



HAL
open science

High Angular Resolution Diffusion MRI: from Local Estimation to Segmentation and Tractography

Maxime Descoteaux

► **To cite this version:**

Maxime Descoteaux. High Angular Resolution Diffusion MRI: from Local Estimation to Segmentation and Tractography. Human-Computer Interaction [cs.HC]. Université Nice Sophia Antipolis, 2008. English. NNT: . tel-00457458

HAL Id: tel-00457458

<https://theses.hal.science/tel-00457458>

Submitted on 17 Feb 2010

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

PhD THESIS

prepared at
INRIA Sophia Antipolis

and presented at the
University of Nice-Sophia Antipolis
Graduate School of Information and Communication Sciences

*A dissertation submitted in partial fulfillment
of the requirements for the degree of*

DOCTOR OF SCIENCE
Specialized in Control, Signal and Image Processing

High Angular Resolution Diffusion MRI: from Local Estimation to Segmentation and Tractography

Maxime DESCOTEAUX

Adviser	Rachid Deriche	INRIA Sophia Antipolis, France
Reviewers	Peter Basser	NICHD, USA
	Cyril Poupon	NeuroSpin / CEA, France
	Carl-Fredrik Westin	Harvard University, USA
Examiners	Nicholas Ayache	INRIA Sophia Antipolis, France
	Habib Benali	INSERM / Pitié-Salpêtrière, France
	Denis Le Bihan	NeuroSpin / CEA, France
Invited members	Alfred Anwander	Max Planck Institute, Germany
	Olivier Faugeras	INRIA Sophia Antipolis, France
	Stéphane Lehéricy	Pitié-Salpêtrière, France

UNIVERSITÉ NICE-SOPHIA ANTIPOLIS - UFR Sciences

École Doctorale STIC

(Sciences et Technologies de l'Information et de la Communication)

THÈSE

pour obtenir le titre de
DOCTEUR EN SCIENCES

de l'UNIVERSITÉ de Nice-Sophia Antipolis

Discipline: Automatique, Traitement du Signal et des Images

présentée et soutenue par

Maxime DESCOTEAUX

IRM de Diffusion à Haute Résolution Angulaire: Estimation Locale, Segmentation et Suivi de Fibres

Thèse dirigée par Rachid DERICHE

Date prévue de soutenance, 5 février 2008

Composition du jury:

<i>Rapporteurs</i>	Peter Basser	NICHD, USA
	Cyril Poupon	NeuroSpin / CEA, France
	Carl-Fredrik Westin	Harvard University, USA
<i>Examineurs</i>	Nicholas Ayache	INRIA Sophia Antipolis, France
	Habib Benali	INSERM / Pitié-Salpêtrière, France
	Denis Le Bihan	NeuroSpin / CEA, France
<i>Membres invités</i>	Alfred Anwander	Max Planck Institute, Germany
	Olivier Faugeras	INRIA Sophia Antipolis, France
	Stéphane Lehéricy	Pitié-Salpêtrière, France

Contents

Contents	v
List of Figures	x
List of Tables	xi
List of Symbols	xiii
Abstract	xv
Résumé	xvii
Acknowledgments	xix
I Introduction	1
1 Introduction	3
2 Introduction (français)	11
II Background	21
3 Neural Tissue and Human Brain White Matter	23
3.1 The Human Brain and Neural Tissue	24
3.2 Organization of the White Matter	27
3.2.1 Important white matter bundles for this thesis	30
3.3 Conclusion	33
4 Principles of Diffusion MRI: Going Beyond the Diffusion Tensor	35
4.1 Basic Principles of Diffusion MRI	36

4.1.1	Brownian Motion of Water Molecules and Fick's Law	37
4.2	NMR and Diffusion MRI	39
4.2.1	Pulse Gradient Spin Echo (PGSE)	40
4.2.2	Diffusion-Weighted Imaging (DWI)	42
4.2.3	Apparent Diffusion Coefficient (ADC) and Trace Imaging	43
4.2.4	Diffusion Tensor Imaging (DTI)	44
4.2.5	High Angular Resolution Diffusion Imaging (HARDI)	46
4.3	Multiple Fiber HARDI Reconstruction Techniques	50
4.3.1	Diffusion Spectrum Imaging (DSI)	51
4.3.2	Single Shell HARD Imaging	53
4.4	Conclusion	60
5	Mathematics on the Sphere: The Spherical Harmonics	61
5.1	Solving the Laplace's Equation	62
5.2	Spherical Harmonics	64
5.3	The Modified Spherical Harmonics Basis	65
5.4	Spherical Function Estimation with Spherical Harmonics	66
5.5	The Funk-Hecke Theorem	68
5.6	Properties of Spherical Harmonics Used in this Thesis	69
III	Methods	71
6	Apparent Diffusion Coefficient Estimation and Applications	73
6.1	Introduction	74
6.2	ADC Profile Estimation from HARDI	76
6.2.1	Fitting the ADC with the Spherical Harmonics	76
6.2.2	Fitting the ADC with a High Order Diffusion Tensor (HODT)	77
6.2.3	A Regularization Algorithm for HARDI Signal/ADC Estimation	78
6.2.4	From SH Coefficients to HODT Coefficients	79
6.3	High Order Anisotropy Measures from ADC Profiles	83
6.3.1	Frank and Chen <i>et al</i> Measures	83
6.3.2	Alexander <i>et al</i> Measure	84
6.3.3	Generalized Anisotropy Measure	84
6.4	Results	85
6.4.1	Optimal Regularization Parameter with the L-Curve Method	85
6.4.2	High Order Anisotropy Measures Results	89
6.4.3	Biological Phantom Results	96
6.4.4	Human Brain HARDI Data Results	96
6.5	Discussion	98
6.5.1	Contributions of this chapter:	100
6.6	Appendix A: Independent Elements of the HODT	102

6.7	Appendix B: Spherical Harmonics to the High Order Tensor	103
7	Analytical Q-Ball Imaging	105
7.1	Introduction	107
7.2	Q-Ball Imaging	108
7.3	Analytical Q-Ball Imaging	109
7.4	Validation and Comparison of QBI methods	114
7.4.1	Numerical QBI Implementation	114
7.4.2	Synthetic Data Generation	114
7.4.3	Computational Complexity Analysis	115
7.4.4	Effect of Spherical Sampling Density	115
7.4.5	Robustness to Noise	115
7.4.6	ODF Shape Comparison	116
7.4.7	Fiber Detection and Angular Resolution	116
7.4.8	Power Spectrum of the Spherical Harmonics Representation . . .	117
7.5	Results	117
7.5.1	Running Time Comparison	117
7.5.2	ODF Shape Comparison and Robustness to Noise	118
7.5.3	ODF Reconstruction for Different Sampling Schemes	120
7.5.4	Fiber Detection and Angular Resolution	121
7.5.5	Power Spectrum of the Spherical Harmonic Representation . . .	127
7.5.6	Biological Phantom Results	129
7.5.7	Human Brain HARDI Data Results	131
7.6	Discussion	134
7.6.1	Contributions of this chapter:	136
7.7	Appendix A: Funk-Radon transform and the diffusion ODF	138
7.8	Appendix B: Rigorous Proof of the Analytical QBI Solution	140
7.9	Appendix C: Exact ODF From the Multiple-Tensor Model	141
8	Segmentation in High Angular Resolution Diffusion MRI	143
8.1	Introduction	145
8.2	Background on Segmentation	145
8.3	Statistical Surface Evolution	148
8.3.1	Distances between ODFs	148
8.3.2	Segmentation by Surface Evolution	150
8.4	Segmentation Results & Discussion	153
8.4.1	Synthetic Simulation Results	153
8.4.2	Biological Phantom Results	155
8.4.3	Human Brain HARDI Data Results	156
8.4.4	Multi-Subject Study on a public HARDI Database	159
8.5	Discussion	162
8.5.1	Contributions of this chapter:	164

9	Improving Q-Ball Imaging: from Diffusion ODF to Fiber ODF estimate	165
9.1	Introduction	167
9.2	Spherical Deconvolution Using Spherical Harmonics	167
9.2.1	Spherical Deconvolution of the q-Ball diffusion ODF	168
9.2.2	Spherical Deconvolution of the raw HARDI Signal	170
9.3	Evaluation of the Spherical Deconvolution Methods	178
9.3.1	Synthetic Data Experiment	178
9.3.2	Real Data Experiment	179
9.4	Results	179
9.4.1	Effect of Varying Signal Generation Parameters	180
9.4.2	Fiber Detection and Angular Resolution	184
9.4.3	Robustness to Noise	186
9.4.4	Real Data Spherical Deconvolution Reconstructions	188
9.5	Discussion	194
9.5.1	Contributions of this chapter:	197
9.6	Appendix A: Diffusion ODF Kernel for Sharpening	198
9.7	Appendix B: Coefficients of A in the FORECAST Solution	199
9.8	Appendix C: Relation between the fODF and FORECAST	200
10	Tractography in High Angular Resolution Diffusion MRI	205
10.1	Introduction	207
10.2	Background on Tractography	208
10.3	Tractography	210
10.3.1	Deterministic Multidirectional ODF Tracking	210
10.3.2	Probabilistic fiber ODF Tracking	212
10.4	Tractography Results	213
10.4.1	Synthetic Simulation Results	214
10.4.2	Biological Phantom Results	215
10.4.3	Human Brain HARDI Data Results	216
10.5	Discussion	223
10.5.1	Contributions of this chapter:	226
IV	Application	227
11	Quantitative Assessment of Transcallosal Fibers	229
11.1	Introduction	230
11.2	Quantifying Lateral Projections of the Corpus Callosum	233
11.2.1	Methods	233
11.2.2	Results	234
11.3	Discussion	235

V Conclusion	237
12 Conclusion	239
13 Conclusion (français)	249
VI Appendix	257
A Synthetic HARDI Data Generation and HARDI Acquisitions	259
A.1 Synthetic Data Generation	259
A.2 Biological Phantom Data	261
A.3 Human Brain Data	261
A.3.1 Data from McGill University, Montréal, Canada	261
A.3.2 Data from the Minnesota University, Minneapolis, USA	262
A.3.3 Data from the Max Planck Institute (MPI), Leipzig, Germany . .	262
A.3.4 Public HARDI Database from NeuroSpin / CEA, Paris, France . .	263
B Publications of the Author Arising from this Work	265
Bibliography	269

List of Figures

1.1	Sketch of the chapters of the thesis	5
1.2	Graphical sketch of the contributions of the thesis	6
2.1	Organisation des chapitres de cette thèse	13
2.2	Contributions de cette thèse	14
3.1	Human brain white and gray matter	25
3.2	The four lobes of the cerebral cortex and the cerebellum	26
3.3	The typical neuron or axonal nerve fiber	27
3.4	The projections fibers	28
3.5	The association fibers	29
3.6	The corpus callosum and its commissural fibers	31
3.7	The corticospinal tract	32
3.8	The corona radiata	32
3.9	The superior and inferior longitudinal fasciculus	33
3.10	The commissural fibers and the anterior commissure	33
4.1	Brownian motion	37
4.2	Brief history of diffusion MRI.	40
4.3	Stejskal-Tanner imaging PGSE sequence	41
4.4	Diffusion-weighted images at different b -values	43
4.5	Diffusion-weighted images with two diffusion gradient directions	44
4.6	Diffusion tensor representation	45
4.7	Field of diffusion tensors and principal diffusion direction	47
4.8	Limitation of DTI in voxels with crossing configurations	48
4.9	Diffusion propagator in isotropic, single fiber and crossing fibers voxels	49
4.10	Simplifying visualization of the diffusion propagator	50
4.11	Sketch of diffusion MRI acquisitions and reconstructions methods	52
4.12	Tessellation of the sphere	53
4.13	The ADC and ODF in a two crossing fibers configuration	54
4.14	Illustration of q-ball imaging and the Funk-Radon transform	57
4.15	q-ball diffusion ODF reconstruction in an axial slice	58

5.1	Spherical harmonics basis up to order 3	65
5.2	Modified spherical harmonics basis up to order 4	67
6.1	Crossing fiber distribution and ADC profiles from DTI and HARDI	75
6.2	Transform from the HODT to the standard SH basis	80
6.3	Transform from the HODT to the modified SH basis	81
6.4	Sketch of the ADC regularization algorithm	82
6.5	L-curves for 1, 2 and 3 fibers with optimal parameter	88
6.6	Synthetic ADC profiles for different estimation algorithms	91
6.7	Two orthogonal fiber crossing synthetic data example	93
6.8	Three orthogonal fiber crossing synthetic data example	94
6.9	ADC reconstruction results on a biological phantom	95
6.10	ADC reconstruction results on real HARDI data	97
6.11	High order anisotropy results on a ROI from the corpus callosum	98
7.1	Laplace-Beltrami regularization reduces ODF estimation errors	119
7.2	Analytical QBI of Hess et al. (2006) for different sampling schemes	121
7.3	Analytical QBI for different spherical sampling schemes and noise level	122
7.4	ODF regularization decreases small perturbations due to noise	123
7.5	ODFs recover multiple fiber crossing in synthetic data	125
7.6	Power spectrum of the spherical harmonics at low b -values	127
7.7	Power spectrum of the spherical harmonics at high b -values	128
7.8	ODFs recover multiple fiber crossing in the biological phantom.	130
7.9	ODFs for the analytical and numerical QBI are qualitatively the same	131
7.10	ODFs recover multiple fiber crossings on a human brain dataset	132
7.11	ODF crossings between the genu of the CC and capsule fibers	133
7.12	ODF crossings in the centrum semiovale	134
7.13	ODF crossings and fanning in the corona radiata	134
7.14	Zeroth-order Bessel function smoothing of the q-ball ODF	140
8.1	Illustration of the distance measure between ODFs	149
8.2	Sketch of distance between crossing and single peaked ODF	150
8.3	Spherical harmonic coefficients in the corpus callosum	151
8.4	Spherical harmonic coefficients in the cortico-spinal tract	152
8.5	Segmentation of the 90° crossing synthetic example	154
8.6	Time evolution of the segmentation flow on the 90° crossing	154
8.7	Segmentation on a synthetic branching example	155
8.8	Segmentation on a synthetic 'U'-fiber example	156
8.9	Segmentation on a biological phantom	157
8.10	Time evolution of the segmentation flow on the biological phantom	157
8.11	Segmentation of the corpus callosum	158
8.12	Segmentation of the cortico-spinal tract	159

8.13	Segmentation of the corpus callosum on the HARDI database	160
8.14	Segmentation of the cortico-spinal tract on the HARDI database	161
9.1	Spherical deconvolution sketch of the HARDI signal	168
9.2	Spherical deconvolution sketch of the diffusion ODF	169
9.3	Regularization incorporated in the FORECAST method	172
9.4	Regularization incorporated in the spherical deconvolution methods . .	175
9.5	Filtered, constrained and super-resolution spherical deconvolution . . .	177
9.6	Effect of varying estimation order on the sharp fiber ODF	180
9.7	Effect of varying SNR on the sharp fiber ODF	181
9.8	Effect of varying estimation separation angle on the sharp fiber ODF . .	182
9.9	Effect of varying estimation volume fraction on the sharp fiber ODF . .	183
9.10	Effect of varying SNR on the detected maxima of the sharp fiber ODF .	186
9.11	Sharp fiber ODF crossings on the biological phantom.	187
9.12	Sharp fiber ODF crossings on the MPI dataset	189
9.13	Sharp fiber ODF crossings in a region of interest of the MPI dataset . .	190
9.14	Sharp fiber ODF crossings on the BIC dataset	191
9.15	Sharp fiber ODF crossing between the genu of the CC and capsule fibers	192
9.16	Sharp fiber ODF crossings in the CMRR dataset	193
9.17	Sharp fiber ODF versus FORECAST	202
9.18	Sharp fiber ODF versus FORECAST represented in SH coefficients . . .	203
10.1	Deterministic tracking in a synthetic branching example	214
10.2	Probabilistic tracking in a synthetic branching example	214
10.3	Deterministic tracking in the biological phantom	215
10.4	Tracking of the superior longitudinal fasciculus	216
10.5	Deterministic and probabilistic tracking of the thalamic radiations . . .	217
10.6	Tracking recovering a known fanning/crossing region	218
10.7	Tracking of the anterior commissure fibers	219
10.8	Deterministic and probabilistic tracking of the tapetum.	220
10.9	Tracking of the projections of the corpus callosum to Broca's area	221
10.10	Tracking comparison of different fiber ODFs in the corpus callosum . . .	222
11.1	DTI tracking misses transcallosal fibers to the lateral cortex	231
11.2	DTI tracking is blocked by the corona radiata and SLF	231
11.3	Diffusion and fiber ODFs in a region with crossing transcallosal fibers .	232
11.4	Fiber ODF tracking of the transcallosal fibers	233
11.5	Transcallosal fibers projecting to ventral and lateral parts of the cortex	234
11.6	Quantification of lateral projections of the corpus callosum	235
12.1	Sub-voxel fiber configurations and their associated diffusion ODF	242
12.2	Schematic tracking in ambiguous sub-voxel configurations	243
12.3	HARDI clustering, tracking and segmentation	247

List of Tables

4.1	Diffusion MRI acquisition techniques	51
4.2	Diffusion MRI acquisition techniques.	60
6.1	Summary of each algorithm used to compare HODT estimation	87
6.2	Laplace-Beltrami regularization reduces the estimation error	90
6.3	Mean generalized anisotropy measure on simulated ADC profiles	91
6.4	Regularization improves classification from anisotropy measures	92
7.1	Summary of the regularized, fast and robust analytical QBI algorithm.	113
7.2	Running times comparison between numerical and analytical QBI	118
7.3	Laplace-Beltrami regularization improves fiber detection	124
7.4	Angular resolution of analytical QBI with/without regularization	126
7.5	Quantitative fiber detection on the biological phantom	131
9.1	The fiber ODF improves fiber detection and angular resolution of QBI	185
9.2	Fraction of crossing voxels in the human brain white matter	192
9.3	Mean difference between our sharp fiber ODF and FORECAST	201

List of Symbols

3D	3 dimensional
5D	5 dimensional
AC	Anterior Commissure
ADC	Apparent Diffusion Coefficient
ATR	Anterior Thalamic Radiations
BIC	Brain Imaging Center
CC	Corpus Callosum
Cg	Cingulum
CHARMED	Composite Hindered And Restricted ModEl of Diffusion
CNS	Central Nervous System
CPT	CorticoPontine Tract
CR	Corona radiata
CRE	Cumulative Residual Entropy
CSF	Cerebral Spinal Fluid
CST	Corticospinal Tract
dMRI	Diffusion MRI
DOT	Diffusion Orientation Transform
DSI	Diffusion Spectrum Imaging
DTI	Diffusion Tensor Imaging
DW	Diffusion-Weighted
DWI	DW Images
FA	Fractional Anisotropy
FMI	Fractional Multifiber Index
FOD	Fiber Orientation Distribution
FRT	Funk-Radon Transform
GA	Generalized Anisotropy
gDTI	generalized DTI
GFA	Generalized Fractional Anisotropy
HARDI	High Angular Resolution Diffusion Imaging
HOT	High Order Tensor
HODT	High Order Diffusion Tensor

IFO	Inferior Fronto-Occipital fasciculus
ILF	Inferior Longitudinal Fasciculus
LB	Laplace-Beltrami
LR	Linear Regression
NMR	Nuclear Magnetic Resonance
MRI	Magnetic Resonance Imaging
ODF	Orientation Distribution Function
PAS	Persistent Angular Structure
PC	Posterior Commissure
PDE	Partial Differential Equation
PDF	Probability Density Function
PGSE	Pulse Gradient Spin Echo
PNS	Peripheral Nervous System
PTR	Posterior Thalamic Radiations
QBI	Q-Ball Imaging
QSI	Q-Space Imaging
RA	Relative Anisotropy
ROI	Region Of Interest
SD	Spherical Deconvolution
SFO	Superior Fronto-Occipital fasciculus
SH	Spherical Harmonics
SHDE	Spherical Harmonics Differential Equation
SHPDE	Spherical Harmonic PDE
SLF	Superior Longitudinal Fasciculus
SNR	Signal to Noise Ratio
STR	Posterior Thalamic Radiations
TAP	Tapetum
TR	Thalamic Radiations
UNC	Uncinate fasciculus
vMF	von Mises-Fisher

Abstract

At the current resolution of diffusion-weighted (DW) magnetic resonance imaging (MRI), research groups agree that there are between one third to two thirds of imaging voxels in the human brain white matter that contain fiber crossing bundles. This thesis tackles the important problem of recovering crossing fiber bundles from DW-MRI measurements. The main goal is to overcome the limitations of diffusion tensor imaging (DTI). It is well-known that imaging voxels where there are multiple fiber crossings produce a non-Gaussian DW signal. This is precisely where DTI is limited due to the intrinsic Gaussian assumption of the technique. Hence, this thesis is dedicated to the development of local reconstruction methods, segmentation and tractography algorithms able to infer multiple fiber crossing from DW-MRI data. To do so, high angular resolution diffusion imaging (HARDI) is used to measure DW images along several directions. Q-ball imaging (QBI) is a recent such HARDI technique that reconstructs the diffusion orientation distribution function (ODF), a spherical function that has its maxima aligned with the underlying fiber directions at every voxel. QBI and the diffusion ODF will play a central role in this thesis.

