142, 2005.
- [29] Yan Cao, Michael I. Miller, Raimond L. Winslow, and Laurent Younes. Large deformation diffeomorphic metric mapping of vector fields. *IEEE Trans. Med. Imaging*, 24:1216–1230, 2005.
- [30] E. D. Castro and C. Morandi. Registration of translated and rotated images using finite fourier transforms. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 9(5):700– 703, September 1987.
- [31] M Cercignani, M Inglese, E Pagani, G Comi, and M. Filippi. Mean diffusivity and fractional anisotropy histograms of patients with multiple sclerosis. *American Journal of Neuroradiology*, 22:952–958, 2001.

- [32] Lin-Ching Chang, Derek K. Jones, and Carlo Pierpaoli. Restore: Robust estimation of tensors by outlier rejection. *Magnetic Resonance in Medicine*, 53:1088–1095, 2005.
- [33] O. Ciccarelli, A.T. Toosy, G.J. Parker, C.A. Wheeler-Kingshott, G.J. Barker, D.H. Miller, and A.J. Thompson. Diffusion tractography based group mapping of major white-matter pathways in the human brain. *NeuroImage*, 19:1545–1555, 2003.
- [34] TE Conturo, NF Lori, TS Cull, E Akbudak, AZ Snyder, JS Shimony, RC McKinstry, H Burton, and ME Raichle. Tracking neuronal fiber pathways in the living human brain. In *The Proceedings of the National Academy of Sciences Online (US)*, volume 96, pages 10422–10427, 1999.
- [35] P. A. Cook, Y. Bai, M. G. Hall, S. Nedjati-Gilani, K. K. Seunarine, and D. C. Alexander. Camino: Open-source diffusion-MRI reconstruction and processing. In *International Society for Magnetic Resonance in Medicine (ISMRM) 14th Scientific Meeting and Exhibition*, May 2006.
- [36] Philip A. Cook, Mark Symms, Philip A. Boulby, and Daniel C. Alexander. Optimal acquisition orders of diffusion-weighted MRI measurements. *Journal of Magnetic Resonance Imaging*, 25(5):1051–1058, 2007.
- [37] W R Crum, T Hartkens, and D L G Hill. Non-rigid image registration: theory and practice. *The British Institute of Radiology*, 77:S140—S153, 2004.
- [38] Kathleen M. Curran and Daniel C. Alexander. Orientation matching for registration of diffusion tensor images. In *SPIE Medical Imaging*, volume 5032, pages 149–156, San Diego, 2003.
- [39] L. S. Davis. Hierarchical generalized hough transform and line segment based generalized hough transforms. *Pattern Recognition*, 15:277–285, 1982.
- [40] M.H. Davis, A. Khotanzad, D.P. Flamig, and S.E. Harms. A physics-based coordinate transformation for 3-d image matching. *IEEE Transactions on Medical Imaging*, 6(3):317–328, 1997.
- [41] Mark Holden Derek L G Hill, Philipp G Batchelor and David J Hawkes. Medical image registration. *Physics in Medicine and Biology*, 46(R1–R45), 2001.

- [42] Qian Dong, Robert C. Welsh, Thomas L. Chenevert, Ruth C. Carlos, Pia Maly-Sundgren, Diana M. Gomez-Hassan, and Suresh K. Mukherji. Clinical applications of diffusion tensor imaging. *Journal of Magnetic Resonance Imaging*, 19(1):6–18, January 2004.
- [43] R.G. Nunes et al. Quantitative simulation of affine registration for correction of eddy current distortions in diffusion-weighted images. In *Proc. Int. Soc. of Magnetic Resonance in Medicine*, 2005.
- [44] A.C. Evans, C. Beil, S. Marrett, C.J. Thompson, and A. Hakin. Anatomical-functional correlation using an adjustable mri-based region of interest atlas with positron emission tomography,. *Journal of Cerebral Blood Flow and Metabolism*, 8:513–30, 1988.
- [45] TL Faber, RW McColl, RM Opperman, JR Corbett, and RM Peshock. Spatial and temporal registration of cardiac SPECT and MR images: methods and evaluation. *Radiology*, 179:857–61, 1991.
- [46] Martin A. Fischler and Robert C. Bolles. Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography. *Communications of the ACM*, 24(6):381–395, 1981.
- [47] Roger Fletcher. *Practical Methods of Optimization*. John Wiley & Sons, New York, second edition, 1987.
- [48] D. A. Forsyth and J. Ponce. *Computer Vision: A Modern Approach*. Prentice Hall, 2002.
- [49] K. J. Friston, J. Ashburner, C. D. Frith, J. B. Poline, J. D. Heather, and R. S. J. Frackowiak. Spatial registration and normalization of images. *Human Brain Mapping*, 3:165–189, 1995.
- [50] James C Gee and Daniel C Alexander. Diffusion-tensor image registration. *Visualization and image processing of tensor field*, 2005.
- [51] S Geman and DE McClure. Statistical methods for tomographic image reconstruction. *Bulletin of the International Statistical Institute*, 1987.
- [52] E. Gmur and H. Bunke. 3-d object recognition based on subgraph matching in polynomial time. *Structural Pattern Analysis*, 1990.
- [53] A. Goshtasby. Image registration by local approximation methods. *Image and Vision Computing*, 6(4):255–261, Nov. 1988.