There are many original contributions in this thesis. First, we propose a robust estimation of the HARDI signal using a closed-form regularization algorithm based on the spherical harmonics. Then, we estimate the apparent coefficient coefficient (ADC) to study HARDI anisotropy measures and to discriminate voxels with underlying isotropic, single fiber and multiple fiber distributions. Next, we develop a linear, robust and analytical QBI solution using the spherical harmonic basis, which is used in a new statistical region-based active contour algorithm to segment important white matter fiber bundles. In addition, we develop a new spherical deconvolution sharpening method that transforms the diffusion q-ball ODF into a fiber ODF. Finally, we propose a new deterministic tractography algorithm and a new probabilistic tractography algorithm exploiting the full distribution of the fiber ODF. Overall, we show local reconstruction, segmentation and tracking results on complex fiber regions with known fiber crossing on simulated HARDI data, on a biological phantom and on multiple human brain datasets. Most current DTI based methods neglect these complex fibers, which might lead to wrong interpretations of the brain anatomy and functioning.

Résumé

La résolution actuelle des IRM pondérées en diffusion suggère qu'il y a entre un et deux tiers des voxels de la matière blanche qui contiennent plusieurs faisceaux de fibres qui se croisent. L'IRM par tenseur de diffusion (DTI) classique est intrinsèquement limitée à ces endroits par son hypothèse de base qu'un seul faisceau traverse chaque voxel de l'image. Le but de cette thèse est donc de développer des techniques d'IRM à haute résolution angulaire (HARDI) pour pouvoir retrouver une ou plusieurs fibres et surmonter aux limites DTI. L'imagerie par q-ball (QBI) est une technique récente qui permet d'estimer la distribution d'orientation des fibres (ODF). La technique de QBI ainsi que l'ODF de diffusion des fibres permettent de retrouver plusieurs directions de fibres en chaque voxel de l'image. Ceux-ci joueront donc un rôle important tout au long de ce document.

Cette thèse propose plusieurs contributions originales. D'abord, nous développons l'estimation robuste du signal HARDI en utilisant une base modifiée d'harmoniques sphériques et un terme de régularisation. Ensuite, nous proposons la modélisation du coefficient de diffusion apparent (ADC) pour étudier les mesures d'anisotropie HARDI et faire la classification des voxels contenant une distribution isotrope, une distribution d'une seule population de fibres et une distribution de plusieurs faisceaux fibres. Nous proposons de plus, le développement d'une solution analytique pour estimer l'ODF de diffusion en QBI ainsi qu'un nouvel algorithme de segmentation de ces images d'ODF obtenues par le QBI. Nous présentons également le calcul de l'ODF de fibres avec une nouvelle méthode de déconvolution sphérique à partir des données QBI. Enfin, nous développons de nouveaux algorithmes de suivi de fibres (tracking) déterministes et probabilistes à partir de l'ODF du q-ball et l'ODF de fibres. Finalement, tous ces nouveaux algorithmes sont testés et appliqués sur des données HARDI simulées, sur un fantôme biologique et sur des données réelles de plusieurs sujets dans des régions complexes avec plusieurs faisceaux qui se croisent. Nos résultats illustrent clairement la valeur ajoutée de nos méthodes HARDI sur la plupart des méthodes courantes en DTI qui négligent ces faisceaux complexes, ce qui peut amener à une mauvaise analyse et interprétation de l'anatomie et des fonctions cérébrales.

Acknowledgments

First and foremost, I would like to thank my supervisor, Prof. Rachid Deriche. He has been a constant presence throughout my three years in France and was able to keep me motivated, challenged and productive at all time. His scientific assistance and moral support go well beyond the expectation of a PhD supervisor. He has given me all the tools to continue in a academic career. His social and human skills have helped our relationship a lot. Not only has he been a great scientific advisor but he has also become a good friend, with whom I enjoy good wine, good Czech beers and BBQs.

I also want to acknowledge and thank Peter Basser, Cyril Poupon and Carl-Fredrik Westin for their careful and constructive review. Your comments and suggestions have contributed to clarify certain aspects and strengthen the manuscript.

I am thankful to my student mentor, Christophe Lenglet. I had the privilege to spend my first year by his side. He has thought me a great deal, advised me well at the beginning of my thesis and has left me a lot of bits and pieces of C++ code. I am also indebted to, Martin de la Gorce, Elaine Angelino and Shaun Fitzgibbons. They were internship students before and during the start of my PhD. They have contributed in some parts of this thesis and I am very grateful. Their names naturally appear on some of our publications. Lastly, thanks to my new partners in the team, Demian Wassermann and Aurobrata Ghosh. You guys are great and I love working with you. The Odyssée diffusion MRI group rocks!

A warm thanks to all people that have contributed in the methodological developments of algorithms in this thesis and that have advised me in my three years of PhD. First, a special thanks goes to my ex-research group at McGill University, Peter Savadjiev, Jennifer Campbell and Kaleem Siddiqi. It was great to keep working with them. A very special “thank you” to Jennifer and Kaleem who also has always been mentors and great advisors in my research in diffusion MRI. Jennifer’s biological phantom was also a great help throughout my thesis. Second, I am also grateful to Alfred Anwander of the Max Planck Institute in Leipzig, Germany. Thank you so much Alfred! It was a great time working with you in the past year. I believe that we have established a great relationship that will last longer and will produce more high quality work. Finally, thanks to Stéphane Lehericy for his help on the cerebral anatomy and for giving us interesting neuroscience problems to test our algorithms.

I also want to thank all our collaborators that have provided datasets for us to test our methods. Thanks to Bruce Pike, V. V. Rymar and Jennifer Campbell of the Brain Imaging Center at McGill University, Canada for the human brain and biological phantom datasets. Thanks also to Thomas Knösche, Timm Wetzels and Alfred Anwander of the Max Planck Institute for Human Cognitive and Brain Sciences, Leipzig, Germany for their human brain datasets. In addition, I would like to thank Cyril Poupon and Jean-François Mangin and colleagues from NeuroSpin / CEA for the HARDI database of 13 subjects. Finally, I would like to thank G. Sapiro of the University of Minnesota and Kamil Ugurbil and Stéphane Lehéricy of the Center for Magnetic Resonance Research in Minneapolis, USA. Without all these datasets, I would not have been able to do this thesis.

Thank you to all Odyssée members and especially to team leader, Olivier Faugeras. Odyssée was a great environment to work in both for the scientific and social quality of the group. At Odyssée, I really felt part of a team. This felt very natural to me from the start. I had the feeling of being part of my ice hockey team back in Canada, where everybody has a specific job towards an ultimate goal. This was great for my productivity. Thanks to François, Adrien, Sylvain, Maria-Jose, Sandrine, Jonathan, Pierre, Théo and others for being colleagues and great friends throughout most of my three years at Odyssée.

I want to thank my squash group in Antibes. A little word in French for them. Un grand merci à ma bande de squasheurs du club d'Antibes. Vous avez été tellement généreux avec moi. Vous êtes comme une famille d'accueil pour moi et j'ai énormément apprécié votre soutien. Plus particulièrement, merci à mon équipe de ligue de France, champion national de la ligue 3 en 2006. Vous m'avez appris beaucoup sur la France et la vie. Vous m'avez aussi permis de faire le fou et de perdre le contrôle. Donc merci, Steph, Phi Phi, Guillaume, Christian et Éric.

Finally, a huge thanks to Maggie for her support and understanding in the last stretch of my PhD thesis. She was an important inspiration in my last accomplishments. I am also very grateful to my family. Thank you P-O, Claude and Hélène for your support and unconditional trust and love since the beginning of my scientific career, even if recently I was across the Atlantic Ocean.

I am most indebted to my mom Hélène for uncountable hours beside me doing homework throughout my life. I would not be the person I am today without your special attention. I would like to dedicate this thesis to you.

I sincerely thank you all,

Max

Part I

Introduction

INTRODUCTION

“All great deeds and all great thoughts have a ridiculous beginning”

–Albert Camus

CONTEXT

How can we get information about the human brain anatomy and in particular, about cerebral white matter? Cerebral dissection used to be the only way to access the neural architecture [Dejerine (1901); Gray (1918); Williams et al. (1997)]. Then, anatomists started using chemical markers to do neuronography [Pribam and MacLean (1953); Selden et al. (1998)]. More recently, neural fiber tractography based on local injection of chemical markers and subsequent observation of the induced propagation yielded high-quality connectivity mapping in the cat and monkey cerebral cortex [Young et al. (1995); Selden et al. (1998)]. As of today, diffusion-weighted (DW) magnetic resonance imaging (MRI) is the unique non-invasive technique capable of quantifying the anisotropic diffusion of water molecules in biological tissues like the human brain white matter.

The great success of DW-MRI comes from its capability to accurately describe the geometry of the underlying microstructure. DW-MRI captures the average diffusion of water molecules, which probes the structure of the biological tissue at scales much smaller than the imaging resolution. The diffusion of water molecules is Brownian under normal unhindered conditions, but in fibrous structure such as white matter, water molecules tend to diffusion along fibers. Due to this physical phenomenon, DW-MRI is able to obtain information about the neural architecture in vivo.

DW-MRI is a recent field of research with a history of roughly twenty years. It is a technique that was introduced in the middle of the 80’s by Le Bihan et al [LeBihan and Breton (1985)], Merboldt et al [Merboldt et al. (1985)] and Taylor et al. [Taylor and Bushell (1985)]. Shortly after the first acquisitions of diffusion-weighted images (DWI) in vivo [Moseley et al. (1990); Osment et al. (1990)], Basser et al [Basser et al. (1993, 1994b,a)] proposed the rigorous formalism of the diffusion tensor (DT) model. Diffusion tensor imaging (DTI) describes the three-dimensional

(3D) nature of anisotropy in tissues by assuming that the average diffusion of water molecules follows a Gaussian distribution. DTI has now proved to be extremely useful to study the normal and pathological human brain [LeBihan et al. (2001); Dong et al. (2004)]. However, the Gaussian assumption is over-simplifying the diffusion of water molecules and thus has some limitations. While the Gaussian assumption is adequate for voxels in which there is only a single fiber orientation (or none), it breaks down for voxels in which there is more complicated internal structure. This is an important limitation, since resolution of DTI acquisition is between 1mm^3 and 3mm^3 while the physical diameter of fibers can be between $1\mu\text{m}$ and $30\mu\text{m}$ [Poupon (1999); Beaulieu (2002)]. Research groups currently agree that there is complex fiber architecture in most fiber regions of the brain [Pierpaoli et al. (1996)]. In fact, it is currently thought that between one third to two thirds of imaging voxels in the human brain white matter contain multiple fiber bundle crossings [Behrens et al. (2007)].

The starting point and motivation of this thesis is therefore to recover fiber crossing information and overcome limitations of the DTI model. To do so, we use high angular resolution diffusion imaging (HARDI) and define new mathematical model-free approaches for the reconstruction of the angular part of the diffusion displacement probability density function (PDF) of water molecules, also called the diffusion orientation distribution function (ODF). This problem is of high complexity in the input HARDI data, in the mathematical tools used to describe the HARDI data and in the computational tools used to process the HARDI data. Hence, we address many different theoretical and computational issues related to the processing of HARDI data. One can simplify the HARDI estimation problem by modeling only the apparent diffusion coefficient (ADC) of water molecules or the diffusion ODF of water molecules. Ultimately, the goal is to recover the underlying fiber distribution so that the white matter architecture can be inferred with segmentation and tractography algorithms. Therefore, this thesis takes the reader from local HARDI reconstruction methods, to fiber bundle segmentation algorithms and fiber tractography algorithms for HARDI data. As the magnetic field strength of scanners increase, as the strength and speed of gradients increase, as new acquisition techniques like parallel imaging appear, and as new and powerful processing tools are developed, the development of HARDI processing methods is of utmost importance.

ORGANIZATION AND CONTRIBUTIONS OF THIS THESIS _____

This thesis is organized in three parts. The Background part describes the white matter cerebral anatomy, the principles of DW-MRI and the mathematical concepts required to understand the thesis. Then, the Methods part describes the theoretical and methodological contributions of the thesis. Finally, the Application part describes the contributions for a specific neuroscientific application. Figures 1.1 and 1.2 provide a graphical depiction of the organization of the thesis. The thesis covers several

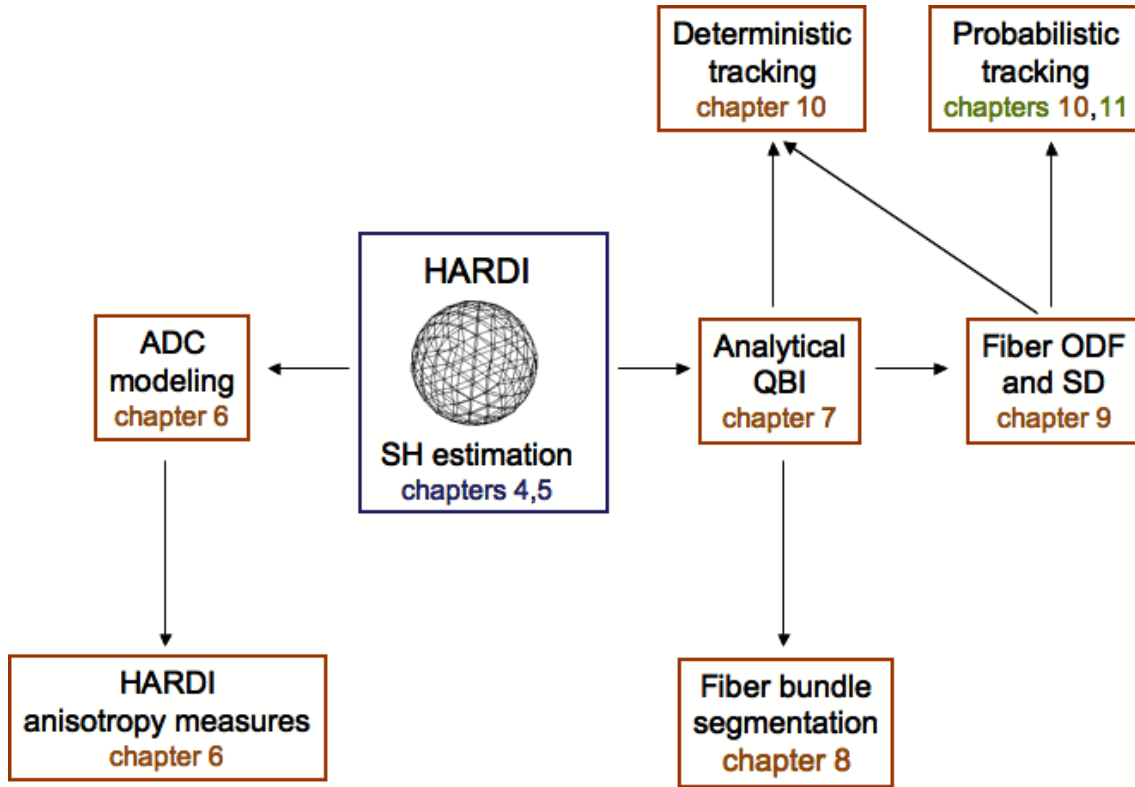


Figure 1.1: Sketch of the chapters of the thesis. Chapters in blue, orange and green correspond to chapters in the Background, Methods and Applications parts respectively. The thesis makes contributions from the local spherical harmonic (SH) estimation of the HARDI signal, to the ADC modeling to recover HARDI anisotropy measures able to classify multiple fiber distributions, to analytical QBI and spherical deconvolution (SD) methods used to develop new deterministic and probabilistic tractography algorithms and fiber bundle segmentation algorithms.

topics: from the local spherical harmonic (SH) estimation of the HARDI signal, to the ADC modeling to recover HARDI anisotropy measures able to classify multiple fiber distributions, to analytical QBI and spherical deconvolution (SD) methods used to develop new deterministic and probabilistic tractography algorithms and fiber bundle segmentation algorithms that describe the human brain white matter architecture. The introduction and conclusion chapters are not represented in Figures 1.1 and 1.2. We now give an overview of the organization and of the contributions of each chapter in turn.

Part I: Background

The Background part describes the white matter cerebral anatomy, the principles of DW-MRI and the mathematical concepts required to understand this thesis.

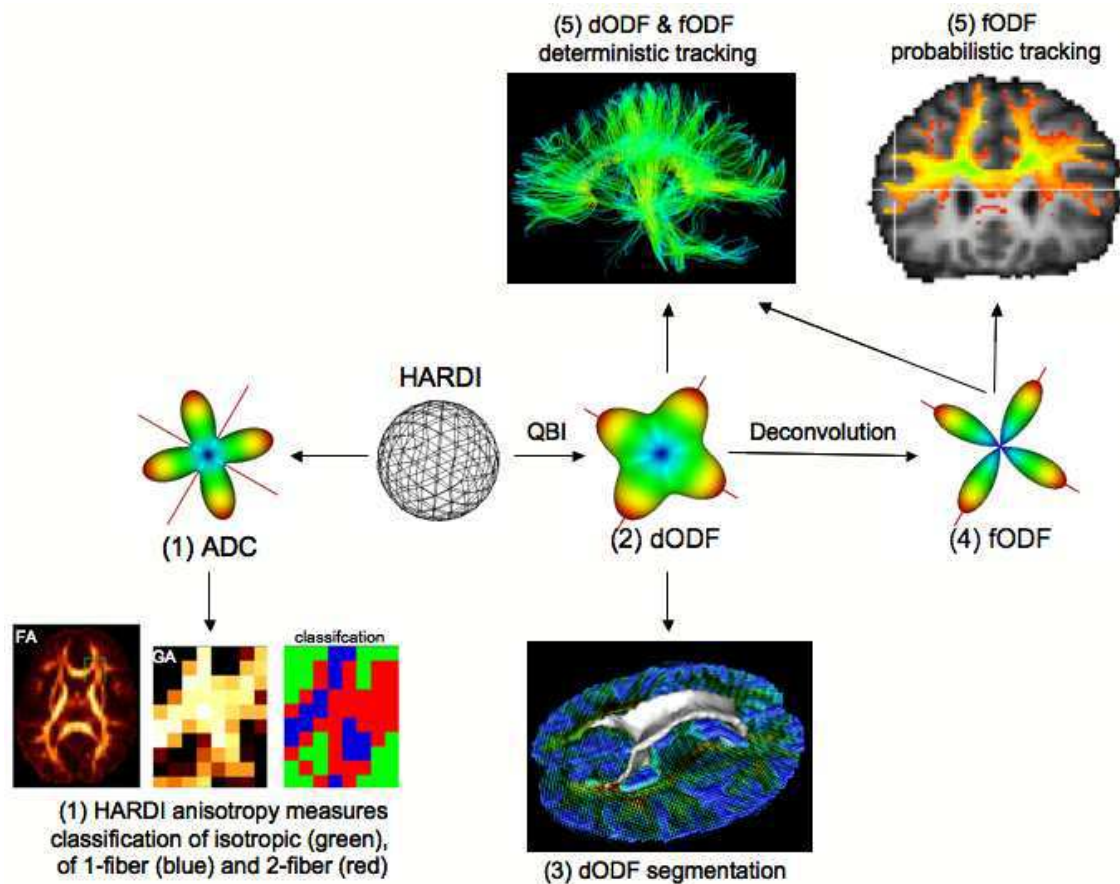


Figure 1.2: Graphical depiction of the contributions of the thesis: (1) The apparent diffusion coefficient (ADC) modeling and HARDI anisotropy measures used to classify voxels with isotropic, single fiber population and multiple fiber configurations. (2) The analytical QBI solution to obtain the diffusion ODF (dODF). (3) The segmentation of fiber bundles with coherent fiber tracts from QBI data. (4) The fiber ODF (fODF) reconstruction from spherical deconvolution of QBI data. (5) The deterministic and probabilistic tractography algorithms based on the q-ball and fiber ODF.

Chapter 3 answers the following question: What are we looking for in the human brain white matter? The principal interest of the thesis is anatomical connectivity. This chapter covers the basic cerebral anatomy of the white matter. In particular, what are the different brain and neural tissues? What is the organization of the white matter? How is the brain connected? What are some of the large fiber bundles in the brain? First, we review the different human brain tissues and neural tissue types, from the neuron, the gray matter, to the white matter. Then, we cover the composition and organization of the human white matter. Finally, we highlight the different fiber tracts and highlight the way different parts of the brain interact.

Chapter 4 answers the following question: How do we image white-matter non-invasively? This chapter covers the principles of DW-MRI and high angular resolution diffusion imaging (HARDI) reconstruction techniques. In particular, what are the physical principles of DW-MRI? What are the fiber reconstruction algorithms to infer microstructure of biological tissues? First, we describe the basic principles of DW-MRI. Next, we review the different DW-MRI acquisition techniques. Properties of each acquisition technique are described separately. Finally, a special emphasis is put on the approaches designed to reconstruct the multiple fiber distributions of water diffusion. Many functions are used in the literature to measure the diffusion properties of the underlying biological tissue. Some methods are model dependent, some model-free, some methods have linear solutions whereas others require non-linear optimization schemes. These different methods are described.

Chapter 5 answers the following question: What is the appropriate mathematical tool to process HARDI data? This chapter covers the mathematical concepts needed to understand the thesis. In particular, what is the appropriate mathematical tool to represent discrete data that lives on the sphere? The SH basis is chosen to be the appropriate tool to represent discrete spherical functions in this thesis. The spherical harmonics are first defined as the solution to the Laplace's equation in spherical coordinates. Then, important properties of the SH basis are described and we see how the SH basis can be used to estimate spherical functions. Finally, the Funk-Hecke theorem is stated. It is an important theorem to solve integrals on the sphere.

Part II: Methods

The Methods part describes the original and most important contributions made in this thesis.

Chapter 6 answers the following question: Given a set of discrete measurements on the sphere, how do we estimate this signal robustly? This chapter covers the HARDI estimation using spherical harmonics and regularization. As a consequence, the apparent diffusion coefficient (ADC) can be modeled and reconstructed robustly. First, we develop a closed-form regularization algorithm to estimate the HARDI signal and the ADC. Second, we develop the link between the SH basis and the high order diffusion tensor (HODT) so that it is possible to compare HARDI anisotropy measures. Finally, we see how the HARDI anisotropy measures can be used to highlight voxels containing multiple fiber distributions.

Chapter 7 answers the following question: Given our new regularized method to obtain the HARDI signal, can it be used to do q-ball imaging (QBI)? This chapter proposes a new analytical QBI solution using our regularized spherical harmonics estimation. In particular, can the Funk-Radon transform be solved in a single step

with spherical harmonics? First, we develop a regularized, fast and robust analytical solution for the QBI reconstruction. To do so, a new corollary of the Funk-Hecke theorem is proved to obtain an elegant mathematical simplification of the Funk-Radon transform and to obtain the analytical QBI solution in a single step. Second, we validate and compare the analytical QBI solution with the numerical QBI solution on synthetic simulations, on a rat biological phantom and on real datasets. Finally, a discussion on its benefits over the other state-of-the-art approaches is done.

Chapter 8 answers the following question: Now that we have robust q-ball reconstructions, can we segment this image of q-ball diffusion ODF into coherent fiber bundles? This chapter covers a new statistical region-based active contour algorithm to segment HARDI data. In particular, how can the segmentation problem be formulated and solved efficiently on HARDI data? What is gained by the HARDI segmentation with respect to the DTI segmentation? Is it possible to validate the segmentation results and make the segmentation automatic? First, a new method to segment HARDI data is developed using our fast, robust and analytical QBI solution. Then, a region-based statistical surface evolution is done on this image of diffusion ODF to efficiently find coherent white matter fiber bundles. Next, our method is shown to be appropriate to propagate through regions of fiber crossings and our results outperform state-of-the-art DTI segmentation methods. Finally, results obtained on synthetic data, on a biological phantom, on real datasets and on all 13 subjects of a public HARDI database show that our method is reproducible and brings a strong added value to DW-MRI segmentation.

Chapter 9 answers the following question: Can we improve angular resolution of QBI? This chapter covers the problem of reconstructing a fiber ODF and not only a diffusion ODF. In particular, can the diffusion ODF be transformed into a sharper fiber ODF? Can this sharper fiber ODF be obtained from QBI? Can it be related to the raw HARDI signal and other spherical deconvolution methods that also estimate a fiber ODF? First, we develop a new spherical deconvolution operation that transforms the q-ball diffusion ODF into a sharper fiber ODF. This deconvolution transformation improves angular resolution and fiber detection of QBI by approximately 15° . Then, our new fiber ODF is shown to behave similarly to the fiber ODF estimated from the filtered, regularized and constrained spherical deconvolution methods of [Tournier et al. (2004, 2007); Sakaie and Lowe (2007)] as well as the fiber ODF estimated with FORECAST [Anderson (2005)]. Next, we extensively compare these methods quantitatively by varying the harmonic order ℓ , the b -value, the sampling density N , the noise level (SNR), the separation angle between fibers and the volume fraction of each fiber. We also evaluate the robustness to noise and the angular resolution of each method. Finally, all fiber ODFs are reconstructed and compared on the biological phantom and the real datasets.

Chapter 10 answers the following question: Now that we have robust fiber ODF estimates, how can we integrate this local information into new deterministic and new probabilistic tractography algorithms? This chapter covers a new deterministic tractography algorithm and a new probabilistic algorithm based on the fiber ODF. In particular, how can the deterministic DTI tracking be generalized? Can deterministic tracking take into account all the maxima of the fiber ODF? Can DTI probabilistic tracking be generalized? How does deterministic and probabilistic tracking compare? First, we develop a new deterministic tracking algorithm and a new probabilistic tracking algorithm are developed to extend the classical DTI tractography algorithms. To do so, our new fiber ODF is used. Then, accurate results of complex fiber bundles with crossing, fanning and branching configurations are obtained from our new tracking algorithms. Finally, an extensive comparison of the new deterministic and the new probabilistic tracking algorithms is performed on a human brain dataset.

Part III: Application

The Application part is short but illustrates the added value and remaining limitations of HARDI-based probabilistic tractography algorithms.

Chapter 11 answers the following question: Now that we have state-of-the-art local reconstruction methods and tractography algorithms, can we track through regions of complex fiber crossings for a neuroscientific application? This chapter covers the tractography of transcallosal fibers. Current DTI based methods neglect these transcallosal fibers, which might lead to wrong interpretations of the brain functions. In particular, can transcallosal fibers be reconstructed with our new fiber ODF-based probabilistic tractography? Are the tracking results consistent and reproducible across subjects? First, we apply our new fiber ODF probabilistic tractography algorithm to reconstruct the transcallosal fibers. Then, we see how we improve the tractography results of transcallosal fibers intersecting with the corona radiata and the superior longitudinal fasciculus. Finally, the tracking results are reproduced and compared on a group of 8 subjects.

Appendix A

Appendix A describes the datasets used in most chapters of this thesis. Our synthetic data simulations is first described in Appendix A.1. Then the acquisition of a biological phantom composed of rat spinal cords is described in Appendix A.2. Finally, the real human brain datasets acquisitions are described in Appendix A.3. We have collaborated with four institutes to obtain these real datasets. The Brain Imaging Center (BIC) in Montreal, Canada, the Max Planck Institute (MPI) in Leipzig, Germany, the Center for Magnetic Resonance Research (CMRR) in Minneapolis, USA

and the Neurospin/CEA in Paris, France. We are very grateful to these institutions for allowing us to use their datasets.

Appendix B

The appendix enumerates the publications from the author arising from this thesis. We have published in major international journals, in important conferences and we have made a special effort in publishing research reports of most of our work.

Software contributions

Finally, we would like to point out that most of the algorithms presented in this manuscript are now available upon request¹ as an extension to the Brainvisa² software platform for visualization and analysis of multi-modality brain data. The ODF and spherical function visualization using spherical harmonics basis will also be soon available as part of the Slicer3³ open source software. This integration work has been done in collaboration with Christophe Lenglet (Siemens Corporate Research, Princeton, USA) and Demian Wassermann (INRIA Sophia-Antipolis, France / Computer Science Department, University of Buenos Aires, Argentina).

¹<https://gforge.inria.fr>

²<http://brainvisa.info>

³<http://www.na-mic.org/Wiki/index.php/Slicer3>

INTRODUCTION (FRANÇAIS)

“Le commencement est la moitié de l’action ”

–Wilfrid Laurier (politicien canadien)

CONTEXTE

Comment peut-on obtenir de l’information sur l’anatomie cérébrale et plus particulièrement sur la matière blanche du cerveau humain? À l’époque, les anatomistes réalisaient des dissections du cerveau [Dejerine (1901)] et des neuronographies à base de strychnine ou autres marqueurs chimiques [Pribam and MacLean (1953)], [Selden et al. (1998)]. Plus récemment, l’estimation des connexions neuronales par injection de marqueurs chimiques et observation de leur propagation a permis d’obtenir des cartes de connectivité de haute qualité du cortex cérébral chez le chat et le singe [Young et al. (1995)], [Selden et al. (1998)]. Actuellement, l’Imagerie par Résonance Magnétique (IRM) de diffusion est la seule technique non-invasive permettant de quantifier l’anisotropie de la diffusion des molécules d’eau dans des tissus biologiques tel que la matière blanche du cerveau de l’homme.

Le phénomène de diffusion est un processus physique macroscopique résultant du mouvement aléatoire des molécules d’eau. Toutes les particules microscopiques subissent en permanence un mouvement de translation et de rotation dû à leur énergie thermique. Ce phénomène est également connu sous le nom de mouvement Brownien. Ce dernier a été formalisé en 1905 par Albert Einstein et caractérisé comme dépendant de la température, de la viscosité du domaine et de la taille des particules. À une échelle macroscopique, la diffusion traduit la tendance des molécules à migrer des régions à forte concentration vers les régions à faible concentration. Le concept essentiel pour l’IRM de diffusion est que tout processus de diffusion dans un milieu biologique reflète directement l’architecture de ce dernier à une échelle microscopique. Ceci s’explique par le fait que le mouvement moléculaire est favorisé dans les directions alignées avec les faisceaux de fibres et contraint dans les directions orthogonales. À chaque voxel, la mesure de ce mouvement le long d’un certain nombre de directions fournit de précieuses informations sur l’orientation locale des fibres. Les

données ainsi obtenues constituent des IRM pondérées en diffusion.

Le développement de l'IRM et ses applications en neuroimagerie ont apporté de grands espoirs au cours des vingt dernières années pour l'exploration non-invasive de l'anatomie et de l'activité du cerveau humain. L'IRM de diffusion est une technique introduite au milieu des années 1980 par Le Bihan et al [LeBihan and Breton (1985)], Merboldt et al [Merboldt et al. (1985)] et Taylor et al [Taylor and Bushell (1985)]. Sensible au mouvement des molécules d'eau, elle permet d'inférer des informations sur l'architecture des tissus biologiques étudiés. Peu après les premières acquisitions d'images caractérisant la diffusion anisotrope des molécules d'eau in vivo [Moseley et al. (1990); Osment et al. (1990)], Basser et al [Basser et al. (1993, 1994b,a)] ont proposé le formalisme rigoureux du modèle du tenseur de diffusion. Celui-ci fournit une forme analytique décrivant précisément la nature tridimensionnelle de l'anisotropie des tissus. Ce modèle donna alors naissance à l'imagerie par tenseur de diffusion (DTI), une technique qui a maintenant un vaste domaine d'applications [Mori and van Zijl (2002); Neil et al. (2002); Horsfield and Jones (2002); Sotak (2002); Lim and Helpert (2002)]. Le tenseur de diffusion encapsule les propriétés de diffusion moyenne des molécules d'eau au sein d'un voxel dont la taille est typiquement de 1 à 3 mm. Effectivement, à cette résolution spatiale, il est maintenant reconnu [Pierpaoli et al. (1996)] qu'il y a entre un et deux tiers des voxels de la matière blanche [Behrens et al. (2007)] qui contiennent plusieurs faisceaux de fibres qui se croisent. De par son hypothèse de base qu'un seul faisceau traverse chaque voxel de l'image, le DTI est donc intrinsèquement limité à ces endroits de croisements de fibres. Par conséquent, la plupart des méthodes courantes de traitement en DTI négligent alors les faisceaux complexes qui se croisent, ce qui peut amener à une mauvaise analyse et interprétation de l'anatomie et des fonctions du cerveau.

Cette thèse est donc dédiée au développement de techniques d'IRM à haute résolution angulaire (HARDI) pour retrouver une ou plusieurs directions de fibres à chaque voxel de l'image. Le but est de surmonter les limites du DTI. Pour ce faire, l'imagerie par q-ball (QBI) est une technique de reconstruction HARDI récente qui permet d'estimer la distribution d'orientation des fibres (ODF). La technique de QBI ainsi que l'ODF de diffusion permettent de retrouver plusieurs directions de fibres. Ces techniques joueront donc un rôle très important tout au long de cette thèse.

Les données HARDI soulèvent depuis moins de dix ans un grand nombre de questions mathématiques et computationnelles, de par la complexité des données, le besoin de modèles et d'algorithmes adéquats pour retrouver la distribution de fibres sous-jacentes ainsi que le besoin de méthodes de traitement du signal adaptées à ces nouvelles images. Dans cette thèse nous nous attaquons à de nombreux problèmes théoriques et computationnels liés au traitement des données HARDI: de la reconstruction locale de l'information de diffusion, à la segmentation et au suivi de fibres (c.f. *tracking* ou *tractographie*). Nous appliquons également certains des outils proposés à l'analyse des connexions anatomiques reliant les hémisphères du cerveau

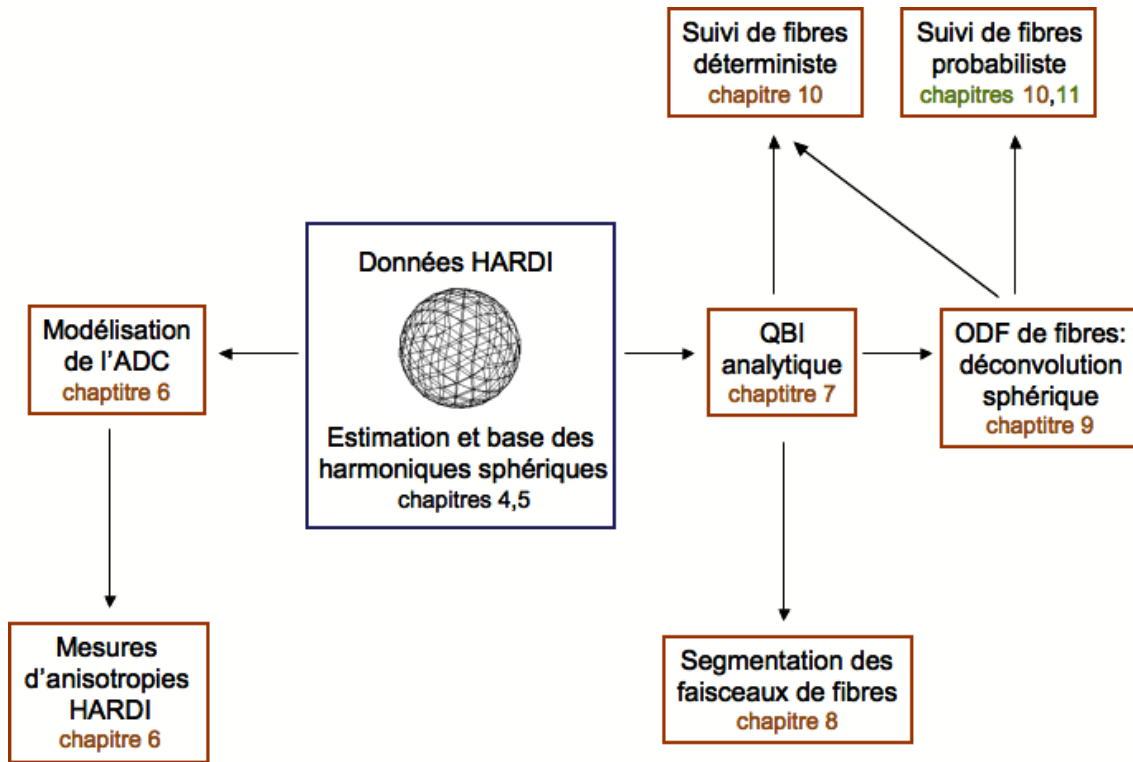


Figure 2.1: Organisation des chapitres de cette thèse. Les chapitres de couleur bleue, orange et verte correspondent respectivement aux chapitres des parties *Préliminaires et Théorie*, *Méthodes* et *Application*.

humain par les fibres du corps calleux. Grâce à l'accroissement de la puissance des champs magnétiques, à l'arrivée de nouvelles techniques d'acquisition ainsi qu'au développement de nouveaux et puissants outils de traitement, notre compréhension de la structure la plus complexe du corps humain s'améliore progressivement. Ces avancées sont cruciales pour le diagnostic rapide et aussi précis que possible des pathologies neurologiques.

ORGANISATION ET CONTRIBUTIONS DE CETTE THÈSE _____

Cette thèse est organisée en trois parties. La première, *Préliminaires et Théorie*, décrit l'anatomie cérébrale, les principes de base de l'IRM de diffusion et les notions mathématiques importantes pour comprendre la thèse. La seconde, *Méthodes*, développe nos contributions théoriques et algorithmiques. Il s'agit de la partie la plus importante de ce document. Enfin, la partie *Applications* présente une application précise de nos outils pour la reconstruction des fibres transcallosales, fibres très importantes dans le fonctionnement du cerveau. Les Figures 2.1 et 2.2 illustrent l'organisation et les grandes lignes de cette thèse.