- [54] Y. Le Guilloux. A matching algorithm for horizontal motion, application to tracking. In *Proceedings of the 8th International IEEE Conference on pattern Recognition*, 1986.
- [55] A. Guimond and C. R. G. Guttman. Deformable registration of DTMRI data based on transformation invariant tensor characteristics. In *the IEEE International Symposium on Biomedical Imaging (ISBI'2)*, Washington, DC. USA.
- [56] A. Guimond, A. Roche, N. Ayache, and J. Meunier. Multimodal brain warping using demons algorithm and adaptive intensity corrections. *IEEE Transactions On Medical Imaging*, 20:58–69, 2001.
- [57] J. V. Hajnal, S. Nadeem, E. J. Soar, A. Oatridge, I. R. Young, and G. M. Bydder. A registration and interpolation procedure for subvoxel matching of serially acquired MR images. *Journal of Computer Assisted Tomography*, 19(2):289–296, Mar.–Apr. 1995.
- [58] Joseph V. Hajnal, Derek L.G. Hill, and David J. Hawkes. *Medical Image Registration*. CRC Press, 2001.
- [59] J. C. Haselgrove and J. R. Moore. Correction for distortion of echo-planar images used to calculate the apparent diffusion coefficient. *Magnetic Resonance in Medicine*, 36:960–964, 1996.
- [60] John C. Haselgrove and James R. Moore. Correction for distortion of echo-planar images used to calculate the apparent diffusion coefficient. *Magnetic Resonance in Medicine*, 36(6):960–964, December 1996.
- [61] Paul Hayton, Michael Brady, Lionel Tarassenko, and Niall Moore. Analysis of dynamic MR breast images using a model of contrast enhancement. *Medical Image Analysis*, 1(3):207–224, 1996/7.
- [62] R. Mark Henkelman. Measurement of signal intensities in the presence of noise in MR images. *Medical Physics*, 12(2):232–233, 1985.
- [63] D. L. Hill, D. J. Hawkes, J. E. Crossman, M. J. Gleeson, T. C. Cox, and E. E. Bracey. Registration of MR and CT images for skull base surgery using point-like anatomical features. *The British Journal of Radiology*, 64(767):1030–1035, November 1991.
- [64] T. Hosey, G. Williams, and R. Ansorge. Inference of multiple fiber orientations in high angular resolution diffusion imaging. *Magnetic Resonance in Medicine*, 54:1480–1489, 2005.