Comme le titre de la thèse le suggère, nous abordons un grand nombre de sujets:

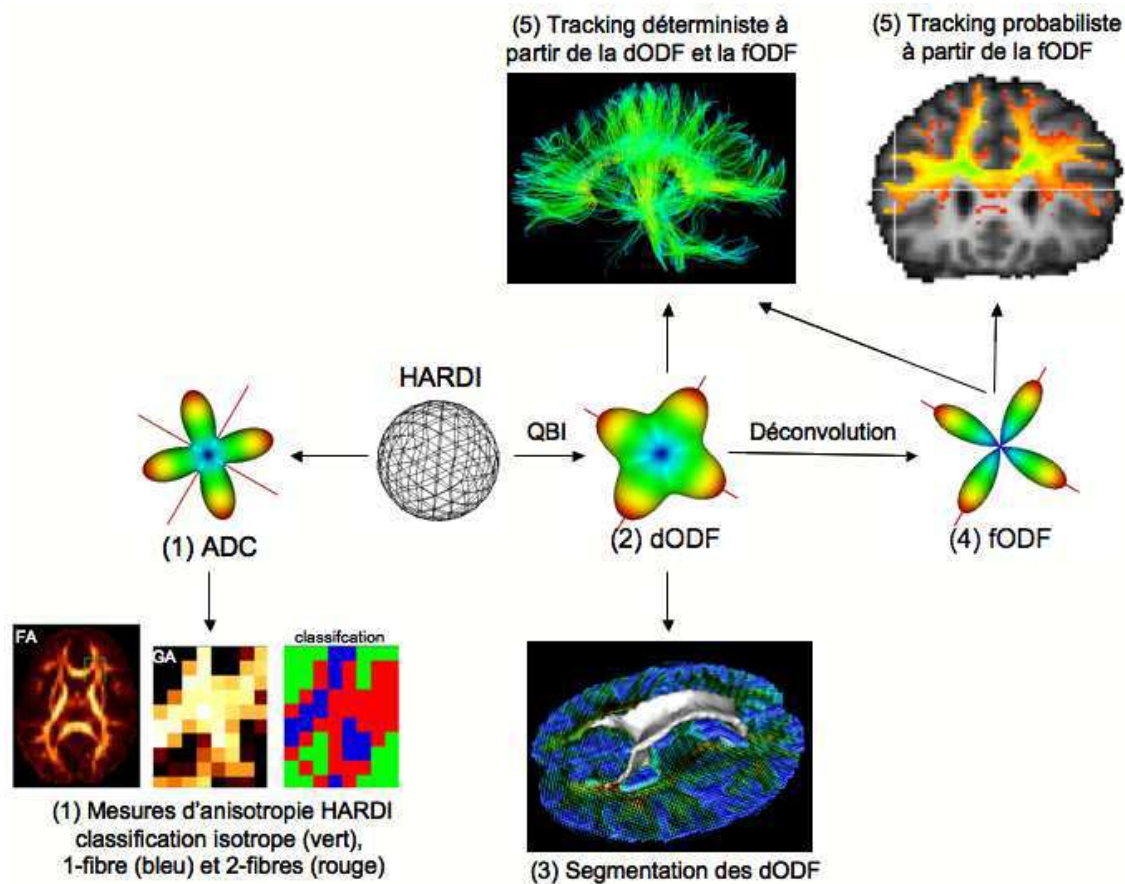


Figure 2.2: Schéma des contributions de cette thèse. (1) L'estimation du coefficient de diffusion apparent (ADC) et l'étude des mesures d'anisotropie HARDI pour la classification des voxels contenant une distribution isotrope, une distribution d'une seule population de fibres et d'une distribution de plusieurs faisceaux fibres. (2) Le développement d'une solution analytique du QBI pour obtenir l'ODF de diffusion (dODF). (3) Nouvel algorithme de segmentation de ces images d'ODF's obtenues par QBI. (4) Le calcul de l'ODF de fibres avec une nouvelle méthode de déconvolution sphérique à partir des données QBI. (5) Nouveaux algorithmes de suivi de fibres (tracking) déterministes et probabilistes à partir de l'ODF du q-ball et de l'ODF de fibres.

l'estimation locale du signal HARDI, la segmentation et le suivi de fibres. D'abord, nous décrivons une méthode d'estimation robuste du signal HARDI en utilisant une base d'harmoniques sphériques et un terme de régularisation. Ensuite, nous présentons la modélisation du coefficient de diffusion apparent (ADC) et le développement de mesures d'anisotropie HARDI pour la classification des voxels contenant une distribution isotrope, une distribution d'une seule population de fibres et d'une distribution de plusieurs faisceaux fibres. De plus, nous proposons une solution analytique pour QBI ainsi qu'un nouvel algorithme de segmentation de ces images d'ODF

obtenues par QBI. Nous présentons également le calcul de l'ODF de fibres avec une nouvelle méthode de déconvolution sphérique à partir des données QBI. Enfin, nous développons de nouveaux algorithmes de suivi de fibres déterministes et probabilistes à partir de l'ODF du q-ball et de l'ODF de fibres. Tous ces nouveaux algorithmes sont testés et appliqués sur des données HARDI simulées, sur un fantôme biologique et sur des données réelles de plusieurs sujets dans des régions complexes où plusieurs faisceaux se croisent, illustrant ainsi la valeur ajoutée de nos méthodes HARDI sur la plupart des méthodes courantes en DTI qui négligent ces faisceaux complexes.

Partie I : Préliminaires et Théorie

Cette partie décrit l'anatomie cérébrale, les principes de base de l'IRM de diffusion et les notions mathématiques importantes pour comprendre la thèse.

Chapitre 3 - Tissu Neuronal et Anatomie Cérébrale de la Matière Blanche :

Comment la matière blanche du cerveau humain est-elle organisée? Ce chapitre présente donc l'essentiel de l'anatomie cérébrale et de la matière blanche nécessaire pour comprendre les résultats de ce document. Nous présentons les différents tissus neuronaux ainsi que les différents lobes cérébraux. Nous décrivons certains faisceaux de fibres importants qui servent à connecter certains lobes entre eux. Nous introduisons également les fibres d'*association*, des fibres de communication entre des régions lointaines au sein d'un même hémisphère, les fibres de *projection*, fibres qui se projettent à partir des noyaux gris centraux jusqu'aux différentes parties du cortex, et les fibres *commissurales*, fibres très importantes connectant les hémisphères et passant par le corps calleux. Ce chapitre servira, tout au long de la thèse, d'élément de comparaison qualitative avec nos résultats de segmentation et de suivi de fibres.

Chapitre 4 - Principes de l'IRM de diffusion: Au Delà du Tenseur de Diffusion :

Comment peut-on imager la matière blanche de façon non-invasive? Ce chapitre présente les principes de base de l'IRM de diffusion et des techniques de reconstruction pour l'imagerie de diffusion à haute résolution angulaire (HARDI). Après un bref historique de la résonance magnétique nucléaire (NMR) et de l'IRM, nous présentons les principes physiques de base de la diffusion moléculaire et du mouvement Brownien. Par la suite, nous décrivons les techniques d'IRM pondérée en diffusion (IRM-PD). Pour ce faire, nous introduisons la séquence de diffusion classique de Stejskal-Tanner [Stejskal and Tanner (1965)] qui permis d'acquérir les premières images pondérées en diffusion le long d'une seule direction [LeBihan and Breton (1985)]. Ces premières acquisitions inspirèrent alors le développement de l'IRM-PD mesurée le long de plusieurs directions de gradient de diffusion. C'est ainsi que les gens commencèrent à faire de l'IRM du coefficient apparent de diffusion (ADC) [Moseley et al. (1990)]. Basser et al [Basser et al. (1993)] proposèrent ensuite le formalisme rigoureux du modèle du tenseur de diffusion, ce

qui donna naissance à l'imagerie par tenseur de diffusion (DTI). Le DTI a besoin d'au moins six images pondérées en diffusion pour reconstruire le tenseur de diffusion. C'est seulement depuis moins de 10 ans que les acquisitions HARDI avec plus d'une soixantaine d'images pondérées en diffusion existent. L'historique du développement de l'IRM de diffusion est illustré à la Figure 4.2. La dernière partie du chapitre présente enfin toutes les méthodes de reconstruction HARDI récentes. Ces méthodes formaient l'état de l'art au moment de la rédaction de ce document. Ce chapitre permet donc de motiver le sujet de cette thèse et de situer les méthodes choisies par rapport aux techniques existantes jusqu'à ce jour.

Chapitre 5 - Mathématiques sur la Sphère: la Base des Harmoniques Sphériques : Quel est l'outil mathématique approprié pour traiter les données HARDI? Ce chapitre décrit les outils mathématiques ainsi que les propriétés mathématiques importantes pour comprendre cette thèse. Les données HARDI sont en fait des mesures physiques qui ont des contraintes. Les données sont discrètes, elles vivent sur la sphère, elles sont symétriques et positives. Il faut donc trouver le bon outil pour estimer au mieux ces données. Pour ce faire, nous utilisons la base des harmoniques sphériques. Comme la transformée de Fourier dans le plan, la transformation en harmoniques sphériques est très naturelle pour traiter les données sphériques. D'abord, le chapitre présente la définition des harmoniques sphériques comme étant la solution de l'équation de Laplace. Ensuite, nous révisons certaines propriétés des harmoniques sphériques qui nous permettront d'estimer, de lisser et de calculer des distances et statistiques entre nos fonctions sphériques. Enfin, nous définissons le théorème de Funk-Hecke. Ce théorème sera d'une grande importance pour le développement de nos solutions analytiques en imagerie par q-ball.

Partie II : Méthodes

Cette partie développe nos contributions théoriques et algorithmiques.

Chapitre 6 - Estimation et Applications du Coefficient Apparent de Diffusion (ADC) : Comment estimons-nous le signal HARDI de façon robuste? Ce chapitre présente une solution linéaire et régularisée pour l'estimation du signal HARDI à partir des harmoniques sphériques. Par conséquent, nous pouvons reconstruire l'ADC avec une représentation en harmoniques sphériques. Nous faisons donc le point sur les techniques d'approximation de l'ADC à partir de modèles de tenseurs d'ordres supérieurs et présentons aussi leur application à la définition de mesures d'anisotropie. Nous développons plus spécifiquement les outils adéquats pour traiter et estimer l'ADC bruité provenant des données HARDI. À partir d'une base modifiée d'harmoniques sphériques et de ses propriétés, nous proposons une nouvelle méthode de régularisation avec l'opérateur de Laplace-Beltrami obtenant des coefficients

de diffusion plus lisses. Nous développons également le lien entre la base des harmoniques sphériques et le tenseur d'ordre supérieur. Il s'avère qu'une simple transformation linéaire est possible entre les deux espaces. Nous validons notre nouvelle approche sur des données synthétiques, sur un fantôme biologique ainsi que sur un cerveau humain. De plus, nous étudions l'état de l'art des mesures d'anisotropie calculées à partir de modèles d'ordres supérieurs et nous évaluons leur capacité à décrire le processus de diffusion. Finalement, nous démontrons que les mesures d'anisotropie HARDI peuvent être utilisées pour la classification des voxels contenant une distribution isotrope, une distribution d'une seule population de fibres et une distribution de plusieurs faisceaux de fibres.

Chapitre 7 - Solution Analytique pour l'Imagerie par Q-Ball (QBI) :

Pouvons-nous utiliser notre méthode robuste d'estimation du signal HARDI pour faire du QBI? Ce chapitre présente une solution linéaire, rapide, régularisée et analytique pour faire du QBI. La reconstruction de la fonction de distribution d'orientation des fibres (ODF) par q-ball est une technique largement répandue dans l'étude des croisements de fibres. L'ODF nous renseigne sur la distribution de probabilité des directions de diffusion des molécules d'eau présentes dans la matière blanche. Cependant, cette méthode ainsi que les autres techniques d'estimation HARDI sont basées sur des solutions numériques complexes et sans aucun processus de régularisation. Nous approximons d'abord le signal par une série d'harmoniques sphériques lissées. Cette formulation nous permet d'obtenir une simplification élégante de la transformée de Funk-Radon en utilisant le théorème de Funk-Hecke. Nous obtenons donc un algorithme robuste et rapide pour l'estimation de l'ODF. L'efficacité et la précision de l'approximation sont évaluées sur des données synthétiques et réelles. Nous faisons varier les paramètres importants dans la formation de l'image et observons les différentes répercussions sur la forme des ODFs. Contrairement au tenseur de diffusion, nous démontrons que l'extraction des maxima de l'ODF permet de retrouver des croisements de fibres dans de nombreux voxels de l'image.

Chapitre 8 - Segmentation d'Images de Diffusion à Haute Résolution Angulaire :

Maintenant que nous avons des ODF de diffusion q-ball, comment pouvons-nous segmenter et extraire certains faisceaux de fibres de la matière blanche? Ce chapitre décrit une nouvelle méthode de segmentation des images HARDI obtenues par QBI. D'abord, l'ODF de diffusion est estimée à l'aide de la base des harmoniques sphériques et de notre nouvelle méthode d'estimation analytique, robuste et rapide. Nous utilisons ensuite une évolution de surface par statistique de région sur cette image d'ODF, pour retrouver des ensembles de faisceaux de fibres cohérents partageant les mêmes caractéristiques. Nous montrons que notre nouvelle méthode reproduit les résultats de l'état de l'art basés sur le tenseur de diffusion et que nous améliorons les résultats de segmentation dans les régions de croisements de faisceaux de fibres,

régions où le tenseur de diffusion est intrinsèquement limité. Enfin, nos résultats sur des données simulées, sur un fantôme biologique, sur des données réelles ainsi que sur la base de données HARDI publique [Poupon et al. (2006)] comportant 13 sujets, démontrent que notre approche est reproductible, automatique et apporte une valeur ajoutée importante pour la segmentation des IRM pondérées en diffusion.

Chapitre 9 - Améliorer l'Imagerie par Q-Ball: de l'ODF de Diffusion à une

Estimation de l'ODF de Fibres : Pouvons-nous améliorer la résolution angulaire du QBI? Ce chapitre décrit une nouvelle méthode de déconvolution sphérique pour transformer l'ODF de diffusion en une ODF de fibres. Cette transformation de *sharpening* augmente la résolution angulaire d'environ 15° et facilite l'extraction des maxima de l'ODF. L'outil mathématique important est encore une fois le théorème de Funk-Hecke. La déconvolution de l'ODF de diffusion par le profil de diffusion ODF estimé à partir des données HARDI réelles nous permet de réduire la partie isotrope de l'ODF et donc de rehausser les compartiments de fibres sous-jacents. Par la suite, nous démontrons que l'ODF de fibres et les distributions obtenues par déconvolution sphérique classique [Tournier et al. (2004); Anderson (2005); Tournier et al. (2007); Sakaie and Lowe (2007)] se comportent de la même manière sur des simulations de données HARDI. Nous étudions également l'effet des différents paramètres d'acquisition. Faisant varier la valeur de b , le nombre de mesures pondérées en diffusion, le bruit, l'angle de séparation entre les compartiments de fibres et leurs proportions dans le voxel, nous pouvons quantifier les limites et avantages de chaque algorithme de déconvolution sphérique. Finalement, tous les algorithmes de déconvolution sphérique ont des comportements très similaires et nous démontrons que l'ODF de fibres produite par la plupart de ces techniques donnent des résultats semblables sur le fantôme biologique et sur nos jeux de données réelles HARDI.

Chapitre 10 - Suivi de Fibres Déterministes et Probabilistes en Imagerie de Diffusion à Haute Résolution Angulaire :

Maintenant que nous avons des estimations robustes de l'ODF de fibres, comment faut-il intégrer cette information locale dans de nouveaux algorithmes de suivi de fibres? Ce chapitre décrit de nouveaux algorithmes de suivi de fibres (tracking) déterministes et probabilistes basés sur l'ODF de fibres. La nouveauté dans le suivi de fibres déterministe est d'utiliser toutes les orientations maximales de l'ODF pour permettre au suivi de fibres de se séparer en plusieurs directions. La nouveauté du suivi de fibres probabiliste est d'utiliser l'ODF de fibres comme distribution de probabilité pour tirer un très grand nombre de fois une orientation aléatoire à chaque pas du suivi. Ainsi, nos résultats de suivi de fibres permettent de retrouver des réseaux complexes de l'architecture neuronale de la matière blanche comportant des croisements, des embranchements et des configurations de fibres en éventail. De plus, nous démontrons que le suivi de fibres fait à partir de l'ODF de fibres est plus complet et de meilleure qualité qu'à partir du tenseur

de diffusion ou de l'ODF de diffusion. Enfin, une autre contribution importante est l'étude poussée et la comparaison des algorithmes de suivi de fibres déterministes et probabilistes sur des données HARDI réelles à partir du DTI, de l'ODF de diffusion et de l'ODF de fibres.

Partie III : Applications

Cette dernière partie présente une application précise de nos outils pour la reconstruction des fibres transcallosales, fibres très importantes dans le fonctionnement du cerveau.

Chapitre 11 - Reconstruction des Fibres Transcallosales à Partir d'un Suivi de Fibres Probabiliste : Maintenant que nous avons des méthodes état de l'art de suivi de fibres, pouvons-nous retrouver des réseaux de fibres complexes pouvant être utilisés dans une application neuroscientifique? Ce chapitre décrit une application intéressante sur le corps calleux et la reconstruction des fibres transcallosales. Ces fibres sont normalement ignorées par les techniques de suivi de fibres en DTI car elles croisent le faisceau supérieur longitudinal ainsi que la couronne rayonnante. Notre suivi de fibres HARDI basé sur l'ODF de fibres nous permet de retrouver ces fibres transcallosales sur une base de données de huit sujets. De plus, nos résultats quantitatifs de suivi de fibres sur plusieurs sujets nous permettent de déterminer les endroits où le suivi de fibres DTI est limité. Ce suivi de fibres nous confère donc une connaissance anatomique plus fine de ces parties du cerveau.

Annexe A - Données HARDI utilisées dans cette thèse :

L'annexe A décrit les jeux de données utilisés tout au long de cette thèse. D'abord, l'annexe A.1 présente notre génération de données synthétiques pour tester et valider nos algorithmes. Nous utilisons le modèle de plusieurs tenseurs de diffusion pour produire des signaux isotropes, des signaux comportant une seule fibre et des signaux comportant plusieurs croisements de fibres. Dans l'annexe A.2, nous décrivons ensuite l'acquisition du fantôme biologique de Campbell et al de l'Université McGill à Montréal, Canada [Campbell et al. (2005)]. Ce fantôme biologique est composé de deux moelles épinière de rat placées en configuration de croisement. Ce jeu de données nous est d'une grande utilité pour valider nos algorithmes de détection de croisements de fibres. Finalement, nous présentons les jeux de données réelles dans l'annexe A.3. Les données ont été acquises au Brain Imaging Center (BIC) à Montréal au Canada, au Max Planck Institute (MPI) à Leipzig en Allemagne, au Center for Magnetic Resonance Research (CMRR) à Minneapolis aux USA et au CEA / NeuroSpin à Paris en France. Nous sommes très reconnaissant envers tous nos collaborateurs d'avoir partagé leurs données avec nous.

Annexe B - Publications de l'auteur issues des travaux de cette thèse :

Cette annexe énumère nos publications issues des contributions de cette thèse. Nous avons publié dans plusieurs journaux internationaux, dans plusieurs conférences de haut niveau et nous nous sommes efforcés de publier nos travaux dans plusieurs rapports de recherche INRIA.

Contributions logicielles

Par ailleurs, nous souhaitons souligner que tous les algorithmes présentés dans cette thèse sont maintenant disponibles, sur demande¹, comme extension de la plateforme logicielle Brainvisa² pour la visualisation et l'analyse d'images cérébrales provenant de différentes modalités. De plus, la visualisation des ODFs et des fonctions sphériques avec leur représentation en harmoniques sphériques sera bientôt disponible dans Slicer³. Ce travail d'intégration a été réalisé en collaboration avec Christophe Lenglet (Siemens Corporate Research, Princeton, USA) et Demian Wassermann (INRIA Sophia-Antipolis, France / Computer Science Department, University of Buenos Aires, Argentina).

¹<https://gforge.inria.fr>

²<http://brainvisa.info>

³<http://www.na-mic.org/Wiki/index.php/Slicer3>

Part II

Background

CHAPTER **3**

NEURAL TISSUE AND HUMAN BRAIN WHITE MATTER

“The human brain is a most unusual instrument of elegant and as yet unknown capacity.”