- [65] R. HUMMEL and S. ZUCKER. On the foundations of relaxation labeling processes. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 1983.
- [66] Petra S Huppi, Stephan E. Maier, Sharon Peled, Gary P. Zientara, Patrick D. Barnes, Ferenc A. Jolesz, and Joseph J Volpe. Microstructural development of human newborn cerebral white matter assessed in vivo by diffusion tensor magnetic resonance imaging. *Pediatric Research*, 44(4):584–590, October 1998.
- [67] Basser P. J. and Pierpaoli C. Microstructural and physiological features of tissues elucidated by quantitative Diffusion Tensor MRI. *Journal of Magnetic Resonance Series*, 111:209–219, 1996.
- [68] K. M. Jansons and D. C. Alexander. Persistent angular structure: new insights from diffusion magnetic resonance imaging data. *Inverse Problems*, 19:1031–1046, 2003.
- [69] Mark Jenkinson, Peter Bannister, Michael Brady, and Stephen Smith. Improved optimization for the robust and accurate linear registration and motion correction of brain images. *NeuroImage*, 17:825–841, 2002.
- [70] Peter Jezzard, Alan S. Barnett, and Carlo Pierpaoli. Characterisation of and correction for eddy current artefacts in echo planar diffusion imaging. *Magnetic Resonance in Medicine*, 39:801–812, 1998.
- [71] Derek K. Jones and Peter J. Basser. Squashing peanuts and smashing pumpkins: How noise distorts diffusion-weighted MR data. *Magnetic Resonance in Medicine*, 52(5):979–993, 2004.
- [72] D.K. Jones, M.A. Horsefield, and A. Simmons. Optimal strategies for measuring diffusion in anisotropic systems by magnetic resonance imaging. *Magnetic Resonance in Medicine*, 42(3):515–525, 1999.
- [73] G. Q. Maure Jr., M. E. Noz, E. M. Lee, and J. H. Shimpf. Correlation methods for tomographic images using two and three dimensional techniques. In S. L. Bacharach, editor, *Proc. of the 9th Conference of Information Processing in Medical Imaging*, pages 266–279, Washington, D.C, June 1985.
- [74] K Krabbe, P Gideon, P Wagn, U Hansen, C Thomsen, and Madsen F. Mr diffusion imaging of human intracranial tumors. *Neuroradiology*, 39:483–489, 1997.

- [75] B.W. Kreher, J. F. Schneider, E. Martin, J. Henning, and K. A. Il'yasov. Effect of spatial resolution of the diffusion tensor images on the fiber tracks detectability. *Magnetic Resonance Materials in Physics, Biology and Medicine*, 15:280, 2002.
- [76] C.D. Kuglin and D.C. Hines. The phase correlation image alignment method. In *Conference on Cybernetics and Society*, pages 163–165, New York, September 1975. IEEE.
- [77] B. A. Landman, J. A. Farrell, S. Mori, P. C. van Zijl, and J. L. Prince. On the coregistration of diffusion weighted images. In *Proc.*, page 2987, Seattle, 2006. ISMRM, 14th Annual Meeting.
- [78] A. Leemans, C.J. Evans, and D.K. Jones. To rotate B or not to rotate B? the importance of reorienting the B-matrix during motion correction of DT-MRI data. In *Proc. Int. Soc. of Magnetic Resonance in Medicine*, page 140, 2008.
- [79] A. Leemans and D.K. Jones. The b-matrix must be rotated when correcting for subject motion in DTI data. *Magnetic Resonance in Medicine*, 61:1336–1349, 2009.
- [80] H. Lester and S. Arridge. A survey of hierarchical non-linear medical image registration. *Pattern Recognition*, 32(1):129–149, 1999.
- [81] DN Levin, CA Pelizzari, GT Chen, CT Chen, and MD Cooper. Retrospective geometric correlation of MR, CT, and PET images. *Radiology*, 169:817–23, 1988.
- [82] Schad LR, Boesecke R, Schlegel W, Hartmann GH, Sturm V, Strauss LG, and Lorenz WJ. Three dimensional image correlation of CT, MR, and PET studies in radiotherapy treatment planning of brain tumors. *Journal of Computer Assisted Tomography*, 11(6):948–54, Nov.-Dec. 1987.
- [83] Holden M, Hill DLG, Denton ERE, Jarosz JM, Cox TCS, Rohlfing T, Goodey J, and Hawkes DJ. Voxel similarity measures for 3D serial MR brain image registration. *IEEE Transactions on Medical Imaging*, 19:94–102, 2000.
- [84] F. Maes, A. Collignon, D. Vandermeulen, G. Marchal, and P. Suetens. Multimodality image registration by maximization of mutual information. *IEEE Transactions on Medical Imaging*, 16(2):187–198, 1997.
- [85] J. B. Maintz and M. A. Viergever. A survey of medical image registration. *Med. Image Anal.*, 2(1):1–36, 1998.