–Stuart Seaton ()

Contents

3.1 The Human Brain and Neural Tissue	24
3.2 Organization of the White Matter	27
3.2.1 Important white matter bundles for this thesis	30
3.3 Conclusion	33

OVERVIEW

What are we looking for in the human brain white matter? The principal interest of the thesis is anatomical connectivity. This chapter covers the basic cerebral anatomy of the white matter. In particular, what are the different brain and neural tissues? What is the organization of the white matter? How is the brain connected? What are some of the large fiber bundles in the brain? First, we review the different human brain tissues and neural tissue types, from the neuron, the gray matter, to the white matter. Then, we cover the composition and organization of the human white matter. Finally, we highlight the different fiber tracts and highlight the way different parts of the brain interact. Overall, this introductory chapter is inspired from atlases, review articles and thesis chapters from [Gray (1918); Williams et al. (1997); Stanisz et al. (1997); Poupon (1999); Beaulieu (2002); Jbabdi (2006); Perrin (2006); Lenglet (2006)], which are great sources for a general understanding of the neural tissue and cerebral anatomy.

Keywords: gray matter, white matter, fiber tracts, projection fibers, association fibers, commissural fibers

Organization of this chapter:

The chapter is organized as follows. We first give a brief overview of the human brain structure and neural tissue in Section 3.1. Then, the organization of the white matter is described in Section 3.2. Finally, the important white matter fiber tracts for this thesis are reviewed in Section 3.2.1.

3.1 THE HUMAN BRAIN AND NEURAL TISSUE

The human brain is extremely complex. We know very little of the complex functioning of the brain but we know quite a lot about its anatomy [Gray (1918)]. The anatomy of the brain has been studied for more than one hundred years. The brain controls the central nervous system (CNS), the peripheral nervous system (PNS) and regulates all human activity. Neural tissue is specialized for the conduction of electrical impulses that convey information or instructions from one region of the body to another. How is this information carried through the white matter and how are the different parts of the brain connected remains unknown. About 98% of neural tissue is concentrated in the brain and spinal cord, the control centers for the nervous system.

At a microscopic scale, neurons transmit signals as electrical impulses which affect their cell membrane potentials. The electrical impulse usually affects the cell membrane potential of one of the neuron's dendrites and then eventually travels along the length of this axon to transmit to other neurons. Axons are often called nerve fibers and bundles of these axons are called fiber tracts. In this thesis, we will




Figure 3.1: Coronal slice taken from a human specimen from the brain museum (www.brainmuseum.org). The white matter is in light white color and the gray matter in purple. (a) shows regions of the ventricles containing CSF. (b) shows a transverse slice taken in the white matter where we see small and larger circular axons tightly packed together. The image was taken in the corpus callosum of a monkey's brain [Stanisz et al. (1997)]. (c) shows a sketch of the axons filled with microtubules and neurofilaments in the longitudinal axis, taken and adapted from [Williams et al. (1997); Mori and van Zijl (2002)]. (d) shows the gray matter as a thin purple layer on the outside of the brain.

be most interested in recovering information about these fiber tracts.

At a larger scale, the human brain is made of different elements such as blood, cerebrospinal fluid (CSF), white matter and gray matter. All these different element classes produce a different signal contrast under magnetic resonance imaging (MRI). Figure 3.1 shows a typical coronal slice of the human brain with some of the tissue classes highlighted. First, the CSF occupies the ventricular system around and inside the brain. It acts as a "cushion" for the cortex, providing also a basic mechanical and immunological protection to the brain inside the skull. Second, the white matter is a major component of the CNS. It is composed of axonal nerve fibers or neurons covered




Figure 3.2: The four lobes of the cerebral cortex and the cerebellum. Figure adapted from [Gray (1918)].

by a myelin sheath giving it its distinctive color. White matter axons are often tightly packed together and highly organized. This is seen in Figure 3.1(b) with a transverse slice of the corpus callosum (CC) of the monkey's brain and in Figure 3.1(c) with a schematic view of the microtubes and neurofilaments that fill the axon. White matter connects various gray matter areas (the locations of nerve cell bodies) of the brain to each other and carries nerve impulses between neurons. Finally, the gray matter is also a major component of the CNS. Gray matter is composed of unmyelinated neurons as opposed to white matter. It has a gray brown color which comes from the capillary blood vessels and the neuronal cell bodies. The cerebral cortex is the most important structure of the gray matter and plays a major role in cognitive functions. It constitutes the outermost layer of the brain and is highly folded to increase its surface in the limited volume of the skull. This folding process creates grooves on the surface of the brain called sulci and ridges called gyri. About two thirds of the cortical surface is buried in those sulci.

The two hemispheres of the brain are separated by a prominent central fissure. Each hemisphere of the cerebral cortex is made of four lobes (Figure 3.2). 1) The frontal lobe, located in front of the central sulcus, play an important role in reasoning, planning, language, memory and motor control. 2) The parietal lobe, behind the central sulcus, are more important to integrate sensory information and process some visuo-spatial stimuli. 3) The temporal lobes, on the most lateral parts of the cortex and below the lateral fissure, are involved in auditory processing as well as language




Figure 3.3: Sketch of a typical neuron or axonal nerve fiber in (a). In (b), we see an electron micrograph in a longitudinal section of a freshly excised garfish. We see tightly packed axons running diagonally across the micrograph. Image (a) is adapted from Wikipedia and image (b) is taken from [Beaulieu (2002)].

and vision related functions. 4) The occipital lobe occupy the rearmost part of the cortex and are the processing centers of visual stimuli.

In this thesis, we are interested in the “wiring” or architecture of the human brain. What are the connections within and between each lobe? How are the different lobes connected with each other? How are they connected to the basal nuclei of the brain? This wiring is done by the white matter connection or fiber tracts. We now describe the organization of the white matter in more detail and describe the important white matter fiber tracts for the rest of this thesis.

3.2 ORGANIZATION OF THE WHITE MATTER _____

As mentioned above, white matter is composed of neurons and axonal nerve fibers. These are covered by a myelin sheath giving it its distinctive color. The human nervous system is made of about 100 billion neurons. A typical neuron is illustrated in Figure 3.3. The axon conducts electrical impulses away from the neuron’s cell body or soma, to the axonal endings and dendrites of other connected neurons. The axons form the wiring of the nervous system and in particular, the white matter architecture. The white matter axons can be distributed diffusely or concentrated in bundles, also called fiber tracts. This is clear from real transverse and longitudinal slices of Figure 3.1(b) and Figure 3.3(b). Three main types of neural tracts are found in the white matter:

Projection tracts establish connections between the cerebral cortex and subcortical structures, such as the basal ganglia and the thalamus. There are two types of

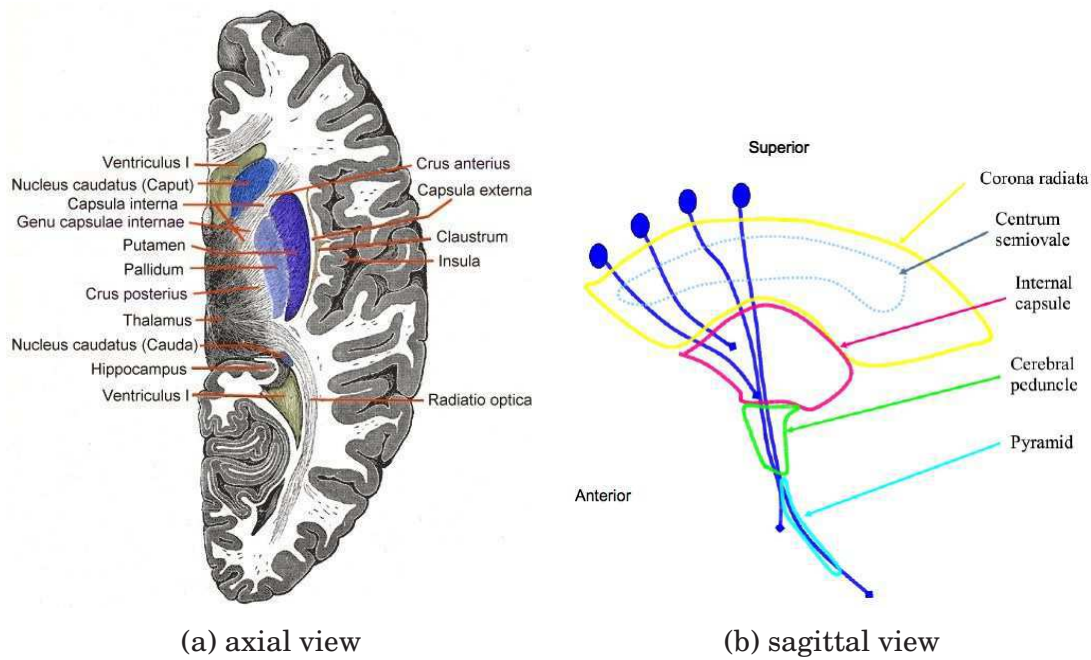




Figure 3.4: The projection fibers pass through and around the basal nuclei (putamen, pallidum, caudate nucleus, thalamus) seen in (a). The projection fibers form parts of the internal, external and extremal capsules and form the coronal radiata, as seen in (b). Figure (a) was taken from [Gray (1918)] and figure (b) was taken from [O'Donnell (2006)].

projections tracts: afferent tracts and efferents tracts. Afferent tracts carry information from different parts of the body to the cerebral cortex. All sensory information, except olfactory, end up in the primary sensory cortex by the means of the optic and acoustic fibers. The optic, acoustic and somatosensory tracts pass in the thalamus before projecting to the cortex. The optic tract projects to the occipital lobe and the acoustic tract to the temporal lobe. There are also numerous thalamo-cortical afferent fibers arising within the thalamus and projecting to the different parts of the cortex.

Efferent tracts carry motor commands from the motor cortex down to the muscles and glands through the lower brain structures and the spinal cord. They reach structure like the basal ganglia, the cerebellum, the brainstem and the spinal cord. The basal ganglia include the striatum (which consist of the putamen, the caudate nucleus and the ventral striatum), the globus pallidus, the subthalamic nucleus and the substantia nigra. Other important efferent fibers include the motor tracts, which occupies a large part of the genu and anterior occipital part of the internal capsule and the corticopontine fibers. Some of these fiber tracts and basal nuclei are illustrated in Figure 3.4.



(a) (b)
Short association fibers



(c) (d)
Long association fibers

Figure 3.5: The association fibers. Short association fibers in (a)-(b) and long association fibers in a sagittal view (c)-(d). Figure (a) and (c) are adapted from [Perrin (2006)] and figure (d) was taken from [Nieuwenhuys et al. (1981)]. In (b), the very short tangential fibers in the gray matter are sketched in region A, the 'U' fibers sketched in region B and the short association fibers in region C.

Association tracts are the communication fibers between different cortical areas within a given hemisphere. They can be divided into two categories: short and long association tracts. Short association tracts build up connections between regions within a given lobe and connect adjacent gyri. The smallest link adjacent cortical zones separated by a sulcus. They are called the 'U'-fibers. The short association fibers lie immediately beneath the gray substance of the cortex of the hemispheres, and connect together adjacent gyri. Long association fibers establish connections between different cerebral lobes passing between more distant parts. The long association fibers include fibers such as the uncinat fasciculus (UNC) from frontal to temporal lobe, the cingulum (Cg) from the cingulate gyrus to the entorhinal cortex, the superior fronto-occipital (SFO) fasciculus from the parietal lobe to frontal lobe, the inferior fronto-occipital (IFO) fasciculus from the occipital lobe to frontal lobe, the superior longitudinal fasciculus (SLF) from the frontal lobe to occipital lobe, the infe-

rior longitudinal fasciculus (ILF) from occipital lobe to temporal lobe and the fornix from the hippocampus to the mammillary body. Figure 3.5 sketches some of these short and long association fibers.

Commissural tracts are bundles of axons connecting a region in one hemisphere to another region of the opposite hemisphere. The important commissural fibers include the transverse fibers of the CC, also called the transcallosal fibers, the anterior commissure (AC) and the posterior commissure (PC). The CC and AC are illustrated in Figures 3.6 and 3.10.

3.2.1 Important white matter bundles for this thesis

We now review the important fiber tracts that will be examined across the rest of this thesis.

Corpus Callosum


The CC is a structure that connects the left and right cerebral hemispheres. It is the largest white matter structure in the brain, consisting of 200-250 million contralateral axonal projections or transcallosal fibers. It is a wide, flat bundle of axons beneath the cortex. Much of the inter-hemispheric communication in the brain is conducted across the CC. The CC is often divided in four parts, as seen in Figure 3.6(a). From most anterior to posterior, the CC is divided into the rostrum, genu, body and splenium.

On either side of the CC, the fibers radiate and pass to the various parts of the cerebral cortex. The transcallosal fibers curving forward from the genu into the frontal lobe constitute the forceps minor (also called anterior forceps) and those curving backward into the occipital lobe constitute the forceps major (also called posterior forceps). Between these two parts is the main body of the fibers which constitute the tapetum (TAP) and extend laterally on either side into the temporal lobe, and cover the central part of the lateral ventricle. This is shown in Figure 3.6. The CC will come in most chapters. In particular, Chapter 11 is entirely dedicated to the recovery of the transcallosal fibers.


Corticospinal Tract (CST)

The corticospinal or pyramidal tract is a massive collection of axons that travel between the cerebral cortex of the brain and the spinal cord. The CST mostly contains the motor axons. The CST is illustrated in Figure 3.7.

Closely related to the CST is another important fiber bundle, the corona radiata (CR). The CR is associated with the CST. In neuroanatomy, the CR refers to the descending bundle or sheet of axons that are involved with the fine coordination of movement. The CR is seen in Figure 3.8. Figure 3.8 also shows the anterior thalamic



(a) CC division in a coronal view



(b) Transcallosal fibers and forceps

(c) CC, TAP and SLF

Figure 3.6: The corpus callosum. In (a) we see the division of the CC in a zoomed sagittal view. In (b)-(c) we see a top view where we see the transcallosal fibers, the forceps and the tapetum (TAP). Figure (a) and (b) are adapted from [Williams et al. (1997); Perrin (2006)] and figure (c) is adapted from [Gray (1918)].

radiations (ATR). The ATR are radiations that pass through the internal capsule. The ATR interconnects the frontal lobe and other cerebral regions such as the pons, the thalamus and the basal ganglia.

Superior and Inferior Longitudinal Fasciculus

The SLF and ILF are important association fiber seen above. The SLF connects the front and back of the cerebral cortex. The ILF connects the temporal lobe and the occipital lobe, running along the lateral walls of the ventricles. Figure 3.9 shows these two fiber bundles.

Figure 3.9 also shows two other important association fibers, the uncinate fasciculus and the cingulum. The UNC passes across the bottom of the lateral fissure, and

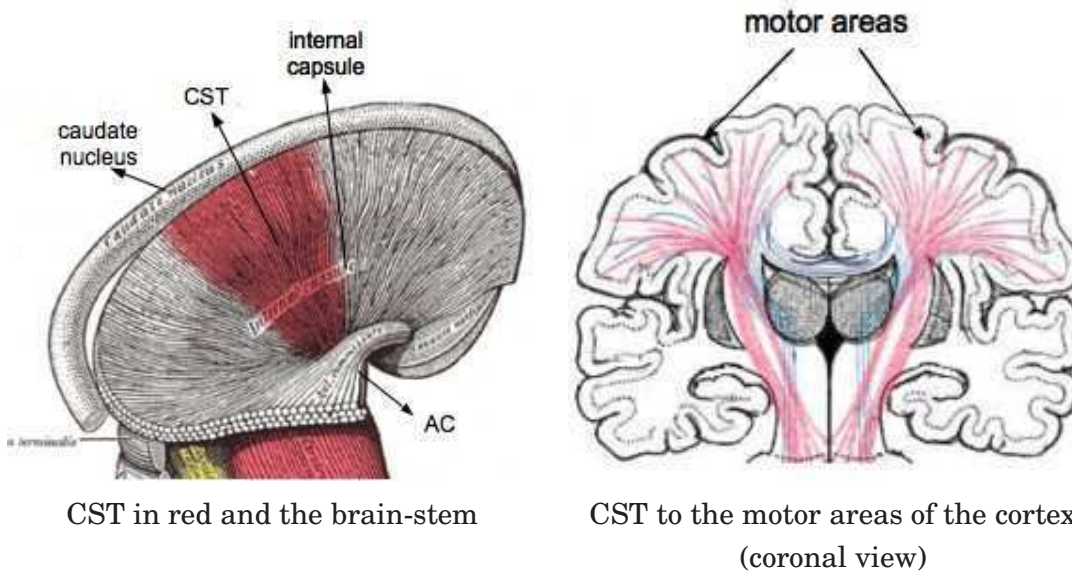


Figure 3.7: The corticospinal tract (CST). Both figures are adapted from [Gray (1918)].




Figure 3.8: The corona radiata (CR), the corticospinal tract CST and the anterior thalamic radiations (ATR) in a sagittal view. Figure adapted from [Williams et al. (1997)].

connects the gyri of the frontal lobe with the anterior end of the temporal lobe. The Cg is a collection of fibers projecting from the cingulate gyrus to the entorhinal cortex in the brain. It allows the communication between components of the limbic system and the cortex.

The Anterior Commissure

The AC is a fiber bundle connecting the two cerebral hemispheres across the middle line. The AC is placed in front of the columns of the fornix. The AC has an oval cross-

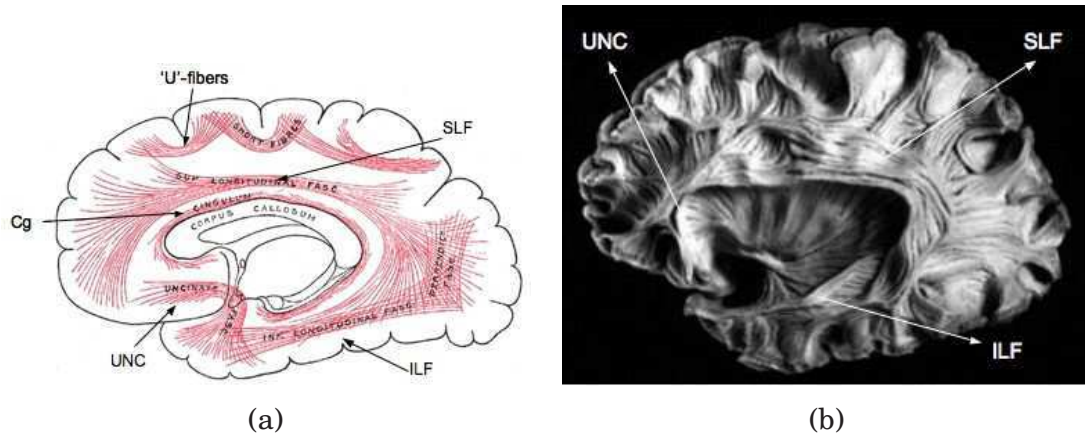


Figure 3.9: Association fibers in details. The superior longitudinal fasciculus (SLF), the inferior longitudinal fasciculus (ILF) and the uncinete fasciculus (UNC) are seen in (a)-(b) in a sagittal view. In (a), we also see short corticocortical fibers ('U' fibers), the CC and the cingulum (Cg). (a) was taken from [Gray (1918)] and (b) adapted from [Williams et al. (1997)].

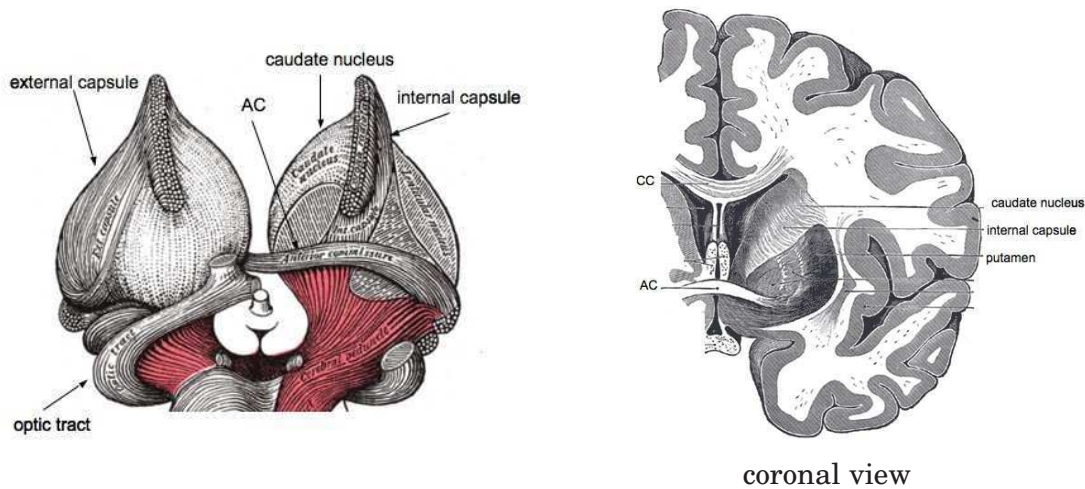


Figure 3.10: Illustration of the commissural fibers and the anterior commissure (AC). Both figures were taken from [Gray (1918)]. In the right figure, we also see the CC, the internal capsule, the caudate nucleus and putamen in a coronal view.

section with a long vertical diameter that measures about 5 mm. Its fibers can be traced laterally and backwards on either side into the substance of the temporal lobe. It thus serves in this way to connect the two temporal lobes. The AC is illustrated in Figure 3.10.

3.3 CONCLUSION

This brief background chapter has introduced the important neural tissue,

brain tissues and brain structures that will be examined in the rest of this thesis. More particularly, in this thesis, we are mostly interested in the white matter, from the local reconstruction of the neural tissue using diffusion MRI to the recovery of information about the wiring of the human brain. We have introduced the necessary nomenclature of small and larger fiber bundles that will come back in most chapters of the thesis.

In the next chapter, we introduce the foundations of diffusion MRI. We will show that diffusion MRI constitutes a powerful non-invasive mean to investigate the architecture of the human brain white matter that we have just described.

PRINCIPLES OF DIFFUSION MRI: GOING BEYOND THE DIFFUSION TENSOR

“A water story”

–Denis Le Bihan

“Water is the driving force of all nature”

–Leonardo da Vinci

Contents

4.1 Basic Principles of Diffusion MRI	36
4.1.1 Brownian Motion of Water Molecules and Fick’s Law	37
4.2 NMR and Diffusion MRI	39
4.2.1 Pulse Gradient Spin Echo (PGSE)	40
4.2.2 Diffusion-Weighted Imaging (DWI)	42
4.2.3 Apparent Diffusion Coefficient (ADC) and Trace Imaging . . .	43
4.2.4 Diffusion Tensor Imaging (DTI)	44
4.2.5 High Angular Resolution Diffusion Imaging (HARDI)	46
4.3 Multiple Fiber HARDI Reconstruction Techniques	50
4.3.1 Diffusion Spectrum Imaging (DSI)	51
4.3.2 Single Shell HARD Imaging	53
4.4 Conclusion	60

OVERVIEW

How do we image white-matter non-invasively? This chapter covers the principles of DW-MRI and high angular resolution diffusion imaging (HARDI) reconstruction techniques. In particular, what are the physical principles of DW-MRI? What are the fiber reconstruction algorithms to infer microstructure of biological tissues? First, we describe the basic principles of DW-MRI. Next, we review the different DW-MRI acquisition techniques. Properties of each acquisition technique are described separately. Finally, a special emphasis is put on the approaches designed to reconstruct the multiple fiber distributions of water diffusion. Many functions are used in the literature to measure the diffusion properties of the underlying biological tissue. Some methods are model dependent, some model-free, some methods have linear solutions whereas others require non-linear optimization schemes. These different methods are described. Overall, this introductory chapter is inspired from review articles and chapters from [Tuch (2002); LeBihan (2003); Campbell (2004); Alexander (2005a); Hagmann et al. (2006b); Lenglet (2006)], which are great sources for a general understanding of the diffusion MRI field.

Keywords: Nuclear magnetic resonance (NMR), magnetic resonance imaging (MRI), diffusion MRI, diffusion tensor imaging (DTI), high angular resolution diffusion imaging (HARDI)

Organization of this chapter:

The chapter is organized as follows. We first review the basic principles of molecular diffusion in Section 4.1 and then diffusion NMR and diffusion MRI in Sections 4.2. We then describe the state-of-the-art multiple fiber HARDI reconstruction algorithms in Section 4.3.

4.1 BASIC PRINCIPLES OF DIFFUSION MRI

Diffusion MRI is a relatively recent field of research with a history of more or less twenty years. Diffusion MRI is of growing interest because it helps understand functional coupling between cortical regions of the brain, which is useful in characterization of neuro-degenerative diseases, in surgical planning and in other medical applications. The great success of diffusion MRI comes from its capability to accurately describe the geometry of the underlying microstructure. To do so, diffusion MRI captures the average diffusion of water molecules, which probes the structure of the biological tissue at scales much smaller than the imaging resolution. The diffusion of water molecules is Brownian under normal unhindered conditions, but in fibrous structure such as white matter, water molecules tend to diffuse along fibers. Due to this physical phenomenon, diffusion MRI is able to obtain information about the neural architecture in vivo. It is also the only imaging modality able to do




Figure 4.1: An image of Brownian motion, done with three different step sizes. The hierarchical structure is clearly visible. More saturated colors represent smaller step sizes. Image under the Gnu Free Documentation License 1.2

so non-invasively. We now review the basic physical principles of diffusion MRI.

4.1.1 Brownian Motion of Water Molecules and Fick's Law

At a microscopic scale, water molecules freely move and collide with each other in an isotropic medium according to Brownian motion [Brown (1828)]. This is illustrated in Figure 4.1.

At a macroscopic scale, this phenomenon yields a diffusion process. In an isotropic medium, the diffusion coefficient D was related by Einstein [Einstein (1956)] to the root mean square of the diffusion distance as

$$D = \frac{1}{6\tau} \langle \mathbf{R}^T \mathbf{R} \rangle, \quad (4.1)$$

where τ is the diffusion time, $\langle \rangle$ denotes an ensemble average and $\mathbf{R} = \mathbf{r} - \mathbf{r}_0$ is the net displacement vector, with \mathbf{r}_0 the original position of a particle and \mathbf{r} its position after the time τ . The scalar constant D , known as the diffusion coefficient, measures the molecule's mobility. In the isotropic case, it depends on the molecule type and the medium properties but not on the direction.

The macroscopic process of diffusion can also be described by Fick's first law, derived by Adolf Fick in 1855 [Fick (1855)]. It relates the concentration difference of the diffusion substance C to a flux. The flux, \mathbf{J} , is proportional to the gradient of the

concentration, ∇C . The proportionality constant D is the diffusion coefficient and the governing equation is given by

$$\mathbf{J} = -D\nabla C. \quad (4.2)$$

In anisotropic biological tissues, a water molecule's mobility is constrained by obstacles formed by surrounding structures, such as the axons in the brain. For example, it is known that the myelin sheath can modulate the anisotropy of the diffusion while the microtubules and neurofilaments do not modify it [Beaulieu (2002)]. The general diffusion displacement probability density function (PDF), also called diffusion propagator, of water molecules is extremely complex and is still unknown today. Hence, historically, simple models of diffusion have been proposed. Amongst these models, the most popular model is certainly the diffusion tensor (DT) model [Basser et al. (1994b)] arising from the transport tensors theory that had been studied in the mathematics literature for a long time [?]. In this simplified model of water diffusion, Einstein's and Fick's law of diffusion can be generalized using the DT model. The scalar diffusion coefficient D can be replaced by a DT, \mathbf{D} . Hence, Einstein's relation 4.1 can be generalized considering the covariance matrix of the net displacement vector \mathbf{R}

$$\mathbf{D} = \begin{pmatrix} \mathbf{D}_{xx} & \mathbf{D}_{xy} & \mathbf{D}_{xz} \\ \mathbf{D}_{xy} & \mathbf{D}_{yy} & \mathbf{D}_{yz} \\ \mathbf{D}_{xz} & \mathbf{D}_{yz} & \mathbf{D}_{zz} \end{pmatrix} = \frac{1}{6\tau} \langle \mathbf{R}\mathbf{R}^T \rangle. \quad (4.3)$$

The diffusion PDF is written as $P(\mathbf{R}, t)$. It represents the probability that a water molecule located at \mathbf{r}_0 will have moved by amount $\mathbf{R} = \mathbf{r} - \mathbf{r}_0$ in time τ . In the diffusion tensor model, if we take the Taylor's expansion of P about \mathbf{R} and τ and ignore the higher order terms, we can use Einstein's relation of Eq. 4.3 to obtain

$$\frac{\partial P(\mathbf{R}, t)}{\partial t} = \mathbf{D}\nabla^2 P(\mathbf{R}, \tau). \quad (4.4)$$

The detailed intermediate steps and mathematical derivations are done in [Campbell (2004)]. The solution to Eq. 4.4 is P , the diffusion PDF of water molecules under the DT model assumption.

In Fick's formulation, the diffusion tensor \mathbf{D} can also be introduced to model the anisotropic diffusion of the biological tissue. We then have,

$$\mathbf{J} = -\mathbf{D}\nabla C \quad (4.5)$$

If we use Fick's Law and the law of conservation of mass, $\partial C/\partial t = -\nabla \cdot \mathbf{J}$, we get

$$\frac{\partial C}{\partial t} = \nabla \cdot (\mathbf{D}\nabla C) = \mathbf{D}\nabla^2 C, \quad (4.6)$$

which is the same as Eq. 4.4 with the diffusion PDF P replaced by the concentration of the medium C . Both the Einstein's Eq. 4.4 and Fick's Eq. 4.6 describe the classical

diffusion equation. In the case D is the identity matrix (isotropic medium), it is called the *heat* equation and in the anisotropic case, it is called the *geometric heat* equation.

Both the isotropic and anisotropic description of diffusion are used in many diffusion MRI applications for different normal and pathological tissue studies. Note that the diffusion coefficient measured in diffusion MRI is not the true diffusion coefficient of water unless the imaging object is pure water. The diffusion is restricted by the structure of the tissue and this is normally called the *apparent diffusion coefficient* (ADC). We now detail how we can acquire DWI, measure the ADC and measure the diffusion PDF using diffusion NMR. Note that in this section we have derived the diffusion equations using the diffusion tensor model. However, for the rest of this chapter we will treat the diffusion PDF, P , in a general model-free fashion.

4.2 NMR AND DIFFUSION MRI ---

A brief history

A long time before MRI and diffusion MRI were invented, NMR was used to measure properties of substances in chemistry to perform NMR spectroscopy.

NMR was simultaneously described in 1946 by Felix Bloch [Bloch (1946)] and by Edward Mills Purcell [Purcell et al. (1946)]. In 1952, they both received the Nobel Prize in Physics for their discovery. The basic principle behind NMR is that after aligning a magnetic nucleus such as 1H (the proton) with a very strong external magnetic field, its response to a perturbation of the alignment by an electromagnetic field is characteristic. Four years after this discovery, in 1950, Herman Carr, proposed to create the first one-dimensional (1D) MR images by introducing a gradient in the magnetic field [Carr and Purcell (1954)]. Shortly after Bloch and Purcell discovery, Hahn published his seminal paper [Hahn (1950)] on the NMR spin echo in which he noted that the random thermal motion of the spins would reduce the amplitude of the observed signal in the presence of a magnetic field inhomogeneity. This is a fundamental notion to understand diffusion MRI.

As soon as 1973, Paul Lauterbur proposed a method [Lauterbur (1973)], based on gradients of magnetic field, to reconstruct two dimensional MR images. Peter Mansfield [Mansfield (1977)] further developed the use of magnetic fields gradients and, by studying the mathematical properties of the MRI signal, proposed a new ultrafast acquisition technique known as the echo-planar technique. In 2003, they jointly received the Nobel prize in Physiology and Medicine for their discoveries on MRI. At the heart of diffusion MRI are the principles of NMR just highlighted but also the gradient spin echo sequence of [Stejskal and Tanner (1965)] that we now describe. A schematic view of our brief history of diffusion MRI is shown in Figure 4.2.

Principles of Nuclear Magnetic Resonance (NMR)

[Bloch, Purcell, Carr, Hahn, Lauterbur, Mansfield, ...]



The Pulse Gradient Spin-Echo (PGSE) sequence

[Stejskal and Tanner (1965)]



Scalar Diffusion-Weighted Imaging (DWI)

DWI

[Wesbey et al (1984), Taylor and
Bushell (1985), Le Bihan and Breton (1985)]

ADC & Trace Imaging

[Moseley et al 1990, Douek et al 1991]



Diffusion Tensor Imaging (DTI)

[Basser et al, (1992-1994)]



Beyond Tensor Imaging

Single shell HARDI

[Tuch et al (1999)]

Diffusion Spectrum Imaging (DSI)

[Callaghan et al (1988),

van Wedeen et al (2000)]

Figure 4.2: Brief history of diffusion MRI.

4.2.1 Pulse Gradient Spin Echo (PGSE)

The Stejskal-Tanner imaging sequence [Stejskal and Tanner (1965)] is used to measure the diffusion of water molecules in a given direction \mathbf{g}_i , $i = 1, \dots, N$. This pulse sequence is illustrated in Figure 4.3. This sequence uses two gradient pulses $\mathbf{g}(t)$ in the direction \mathbf{g} , of duration time δ , to control the diffusion-weighting. They are placed before and after a 180° degrees refocusing pulse. More specifically, a first 90° degrees RF is applied to flip the magnetization in the transverse plane. The first gradient pulse causes a phase shift of the spins whose position are now a function of time. Spin position is in fact assumed to stay constant during time δ . Finally, the 180° pulse combined with the second gradient pulse induces another phase shift. It is applied after a time Δ separating the two gradient pulses. This pulse cancels the first phase

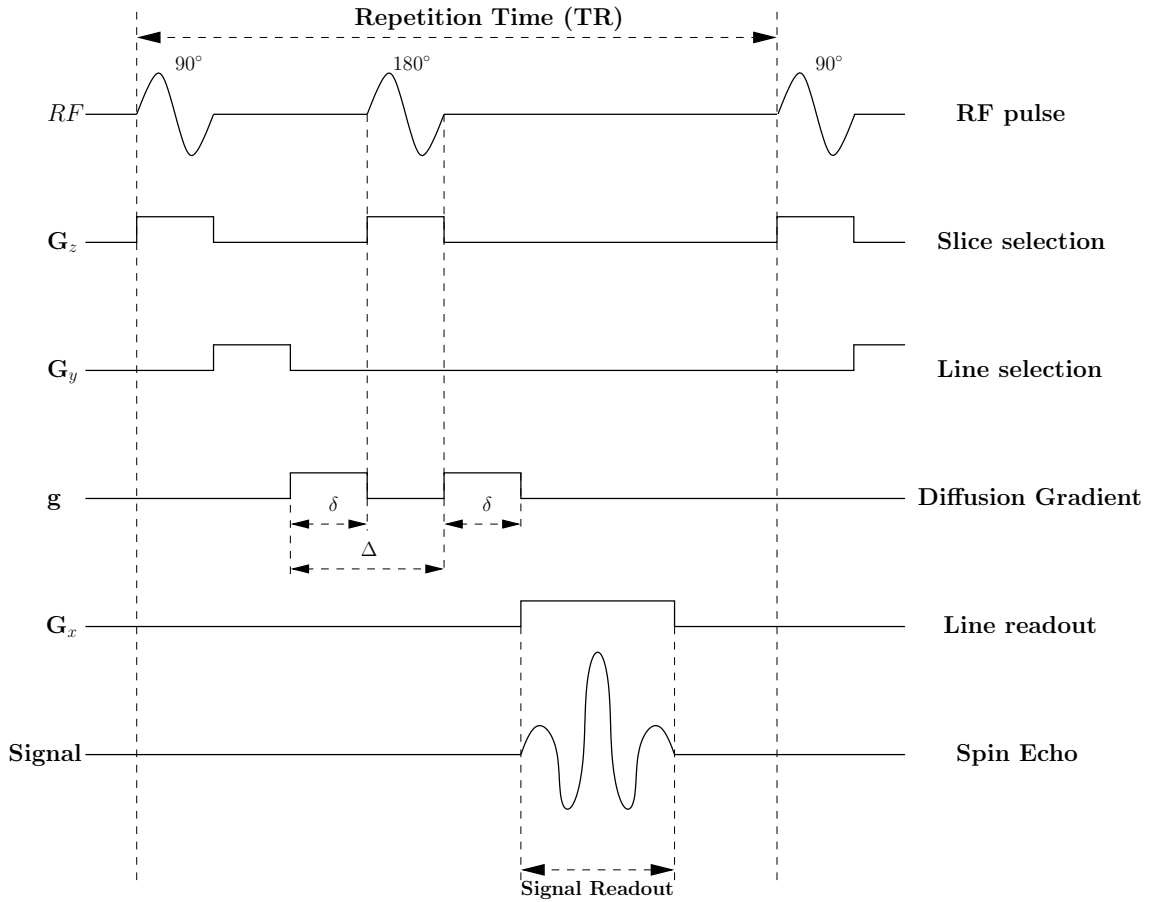


Figure 4.3: Schematic Stejskal-Tanner imaging PGSE sequence. RF pulses could more realistic.

shift only for static spins. On the other hand, spins under Brownian motion during the time period Δ separating the two pulses undergo different phase shifts by the two gradient pulses, resulting in a T2 signal attenuation [Cercignani and Horsfield. (2001)].

By assuming the pulses to be infinitely narrow (*narrow pulse approximation*), i.e. if the gradient pulse duration δ is short enough for the diffusion of the water molecule to be negligible during that time, [Stejskal and Tanner (1965)] showed that the signal attenuation $S(\mathbf{q}, \tau)$ is expressed as the 3-dimensional (3D) Fourier transform \mathcal{F} of the ensemble average propagator P ,

$$\frac{S(\mathbf{q}, \tau)}{S_0} = \int_{\mathbb{R}^3} P(\mathbf{r}|\mathbf{r}_0, \tau) \exp(-2\pi i \mathbf{q}^T \mathbf{R}) d\mathbf{r} = \mathcal{F}[P(\mathbf{r}|\mathbf{r}_0, \tau)], \quad (4.7)$$

where the value of \mathbf{q} is given by $\mathbf{q} = \gamma \delta \mathbf{G} / 2\pi$, with γ the nuclear gyromagnetic ratio for water protons, \mathbf{G} the applied diffusion gradient vector, S_0 is the baseline image acquired without any diffusion gradients (also called $b = 0$ image) and $P(\mathbf{r}|\mathbf{r}_0, \tau)$ is the diffusion PDF or diffusion propagator of water molecules introduced earlier. This P is ultimately the function we are looking to reconstruct in diffusion MRI. Intuitively,

one has to sample the diffusion PDF along many \mathbf{q} vectors to be able to reconstruct the diffusion PDF. The space of all possible 3D \mathbf{q} vectors is called *q-space*. This is the idea behind q-space imaging [Callaghan (1991)].

If the diffusion PDF is assumed to be Gaussian, one can work out the Fourier integral in Eq. 4.7 analytically. The Stejskal-Tanner signal attenuation equation then becomes

$$S(\mathbf{q}, \tau) = S_0 e^{-\tau \mathbf{q}^T \mathbf{D} \mathbf{q}}, \quad (4.8)$$

where $D(\mathbf{q}) = \mathbf{q}^T \mathbf{D} \mathbf{q}$ is the ADC. The signal attenuation is also often written with respect to unit vector, $\mathbf{g} = \mathbf{q}/|\mathbf{q}|$, and it is common to introduce the *b-value*, $b = \tau |\mathbf{q}|^2$. We thus obtain a signal attenuation with respect to the *b-value* given by

$$S(\mathbf{q}, \tau) = S_0 e^{-\tau |\mathbf{q}|^2 \mathbf{g}^T \mathbf{D} \mathbf{g}} \iff S(b, \mathbf{g}) = S_0 e^{-b \mathbf{g}^T \mathbf{D} \mathbf{g}} = S_0 e^{-b \text{ADC}}. \quad (4.9)$$

This is the most common formulation of the Stejskal-Tanner equation under a Gaussian assumption. Thus, for large *b-values*, the true signal quickly falls off, while the background noise is relatively unaffected, resulting in very noisy data measurements. Figure 4.4 shows DW images for different *b-values*. We note the importance of the *b-value*. One has to appropriately tune the *b-value* to avoid either a very low signal attenuation when *b* is too small or a poor signal-to-noise ratio (SNR) when *b* is too high.