- [86] H. Maitre and Y. Wu. Improving dynamic programming to solve image registration. *Pattern Recognition*, 20(4):443–462, 1987.
- [87] L.E. Malvern. *Introduction to the Mechanics of a Continuous Medium*. Prentice-Hall, 1969.
- [88] J. F. Mangin, C. Poupon, C. Clark, D. Le Bihan, and I. Bloch. Distortion correction and robust tensor estimation for MR diffusion imaging. *Medical Image Analysis*, 6(3):191–198, September 2002.
- [89] Susana M. Maniega, Mark E. Bastin, and Paul A. Armitage. Effects of random subject rotation on optimised diffusion gradient sampling schemes in diffusion tensor MRI. *Magnetic Resonance Imaging*, 26:451–460, 2008.
- [90] P Mansfield. Multi-planar image formation using nmr spin echoes. *Journal of Physics C*, 10:L55–L58, 1977.
- [91] Bastin ME. Correction of eddy current-induced artifacts in diffusion tensor imaging using iterative cross-correlation. *Magnetic Resonance Imaging*, 17:1011–1024, 1999.
- [92] A. Melbourne, D. Atkinson, M.J. White, D. Collins, M. Leach, and D. Hawkes. Registration of dynamic contrast-enhanced MRI using a progressive principal component registration (PPCR). *Physics in Medicine and Biology*, 2007.
- [93] E. E. Milios. Shape matching using curvature processes. *Computer Vision, Graphics, and Image Processing*, 47:203–226, 1989.
- [94] D Miller, R Grossman, S Reingold, and H. Mcfarland. The role of magnetic resonance techniques in understanding and managing multiple sclerosis. *Brain*, 121:3–24, 1998.
- [95] O. Monga and S. Benayoun. Using partial derivatives of 3d images to extract typical surface features. *Vision Image Understanding*, 61:171–189, 1995.
- [96] D. S. Moore and G. P. McCabe. *Introduction to the Practice of Statistics*. New York: W. H. Freeman, 3 edition, 1999.
- [97] S Mori, BJ Crain, VP Chacko, and PCM. Zijl. Three dimensional tracking of axonal projections in the brain by magnetic resonance imaging. *Annals of Neurology*, 45:265–269, 1999.

- [98] S Mori, BJ Crain, and PC Zijl. 3d brain fiber reconstruction from diffusion MRI. In *Proceedings of International Conference on Functional Mapping of the Human Brain, Montreal, 1998*.
- [99] S. Mori, M. Kraut, P.C. van Zijl, M. Solaiyappan, H. Brem, and M. Pomper. Diffusion tensor imaging and fiber tracking of brain tumours. In *Proc. International Society Magnetic Resonance in Medicine*, 2001.
- [100] Susumu Mori and Peter C. M. van Zijl. Fiber tracking: principles and strategies - a technical review. *NMR in Biomedicine*, 15:468–480, 2002.
- [101] M Moseley, J Kucharczyk, and J Mintorovitch. Diffusionweighted mr imaging of acute stroke: correlation with t2-weighted and magnetic susceptibility-enhanced mr imaging in cats. *AJNR American Journal of Neuroradiology*, 11:423–429, 1990.
- [102] Thomas Netsch and Arianne van Muiswinkel. Quantitative evaluation of image-based distortion correction in diffusion tensor imaging. *IEEE Transactions on medical imaging*, 23(7):789–798, July 2004.
- [103] C. H. Neuman. Spin echo of spins diffusion in a bounded medium. *Journal of Chemical Physics*, 60:4508–4511, 1974.
- [104] Y. Ohta, K. Takano, and K. Ikeda. A high-speed stereo matching system based on dynamic programming. In *In Proceedings of the International Conference in Computer Vision*, 1987.
- [105] J.L. Ostuni and A.K. Santha. Analysis of interpolation effects in the reslicing of functional MR images. *Journal of computer assisted tomography*, 21(5):803–810, 1997.
- [106] G. Paar and W. G. Kropatsch. Hierarchical cooperation between numerical and symbolic image representations. *Structural Pattern Analysis*, 1990.
- [107] H. J. Park, M. Kubicki, M. E. Shenton, A. Guimond, R. W. McCarley, S. E. Maier, R. Kikinis, F. A. Jolesz, and C.-F. Westin. Spatial normalization of diffusion tensor MRI using multiple channels. *NeuroImage*, 20(4):1995–2009, 2003.
- [108] GJ Parker. Tracing fiber tracts using fast marching. In *Proceedings of International Society of Magnetic Resonance in Medicine*, page 85, 2000.
- [109] C. Pierpaoli, P. Jezzard, P. J. Basser, A. Barnett, and G. Di Chiro. Diffusion Tensor MR imaging of the human brain. *Radiology*, 201(3):637–648, 1996.