To conclude, it is important to remember that the ratio $S(\mathbf{q}, \tau)/S_0$ has a Fourier relationship with the diffusion PDF, which is not necessarily Gaussian. This is a key observation that is at the heart of q-space imaging using high angular resolution diffusion imaging, since it potentially gives access to the complex diffusion profile of water molecules at each voxel. However, the actual computation of the inverse Fourier transform to obtain the diffusion PDF is difficult in practice and has given rise to many research developments of alternative acquisition and reconstruction techniques. Historically, people have first measured scalar-valued DWI, then, tensor-valued DWI and finally 3D images of the the full diffusion PDF at each voxel. We now describe these different diffusion MRI acquisition techniques.

4.2.2 Diffusion-Weighted Imaging (DWI)

It was in 1984-1985 that diffusion NMR and imaging techniques were put together to make the first diffusion MRI acquisitions. Wesbey et al were the first to demonstrate DWI, but their sequence was not clinically feasible [Wesbey et al. (1984b,a)]. A little later, the first DWI acquisition was done by Taylor and Bushell using a hen's egg as a phantom in a small bore magnet [Taylor and Bushell (1985)]. Then, Le Bihan et al did the first DWI acquisition in vivo of the human brain using a whole-body scanner [LeBihan and Breton (1985); LeBihan et al. (1986)] and introduced the famous *b-value*. At that time, a DWI was simply the unprocessed result of the application of the PGSE sequence in one gradient direction (Figure 4.5). This image corresponded to a single point in q-space [Hagmann et al. (2006b)].




Figure 4.4: Diffusion-weighted images for different b -values. DWI taken from one of the subjects of the HARDI database described in Appendix A.3.4.

4.2.3 Apparent Diffusion Coefficient (ADC) and Trace Imaging

It was not long after the first DWI acquisition, in 1990, that diffusion anisotropy measurements were first done in more than one direction by [Moseley et al. (1990)] in cat brain. It was known that diffusion truly was a 3D process with a faster diffusion presumed in the direction of fibers than in perpendicular directions. Hence, there was no reason that the diffusion in biological tissue was the same in all directions.

Assuming a Gaussian diffusion PDF, the log version of the Stejskal-Tanner equation with respect to the b -value (Eq. 4.9) was written as

$$\ln \left(\frac{S(b, \mathbf{g})}{S_0} \right) = -b \cdot \text{ADC}. \quad (4.10)$$

[Moseley et al. (1990)] proposed to measure this ADC along the x and z axis and take the anisotropy index $\text{ADC}_z / \text{ADC}_x$ to characterize the level of anisotropy in the tissue. Hence, at least two DWI acquisitions were required at that point of time. [Douek et al. (1991)] then suggested that diffusion MRI and measures of the ADC





Figure 4.5: Axial slice of diffusion-weighted images (DWI) with two different diffusion gradient directions (red arrows). MR signal attenuation is found in regions having fibers mostly aligned with diffusion gradient direction (yellow arrows).

along two directions could be used to determine orientations of fiber bundles in the white matter. Note that [Wesbey et al. (1984b,a)] had also proposed the formula to measure the ADC in one direction of each voxel, but their sequence was not clinically feasible.

However, these ADC measures and anisotropy index were very dependent on the direction of the gradient encoding used in the acquisition. Figure 4.5 shows examples of diffusion weighted images acquired with two different directions. It illustrates the direction specific attenuation related to white matter fibers orientation. The ADC and the proposed anisotropy measures were then confounded to be rotationally variant. Therefore, it was starting to be clear that rotationally invariant measures were needed. It was also clear that a mathematical model with invariant properties was required to describe the 3D displacement distribution of water diffusion. As we will see in the next section, DTI, proposed by Peter Basser, was introduced to overcome this limitation and rotationally invariant measures were defined [Basser et al. (1994b); Pierpaoli and Basser (1996); Pierpaoli et al. (1996)].

4.2.4 Diffusion Tensor Imaging (DTI)

In 1992, Basser et al proposed to use a second order symmetric and positive-definite tensor to model the intrinsic diffusion properties of biological tissues [Basser et al. (1992, 1993, 1994b,a); Basser (2002); Basser and Jones (2002)]. This is the same diffusion tensor (DT) as encountered earlier in Einstein's Eq. 4.3 and Eq. 4.5 for anisotropic diffusion. In fact, if one solves for P in Eq. 4.4, the solution obtained is the Gaussian diffusion PDF. In this case, the probability P to find a molecule, ini-



H₂O Brownian motion along the fibers [Poupon (1999)] Eigen decomposition of the DT Ellipsoidal visualization of the DT

Figure 4.6: Diffusion tensor representation.

tially at position \mathbf{r}_0 , at \mathbf{r} after a delay τ is then given by

$$P(\mathbf{r}|\mathbf{r}_0, \tau) = \frac{1}{\sqrt{(4\pi\tau)^3 |\mathbf{D}|}} \exp\left(-\frac{1}{4\tau}(\mathbf{r} - \mathbf{r}_0)^T \mathbf{D}^{-1}(\mathbf{r} - \mathbf{r}_0)\right), \quad (4.11)$$

where $|\mathbf{D}|$ is the determinant of the DT, \mathbf{D} . In other words, DTI approximates the diffusion PDF by a 3-variate normal distribution with zero mean. Hence, \mathbf{D} can be viewed as the covariance matrix describing the Brownian motion of water molecules at each imaging voxel.

Since \mathbf{D} is symmetric (Eq. 4.3), it has six unknown coefficients that we need to estimate. Hence, DTI needs at least six DW images and one unweighted diffusion image ($b = 0$ s/mm²) to solve the system of equations. Typically, a b -value of 1000 s/mm² is used with 7 to 60 gradient directions. The DT estimation problem has been the subject of much research and several theses in the past decade. DTI estimation methods go from classical linear and non-linear least-squares [Basser et al. (1994a)] to much more sophisticated Riemannian frameworks [Lenglet (2006); Arsigny (2006); Fillard et al. (2007a)] that forbid degenerate tensors. There are trade-offs between speed, robustness to noise, symmetric diffusion guaranty of the DT and positive definite guaranty of the DT that one must consider when estimating the diffusion tensor. [Lenglet (2006)] and [Arsigny (2006)] give extensive details on this subject and [Fillard et al. (2007a)] reviews these contributions and introduces the latest state-of-the-art tensor estimation algorithms for clinical applications and to study brain variability [Fillard et al. (2007b)].

Most importantly, the DT is a rich mathematical tool with interesting properties that one can exploit for diffusion MRI visualization and analysis. The DT is system-

atically decomposed into its three eigenvalues $\lambda_1, \lambda_2, \lambda_3$, where $\lambda_1 \geq \lambda_2 \geq \lambda_3$, and corresponding eigenvectors $\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3$. The largest eigenvalue λ_1 gives the principal direction of the DT \mathbf{e}_1 and the other two eigenvectors span the orthogonal plane to it. This is illustrated in Figure 4.6. Note that we can have isotropic tensors when $\lambda_1 = \lambda_2 = \lambda_3$ and planar tensors when $\lambda_1 = \lambda_2$.

From this eigenvalue decomposition, several rotationally invariant quantities can be extracted such as the trace (equal to $\text{ADC}_x + \text{ADC}_y + \text{ADC}_z$), the mean diffusivity $\bar{\lambda} = \text{trace} / 3$, the fractional anisotropy (FA) and other measures [Westin et al. (2002)]. A very popular and extensively studied measure is the FA [Pierpaoli and Basser (1996)], defined as

$$\text{FA} = \sqrt{\frac{3}{2}} \sqrt{\frac{(\lambda_1 - \bar{\lambda})^2 + (\lambda_2 - \bar{\lambda})^2 + (\lambda_3 - \bar{\lambda})^2}{\lambda_1^2 + \lambda_2^2 + \lambda_3^2}}. \quad (4.12)$$

Most often the FA map itself is used to visualize the regions of anisotropy. DTs are also visualized with a Red-Blue-Green (RGB) colormap [Pajevic and Pierpaoli (1999)]. The RGB map encodes the x-y-z coordinates of the principal eigenvector \mathbf{e}_1 . The RGB map gives a good orientation description of the DT field. The areas of red, blue and green color indicate tensors aligned with x, y, z directions respectively. One can also visualize the principal eigenvector or the DT ellipsoids themselves. Figure 4.7 shows diffusion tensors, their corresponding principal diffusion directions, the FA map and RGB map in an axial slice of the MPI dataset described in Appendix A.3.3.

Even though the DT is a powerful tool which has earned success in many clinical applications [Mori and van Zijl (2002); Neil et al. (2002); Horsfield and Jones (2002); Sotak (2002); Lim and Helpert (2002)], tractography studies (see Chapter 10) and segmentation studies (see Chapter 8), the DT model is intrinsically limited when imaging voxels with multiple fiber populations crossing, branching, fanning or kissing¹[Basser et al. (2000)]. The DT model is limited because of the Gaussian PDF assumption and the limited number of degrees of freedom in the model. Hence, we need higher order models to be able to describe non-Gaussian distributions. This limitation of the DT model is illustrated in Figure 4.8 for two orthogonally crossing fibers. The expected fiber distribution has two maxima whereas the reconstructed DT profile is planar-like with no preferred diffusion direction.

4.2.5 High Angular Resolution Diffusion Imaging (HARDI)

The idea now is to sample q-space along as many directions and q-magnitudes as possible in order to reconstruct the true diffusion PDF. This true diffusion PDF is model-free and can recover the diffusion of water molecules in any underlying fiber population. For example, Figure 4.9 illustrates the expected diffusion PDF in the case of an isotropic imaging voxel, a single fiber imaging voxel and two crossing fibers

¹Figure 12.1 illustrates these different crossing sub-voxel configurations.




Figure 4.7: FA, RGB, field of diffusion tensors and principal diffusion direction of the DTs in an axial slice. The DTs are colored with respect to the FA map, blue to red corresponds to isotropic to anisotropic tensors. The DTs and e_1 vectors are shown where $FA > 0.1$ in half the axial slice.

imaging voxel. Note that we no longer have a scalar-valued or tensor-valued image but we now have a 3D image of 3D diffusion distributions. Hence, one can imagine that there are technical requirements and trade-offs that one must make in HARDI acquisition. In particular, HARDI depends on the number of measurements N and the gradient strength (b -value), which will directly affect acquisition time and signal to noise ratio in the signal. Typically, there are two strategies used in HARDI: 1) sam-




Figure 4.8: Limitation of DTI in voxels with crossing configurations. DTI cannot resolve imaging voxels containing multiple fiber crossings.

pling of the whole q-space 3D Cartesian grid or 2) single shell spherical sampling.²

In the first case, a large number of q-space points are taken over the discrete grid ($N > 200$) and the inverse Fourier transform of the measured DWI signal is taken to obtain an estimate of the diffusion PDF P . This is Diffusion Spectrum Imaging (DSI) [Wedeen et al. (2000)] in the HARDI literature and the theory of DSI goes back to the development of QSI by [Callaghan et al. (1988); Callaghan (1991)]. The method requires very strong imaging gradients ($500 \leq b \leq 20000 \text{ s/mm}^2$) and a long time for acquisition depending on the number of sampling directions (see Table 4.1). The visualization of 3D diffusion PDF at every voxel is computationally intensive. Hence, people either take an isosurface of the diffusion PDF for a certain radius r or the diffusion orientation distribution function (ODF) is computed. The diffusion ODF contains the full angular information of the diffusion PDF and is defined as

$$\Psi(\theta, \phi) = \int_0^\infty P(r, \theta, \phi) dr, \quad (4.13)$$

where (θ, ϕ) obey physics convention in this thesis ($\theta \in [0, \pi], \phi \in [0, 2\pi]$). These data representations and data reductions for visualization are seen in Figure 4.10. The diffusion ODF will play a central part of the thesis and is at the heart of Q-Ball Imaging (QBI).

In the second case, a discrete uniform sampling of the sphere is done for a certain radius in q-space (given by the b -value). The signal attenuation is thus measured on a single shell of q-space. In this chapter, we refer to this acquisition as single shell HARDI [Tuch et al. (1999)].³ The idea is that the radial information of the diffusion

²There is very recent development in multiple-shell acquisition schemes. See [Khachaturian et al. (2007)].

³HARDI acquisitions as described by D. Tuch in [Tuch (2002)] are generally used to mean acquisi-




Figure 4.9: Expected diffusion PDF in isotropic, single fiber and two crossing fibers from images taken from [Hagmann et al. (2006b)] and the brain museum coronal slice of the previous chapter (www.brainmuseum.org/Specimens).

PDF can be discarded if one is interested in fiber directions. Thus, most single shell HARDI techniques aim at reconstructing the diffusion ODF or variants of this function in order to have a function whose maxima are aligned with the underlying fiber structure. More than 60 measurements are desirable and medium gradient strengths are acceptable although strong gradients give better diffusion ODF reconstructions. Typically, $60 \leq N \leq 200$, $b \geq 3000 \text{ s/mm}^2$ is used and acquisition time is between 10 and 20 minutes.

Summary Discussion

We have just described a brief (not exhaustive) overview of the history of diffusion NMR and diffusion MRI: from the physics of molecular diffusion, the PGSE sequence,

tions on a single shell in q-space to simplify sampling uniformly the whole 3D q-space. In this thesis, HARDI is more general and refers to any high angular sampling of q-space for many q-magnitudes. We prefer to specifically say single-shell and multiple-shell HARDI for acquisitions done on a single and multiple spherical shells.




Figure 4.10: Simplifying visualization of the 3D diffusion PDF or 3D displacement distribution by either taking an isosurface of the distribution or computing the diffusion ODF. Image taken from [Hagmann et al. (2006b)].

to diffusion MRI. Overall, we saw the evolution of diffusion MRI from simple scalar DWI, to tensor DWI or DTI and beyond to HARDI techniques able to recover the 3D diffusion PDF and diffusion ODF of water molecules in biological tissues. This brief history of diffusion MRI is summarized in Figure 4.2. We also summarized the principal diffusion MRI acquisition techniques and enumerated some of their technical requirements in Table 4.1.

4.3 MULTIPLE FIBER HARDI RECONSTRUCTION TECHNIQUES

The goal of HARDI is to capture multiple fiber directions within the same imaging voxel. HARDI acquisitions are currently being improved every day with better material and better reconstruction algorithms. This thesis is on the latter - improving software development for processing HARDI data. We will now review the existing multiple fiber reconstruction algorithms. Some HARDI reconstruction methods are model dependent, some model-free, some have linear solutions whereas others require non-linear optimization schemes. A schematic view of the diffusion MRI techniques and major multiple fiber HARDI reconstruction algorithms is shown in Figure 4.11 and Table 4.2. We now describe each method in turn.

dMRI modality	Gradient strength (sec / mm ²)			# of measurements N	Acquisition time (min)*	
	weak	medium	strong		fast	slow
DWI	✓ ≤ 1000			$N = 1$	✓ 1-3	
Trace & ADC	✓ ≤ 1000			$2 \leq N \leq 4$	✓ 2-4	
DTI	✓ ≤ 1000			$N \geq 7$	✓ 3-6	
Single shell HARDI		✓	✓	$N \geq 60$	✓	✓ 10-20
DSI			✓ $b > 8000$	$N \geq 200$ (more the better)		✓ 15-60

* Assuming 30 axial slice with thickness of approximately 3 mm each.

Table 4.1: Diffusion MRI acquisition techniques. Acquisition times are assuming thickness of 3 mm as reported in [Hagmann et al. (2006b)]. However, note that is too coarse for most application and typically, a slice thickness of 2 mm is chosen. In that case, one has to expect longer acquisition times. We also assume that a non-diffusion weighted measurement must be done for Trace and ADC imaging and for DTI.

4.3.1 Diffusion Spectrum Imaging (DSI)

Eq. 4.7 suggests the fairly straightforward means of extracting the diffusion PDF from measurements in q-space by measuring the signal on a Cartesian grid of points in q-space and then taking the 3D inverse Fourier transform to obtain an approximated PDF. This technique is called q-space imaging (QSI) [Callaghan et al. (1988); Callaghan (1991)] or diffusion spectrum imaging (DSI) [Wedeen et al. (2000)]. QSI was originally done only in the radial dimension taking 1D inverse Fourier transforms [Callaghan et al. (1988); King et al. (1994, 1997); Assaf et al. (2000); Cohen and Assaf (2002)]. Then, QSI was done with a full 3D Fourier inversion with high radial and high angular measurements of the diffusion PDF. At that point, the technique was named DSI [Wedeen et al. (2000)]. In the context of diffusion MRI, full human brain DSI acquisitions have been developed and used with some success (see [Tuch (2002); Lin et al. (2003)]).