- [110] William H. Press, Brian P. Flannery, Saul A. Teukolsky, and William T. Vetterling. *Numerical Recipes in C : The Art of Scientific Computing*. Cambridge University Press, October 1992.
- [111] K. E. PRICE. Relaxation matching techniques - a comparison. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 1985.
- [112] William S. Price. Pulsed-field gradient nuclear magnetic resonance as a tool for studying translational diffusion: Part 1. basic theory. *Concepts in Magnetic Resonance*, 9(5):299–336, 1997.
- [113] Ajit Rajwade, Arunava Banerjee, and Anand Rangarajan. New method of probability density estimation with application to mutual information based image registration. In *CVPR '06: Proceedings of the 2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 1769–1776, Washington, DC, USA, 2006. IEEE Computer Society.
- [114] S. RANADE and A. ROSENFELD. Point pattern matching by relaxation. *Pattern Recognition*, 12:269–275, 1980.
- [115] Ananth Ranganathan. The levenberg-marquardt algorithm, 2004.
- [116] Fazlollah M. Reza. *An Introduction to Information Theory*. New York, 1994.
- [117] P Ricci, J Burdette, A Elster, and D. Reboussin. A comparison of fast spin echo, fluid-attenuated inversion-recovery, and diffusionweighted mr imaging in the first 10 days after cerebral infarction. *American Journal of Neuroradiology*, 20:1535–1542, 1999.
- [118] Peter A. Rinck. *Magnetic resonance in medicine : the basic textbook of the European Magnetic Resonance Forum*. Blackwell Scientific, 3rd edition, 1993.
- [119] G. K. Rohde, A. S. Barnett, P. J. Basser, S. Marengo, and C. Pierpaoli. Comprehensive approach for correction of motion and distortion in diffusion-weighted MRI. *Magnetic Resonance in Medicine*, 51:103–114, 2004.
- [120] Pajevic S. and Pierpaoli C. Color schemes to represent the orientation of anisotropic tissues from diffusion tensor data: application to white matter fiber tract mapping in the human brain. *Magnetic Resonance Imaging*, 42:526–540, 1999.
- [121] A. Sanfeliu. Matching complex structures: The cyclic-tree representation scheme. *Structural Pattern Analysis*, 1990.

- [122] K. K. Seunarine and D. C. Alexander. *Diffusion MRI: From quantitative measurement to in-vivo neuroanatomy*, chapter Multiple Fibres: Beyond the diffusion tensor. Elsevier, 2009.
- [123] R. Shams, R.A. Kennedy, P. Sadeghi, and R. Hartley. Gradient intensity-based registration of multi-modal images of the brain. In *IEEE 11th International Conference on Computer Vision*, 2007.
- [124] L. G. SHAPIRO and R. M. HARALICK. Matching relational structures using discrete relaxation. *Syntactic and Structural Pattern Recognition: Theory and Applications*, 1990.
- [125] S.M. Smith, M. Jenkinson, M.W. Woolrich, C.F. Beckmann, T.E.J. Behrens, H. Johansen-Berg, P.R. Bannister, M. De Luca, I. Drobnjak, D.E. Flitney, R. Niazy, J. Saunders, J. Vickers, Y. Zhang, N. De Stefano, J.M. Brady, , and P.M. Matthews. Advances in functional and structural MR image analysis and implementation as FSL. *NeuroImage*, 23(S1):208–219, 2004.
- [126] Jan A. Snyman. *Practical Mathematical Optimization: An Introduction to Basic Optimization Theory and Classical and New Gradient-Based Algorithms*. Number ISBN 0-387-24348-8. Springer, 2005.
- [127] E. O. Stejskal and J. E. Tanner. Spin diffusion measurements: Spin echoes in the presence of a time-dependent field gradient. *The Journal of Chemical Physics*, 42(1):288–292, 1965.
- [128] C. Studholme, D.L.G. Hill, and D.J. Hawkes. An overlap invariant entropy measure of 3d medical image alignment. *Pattern Recognition*, 32:71–86, 1999.
- [129] Edith V. Sullivan, Elfar Adalsteinsson, Maj Hedehus, Catherine Ju, Michael Moseley, Kelvin O. Lim, and Adolf Pfefferbaum. Equivalent disruption of regional white matter microstructure in ageing healthy men and women. *NeuroReport*, 12(1):99–104, January 2001.
- [130] P M Thompson, J N Gledd, R P Woods, D MacDonald, A C Evans, and A W Toga. Growth patterns in the developing brain detected by using continuum mechanics tensor maps. *Nature*, pages 190–3, 2000.
- [131] A Tievsky, T Ptak, and Farkas J. Investigation of apparent diffusion coefficient and diffusion tensor anisotropy in acute and chronic multiple sclerosis lesions. *American Journal of Neuroradiology*, 20:1491–1499, 1999.

- [132] Peters TM, Clark J A, Olivier A, Marchand E P, Mawko G, Dieumegarde M, Muresan L, and Ethier R. Integrated stereotaxic imaging with CT, MR imaging, and digital subtraction angiography. *Radiology*, 161(3):821–6, December 1986.
- [133] J.-Donald Tournier, Fernando Calamante, David G. Gadian, and Alan Connelly. Direct estimation of the fiber orientation density function from diffusionweighted MRI data using spherical deconvolution. *NeuroImage*, 23:1176–1185, 2004.
- [134] D. S. Tuch. *Diffusion MRI of complex tissue structure*. Doctor of philosophy in biomedical imaging, Massachusetts Institute of Technology, 2002.
- [135] David S. Tuch. Q-ball imaging. *Magnetic Resonance in Medicine*, 52:1358–1372, 2004.
- [136] DS Tuch, MR Wiegell, TG Reese, JW Belliveau, and V Wedeen. Measuring cortico-cortical connectivity matrices with diffusion spectrum imaging. In *Proceedings of International Society of Magnetic Resonance in Medicine*, page 502, 2001.
- [137] S Warach, D Chien, W Li, M Ronthal, and R. Edelman. Fast magnetic resonance diffusion-weighted imaging of acute human stroke. *Neurology*, 42:1717–1723, 1992.
- [138] V. J. Wedeen, T. G. Reese, D. S. Tuch, J.-G. Dou, R. M. Weiskoff, and D. Chessler. Mapping fiber orientation spectra in cerebral white matter with fourier-transform diffusion MRI. In *7th Annual Meeting of the ISMRM*, page 321, 1999.
- [139] J B West, J M Fitzpatrick, M Y Wang, B M Dawant, C R Jr Maurer, R M Kessler, and R J Maciunas. Retrospective intermodality techniques for images of the head: surface-based versus volume-based. *IEEE Transactions on Medical Imaging*, 18:144–150, 1999.
- [140] Jay et. al. West. Comparison and evaluation of retrospective intermodality brain image registration techniques. *Journal of Computer Assisted Tomography*, 21(4):554–568, 1997.
- [141] III William M. Wells, Paul Viola, Hideki Atsumi, Shin Nakajima, and Ron Kikinis. Multi-modal volume registration by maximization of mutual information. *Medical Image Analysis*, 1(1):35–51, March 1996.
- [142] R. P. Woods, S. T. Grafton, C. J. Holmes, S. R. Cherry, and J. C. Mazziotta. Automated image registration: I. general methods and intrasubject, intramodality validation. *Journal of Computer Assisted Tomography*, 22:139–152, 1998.

- [143] Roger P. Woods, Scott T. Grafton, Colin J. Holmes, Simon R. Cherry, and John C. Mazziotta. Automated image registration: I. general methods and intrasubject, intramodality validation. *Journal of Computer Assisted Tomography*, 22(1):139–152, 1998.
- [144] Dongrong Xu, Susumu Mori, Dinggang Shen, Peter C.M. van Zijl, and Christos Davatzikos. Spatial normalization of diffusion tensor fields. *Magnetic Resonance in Medicine*, 50(1):175–182, Jun. 2003.
- [145] R. Xue, P.C.M. Zijl, B.J. Crain, M. Solaiyappan, and S. Mori. In vivo three-dimensional reconstruction of rat brain axonal projections by diffusion tensor imaging. *Magnetic Resonance in Medicine*, 42:1123–1127, 1999.
- [146] Hui Zhang, Paul A. Yushkevich, Daniel C. Alexander, and James C. Gee. Deformable registration of diffusion tensor MR images with explicit orientation optimization. *Medical Image Analysis*, 10(5):764–785, 2006.
- [147] B. Zitova and J. Flusser. Image registration methods: A survey. *Image and Vision Computing*, 21(11):977–1000, 2003.