However, DSI is restricted by severe technical limitations. First, in order to resolve features in the PDF of the order of some scale $1/a$, it requires a box of side length $> a$ in q-space. In practice, this requires many measurements and very large b -values compared to those used on conventional scanners. As we want δ small to satisfy the narrow pulse approximation, the gradients G must be very high which creates eddy current distortions, and even can induce harmful electric fields in the subject. This

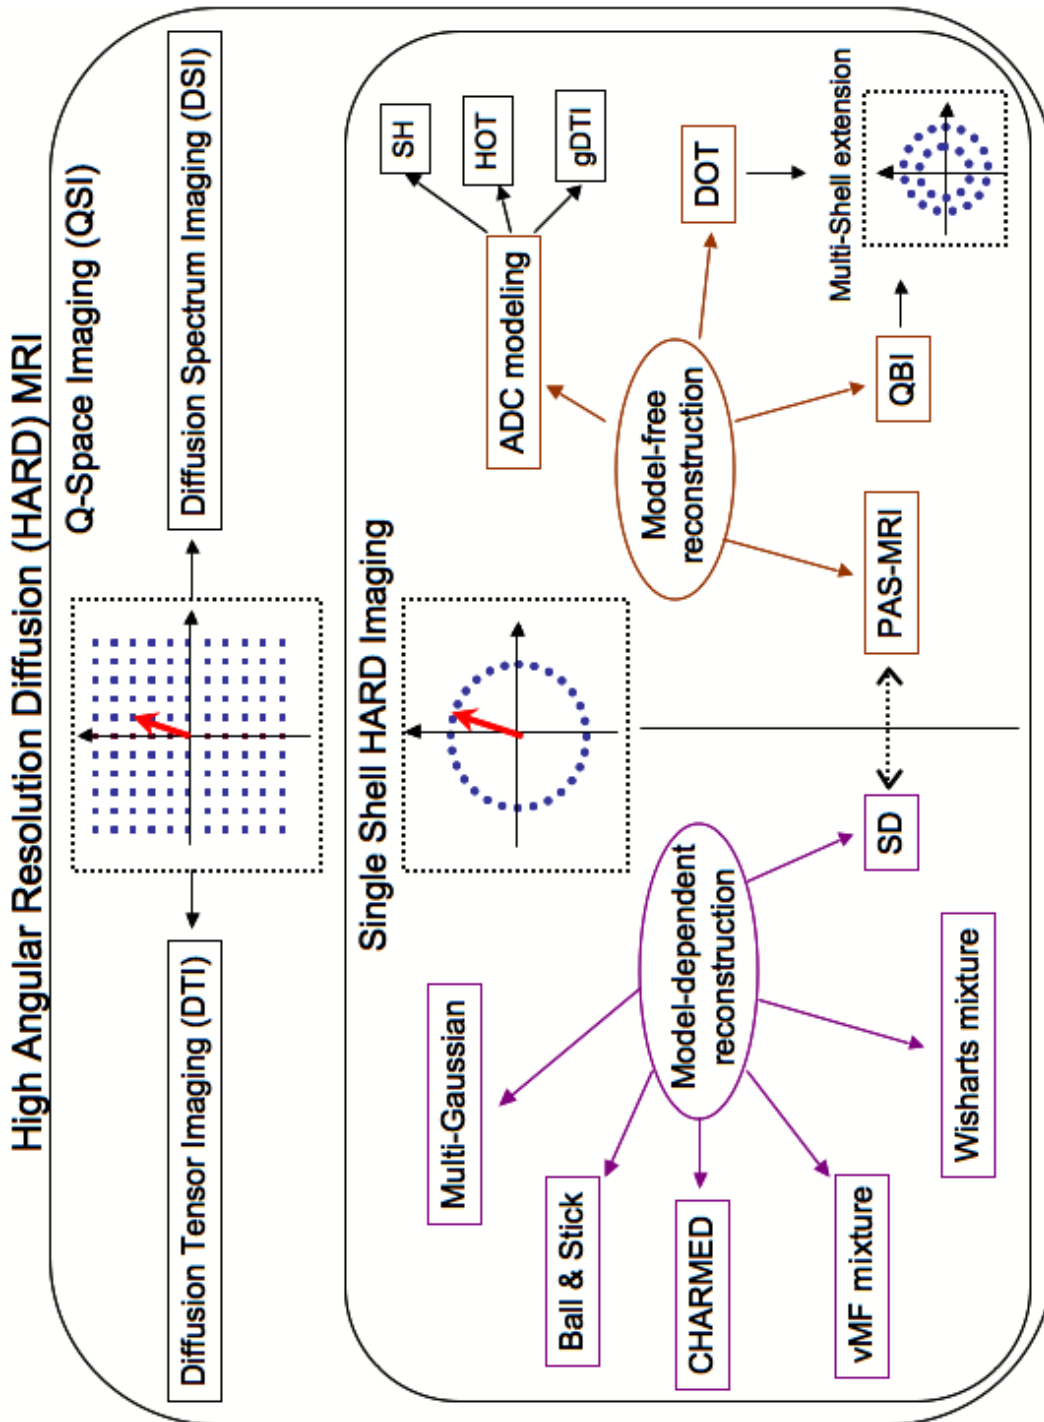


Figure 4.11: Sketch of diffusion MRI acquisition and reconstruction methods.

results not only in an engineering limitation as to the maximal b -values attainable, but also in a SNR problem. The latter arises from the fact that the measured quantity is signal attenuation. This results in a major practical problem because of the number of samples required to tightly fill in a large 3D Cartesian grid. This problem is mainly one of imaging time, which increases like N^3 as the size of the grid is




Figure 4.12: Discrete samplings of the sphere for different numbers of gradient directions corresponding to order 3rd and order 4th tessellation of the sphere respectively.

increased, and makes it less practicable to obtain very high resolution images using the q-space method.

As a result of the DSI limitations, other techniques have been developed to attempt to extract the desired diffusion PDF and diffusion ODF information in a more efficient way. One such clinically feasible approach is single shell (HARDI) [Tuch et al. (1999, 2002); Tuch (2002)].

4.3.2 Single Shell HARD Imaging

The idea of single-shell HARDI is to sample a single sphere with N discrete gradient directions as illustrated in Figure 4.12. Then, the signal attenuation is measured along each of the N directions. In practice, one can sample only the hemisphere as the diffusion PDF is assumed to be symmetric. Given these discrete measurements over the surface of the sphere, several methods have been proposed to extract estimates of the ADC, the diffusion ODF and the diffusion PDF. The values of b required in single shell HARDI are relatively small ($b = 4000 \text{ s/mm}^2$ used in [Tuch (2002)]) compared to the maximal values of b required for the use the DSI technique ($b = 20000 \text{ s/mm}^2$ used in [Wedeen et al. (2000)] and $b = 60000 \text{ s/mm}^2$ used in [Ozarslan and Mareci (2003a)]). Therefore, the advantage of single shell HARDI is that samples are only taken on a single sphere in q-space and thus, the imaging time is much smaller than that of the DSI despite high angular resolution measurements. Moreover, due to a relatively smaller b -value, the SNR is greatly improved. We now perform a survey of the major existing high order reconstruction methods from single shell HARDI.

Apparent Diffusion Coefficient (ADC) Modeling

The first attempts to characterize multiple fiber configurations were done by modeling the apparent diffusion coefficient [Frank (2002); Alexander et al. (2002)] without assuming a Gaussian diffusion. From HARDI signal measurements, we can solve for the ADC using Eq. 4.10. That is, from N samples of the signal on the sphere, one computes the function $\log(S(b)/S_0)$. As seen in Figure 4.13, even though the




Figure 4.13: The ADC and ODF in a two crossing fibers configuration. The ADC and diffusion ODF have multiple maxima, which reveals the intravoxel multiple fiber population. The maxima of the ADC are not aligned with the underlying fiber distribution whereas for the ODF, the maxima are aligned.

ADC reveals information about intravoxel of multiple directions, the maxima of the ADC do not agree with the underlying fiber population. This was shown in phantoms [von dem Hagen and Henkelman (2002)] and in vivo [Tuch et al. (2002)] as well as analytically with a cylindrical diffusion model [von dem Hagen and Henkelman (2002)]. This is due to the nature of the ADC measurement which is the projection of spin displacements onto the diffusing gradient axis. Hence, in the presence of two perpendicular fiber fibers, the greatest diffusion coefficient and hence signal drop occurs, surprisingly, at an angle midway between the fibers and *not* in the direction of either of the fibers, illustrated in [von dem Hagen and Henkelman (2002)]. Hence, the ADC cannot be used directly for fiber tractography. Nonetheless, the ADC can be useful to identify the presence of multiple fiber compartments without actually computing the orientations themselves. The ADC profile can be modeled with spherical harmonics (SH), high order tensors (HOT) and similar generalized DTI (gDTI) approaches [Frank (2002); Alexander et al. (2002); Ozarslan and Mareci (2003a); Zhan et al. (2003); Hirsch et al. (2003); Liu et al. (2004); Chen et al. (2004b)]. Estimation techniques as well as high order anisotropy measures that can be extracted from ADC modeling will be seen in great details in Chapter 6.

Mixture Models

Multi-Gaussian Modeling It is a simple extension of the DTI model to assume that a mixture of Gaussians can describe the diffusion PDF, P . [Tuch et al. (2002)] proposed to model the HARDI signal as a finite mixture of n Gaussians to obtain the diffusion PDF as

$$P(\mathbf{r}) = \sum_{i=1}^n \left(\frac{1}{(4\pi\tau)^3 |\mathbf{D}_i|} \right)^{1/2} \exp \left(\frac{-\mathbf{r}^T \mathbf{D}_i^{-1} \mathbf{r}}{4\tau} \right). \quad (4.14)$$

In [Tuch et al. (2002)], an iterative gradient descent is used for the reconstruction of the PDF which is unstable and sensitive to the starting point. The multiple Gaussian modeling also depends on a possibly over-simplified and incorrect fiber model when fixing the number of fiber compartments. However, [Blyth et al. (2003)] uses the voxel-classification of [Alexander et al. (2002)] to automatically select the number of compartments at every voxel. Typically, people manually set $n = 2$. However, it is reported [Tuch (2002); Alexander (2005b)] that the solution can be quite unstable and considerably affected by the number of measurements and SNR. To make the numerical solution more stable, one can make assumptions about the estimated DT and enforce constraints on the multiple compartment fitting problem. In fact, one can impose symmetry of eigenvalues, force certain magnitude and ratios of eigenvalues or impose positive definiteness of the DT [Alexander et al. (2001); Tuch et al. (2002); Tuch (2002); Blyth et al. (2003); Chen et al. (2004a); Maier et al. (2004); Peleda et al. (2006)]. Finally, a very recent multi-Gaussian extension uses a diffusion basis function of Gaussians [Ramirez-Manzanares et al. (2007)] to recover multiple crossing fibers.

Note that multi-Gaussian modeling is also called multi-tensor modeling. In this thesis, we choose this technique to generate synthetic datasets because it has an analytical PDF and an analytical ODF expression, which facilitates validation against a ground truth. Appendix A.1 describes our synthetic multi-tensor simulations.

Ball & Stick Model A similar approach to multi-Gaussian modeling is the *ball & stick* mixture model. It assumes that water molecules in an imaging voxel belong to one of two populations, a restricted population within or near fiber structures and a free population that is not affected by fiber structure barriers. An anisotropic Gaussian distribution P_r is used for the restricted population and an isotropic Gaussian distribution P_f is used for the free population. The diffusion PDF is then given by $P = \alpha P_f + (1 - \alpha)P_r$, where α represents the volume fraction of the free population. The approach extends to a mixture of restricted compartments and is thus able to recover multiple fiber compartments [Hosey et al. (2005); Behrens et al. (2007)]. However, note that this method has room for improvement with physically more meaningful compartment models. Restricted diffusion does not behave like a stick and the extra-axonal compartment does not behave like a ball.

Composite and hindered restricted model of diffusion (CHARMED) A similar approach to the multi-Gaussian and ball & sticks models, with slightly more complex compartment modeling, is the technique proposed in [Assaf et al. (2004); Assaf and Basser (2005)]. This technique assumes a highly restricted compartment that is non-Gaussian and a hindered compartment that is approximately Gaussian. Hence, the solution combines a QSI non-Gaussian modeling and a DTI Gaussian modeling. The restricted compartment uses a Neuman’s model for restricted dif-

fusion in a cylinder [Neuman (1974)]. The approach can also be formulated as a mixture of restricted compartments and is thus able to recover multiple fiber compartments.

The multi-Gaussian, Ball & Stick and CHARMED models all suffer from the same shortcomings regarding model selection and numerical implementation. One must select the number of compartments a priori, one must use non-linear optimization to solve for the parameters and the methods are sensitive to noise and to the number of measurements.

Spherical Deconvolution (SD)

SD methods generalize the mixture modeling methods of the previous section by assuming a distribution of fiber orientations rather than a discrete number of them. Hence, the limitation of the number of compartment selection n is overcome in the formulation of the problem. The original SD method was proposed in [Tournier et al. (2004)]. The idea is to view the HARDI signal as the convolution of the response function produced by a single fiber with the expected true fiber distribution. The original SD method is linear and uses spherical harmonics and rotational harmonics to parameterize the signal and fiber response function respectively. Spherical deconvolution methods are able to recover fibers crossing but suffer from severe instabilities for high harmonic orders [Tournier et al. (2004); Alexander (2005b)]. The original SD method of [Tournier et al. (2004)] is improved in [Alexander (2005b)] using a non-linear maximum entropy implementation of the spherical deconvolution. Other linear and non-linear SD methods have quickly appeared in the literature to better deal with the SD instabilities, noise and negative diffusivities appearing in the deconvolution process [Anderson (2005); Ramirez-Manzanares et al. (2007); Dell'Acqua et al. (2007); Kaden et al. (2007); Jian and Vemuri (2007b)]. Spherical deconvolution methods will be reviewed and some implemented in great details in Chapter 9.

Just as in mixture models, a disadvantage of deconvolution methods is that a fiber response function needs to be assumed a priori. In [Alexander (2005b)], the response function is a standard Gaussian function whereas in [Tournier et al. (2004); Anderson (2005)], the response function is chosen to be a symmetric diffusion tensor estimated from real datasets.

Mixture of Wishart Distributions

In [Jian and Vemuri (2007a); Jian et al. (2007)] the case of multiple fiber bundles is handled in a similar way to the spherical deconvolution methods. The novelty is that each fiber bundle is represented by a Wishart distribution. Hence, the signal is thought to be given by a Laplace transform defined on the manifold of symmetric




Figure 4.14: Funk-Radon Transform $\mathcal{G}[S]$ illustrated for the HARDI signal S with 1 fiber (left) and two orthogonal fibers (right). The thin lines are the true underlying fiber orientations and the thicker tubes are the detected maxima. These functions are spherical functions and for visualization purposes, the radius of the respective spheres are scaled by the corresponding value on the surface.

positive definite tensors. This leads to a reformulation of DTI in the presence of a single orientation but is also able to account for multiple fibers crossing. The solution is obtained with a non-linear Levenberg-Marquardt optimization technique.

q-Ball Imaging (QBI)

q-ball imaging [Tuch (2002, 2004)] has the advantage of being model-independent. QBI showed that it was possible to reconstruct a smoothed version the diffusion ODF directly from single shell HARDI acquisition with the Funk-Radon transform (FRT). Intuitively, the FRT value at a given spherical point is the great circle integral of the signal on the sphere defined by the plane through the origin perpendicular to the point of evaluation. The FRT is qualitatively illustrated in figure 4.14.

The ODF is intuitive because it has its maximum(a) aligned with the underlying population of fiber(s). Hence, it is a more interesting function for tractography than the ADC. The original QBI has a numerical solution [Tuch (2004)] and more recent methods [Anderson (2005); Hess et al. (2006); Descoteaux et al. (2007a)] have introduced an analytical spherical harmonic reconstruction solution that is faster and more robust to noise. QBI is at the heart of this thesis and will come back often. Review of the existing methods and discussion will be given in great details in Chapter 7.

Figure 4.15 shows the q-ball diffusion ODF reconstruction in the same mid-axial slice as Figure 4.7. Where the DT is planar and greenish in Figure 4.7, the diffusion ODF has multiple maxima in Figure 4.15.

von Mises-Fisher (vMF) Mixture Model of the Diffusion ODF
 In [McGraw et al. (2006a)], the diffusion ODF is modeled with a mixture of




Figure 4.15: q-ball diffusion ODF reconstruction in the same mid-axial slice as Figure 4.7. The ODFs are visualized only where $FA > 0.1$. When the DT is planar with greenish color in Figure 4.7, the diffusion ODF has multiple maxima here.

Von Mises-Fisher distributions. As with a spherical harmonic basis, this ODF representation allows to represent the diffusion ODF with a few parameters instead of having to deal with all spherical values. Although one must solve for the Von Mises-Fisher distribution parameters with non-linear methods, the formulation allows the definition of closed-form Riemannian distances between diffusion ODFs, which is useful for interpolation and segmentation applications [McGraw et al. (2006b)].

Multiple Shell HARD Imaging It is interesting to point out a very recent work that shows that the QBI can be extended to multiple shell HARDI data. In [Khachaturian et al. (2007)], a two-shell acquisition is used to improve angular resolution of QBI. The method combines a low b -value and a high b -value acquisition - one at low b -value $b = 700 \text{ s/mm}^2$ with 70 gradient directions and the other with a higher b -value $b = 3200 \text{ s/mm}^2$ and 262 directions. The reconstruction combines the benefits of the high SNR of the low b -value shell and high angular contrast

and peak separation of the high b -value shell. With hardware improvements of current scanners, these multiple shell HARDI acquisitions might become more popular. Multi-shell HARDI is a trade-off between a DSI acquisition and a single shell QBI acquisition.

Persistent Angular Structure MRI (PAS-MRI)

Another model-independent method to estimate an ODF was proposed in [Jansons and Alexander (2003)]. The method reconstructs the radially persistent angular structure (PAS), \hat{p} , of the diffusion PDF. Assuming independence of the angular and radial structure of the diffusion PDF, the problem is formulated as

$$P(\mathbf{r}) = \frac{1}{r_0^2} \hat{p}(\hat{\mathbf{r}}) \delta(|\mathbf{r}| - r_0). \quad (4.15)$$

This formulation forces probabilities to be non-zero only on a spherical shell of radius r_0 , which can be a questionable assumption. Plugging Eq. 4.15 in Eq. 4.7 and using a maximum entropy cost function, the PAS function $\hat{p}(\hat{\mathbf{r}})$ can be reconstructed by fitting the raw data with a non-linear iterative algorithm. Hence, the PAS is the function on the sphere that best describes the signal measurements and that best represents the relative mobility of spins in each direction. Although it is not obvious why the peaks of the PAS should agree with the underlying fiber distribution, the method is accurate [Alexander (2005b)] and produces ODFs that look sharper than the q-ball estimation of the diffusion ODF. However, the reconstruction is extremely heavy computationally and can take several hours/days to reconstruct a full brain dataset of PAS functions. Recent efforts [Seunarine and Alexander (2006)] have been done to propose a linearized solution to PAS-MRI. This is based on the fact that PAS-MRI is a special case of spherical deconvolution methods [Alexander (2005b)]. This relation is indicated by a dotted arrow connection between SD and PAS-MRI in Figure 4.11.

Diffusion Orientation Transform (DOT)

The diffusion orientation transform (DOT) proposed by [Ozarslan et al. (2006)] is yet another model-independent reconstruction algorithms. The DOT is a function that maps the ADC profile to the diffusion PDF. From this diffusion PDF, the ODF can be obtained by integrating the radial component of the PDF (Eq. 4.13). The key idea is to note that the Fourier transform can be done using the Rayleigh expansion of a plane wave in spherical coordinates. This was also used in [Ozarslan et al. (2004a)] to fit high-order tensors to the HARDI measurements. The DOT has the advantage of being based on the exact PDF of Eq. 4.7 whereas the QBI solution is a smoothed version of the diffusion ODF using the Funk-Radon transform approximation. However, [Ozarslan et al. (2006)] use a {mono,bi,tri}-exponential decay assumption of the signal whereas in QBI, the method does not assume anything about the signal attenuation. Finally, similar to QBI, it is important to point out that the DOT has a

Reconstruction methods	Model		Solution is	
	dependent	independent	linear	non-linear
DTI	✓		✓	✓
DSI		✓	✓	✓
QBI		✓	✓	
DOT		✓	✓	
PAS		✓	✓	✓
SD	✓		✓	✓
Wisharts Mixture	✓			✓
vMF Mixture	✓			✓
CHARMED	✓			✓
Ball & Sticks	✓			✓
Multi-Gaussian	✓			✓

Table 4.2: Diffusion MRI acquisition techniques.

possible multiple shell HARDI extension with the bi and tri exponential fit suggestion [Ozarslan et al. (2006)]. However, as of today, the bi- and tri-exponential decay modeling has only been done on synthetic numerical phantoms.

4.4 CONCLUSION

We have presented a brief history of diffusion MRI: from the physics of molecular diffusion, the PGSE sequence, to diffusion MRI (Figure 4.2). We have also presented the state-of-the-art multiple fiber reconstruction algorithms. The methods were grouped into two classes: model independent and model dependent (Figure 4.11 and Table 4.2) reconstruction methods. Overall, we saw the evolution of diffusion MRI from simple scalar DWI, to tensor DWI or DTI and beyond to HARDI techniques able to recover the 3D diffusion PDF and diffusion ODF of water molecules in biological tissues. In this thesis, we focus on techniques beyond the diffusion tensor imaging because they are able to recover complex multiple fiber distributions. We will be most interested in ADC modeling, QBI and spherical deconvolution from HARDI measurements in the next chapters. One has to keep in mind that at the start of this thesis, DTI was very well studied by C. Lenglet [Lenglet (2006)] in the Odyssee Project Team and HARDI reconstruction algorithms were at their beginning. It was thus important to first understand the HARDI data and the crossing fiber problem, then, find the right mathematical tools to process HARDI data and finally propose new methods to overcome the limitations of DTI. In this thesis, we use spherical harmonics as our mathematical tool and develop linear and fast reconstruction methods for ADC modeling, for analytical QBI and for spherical deconvolution reconstruction to obtain information on the underlying complex fiber distribution.

MATHEMATICS ON THE SPHERE: THE SPHERICAL HARMONICS

“The differential equations of the propagation of heat express the most general conditions, and reduce the physical questions to problems of pure analysis, and this is the proper object of theory.”

–Joseph Fourier

Contents

5.1 Solving the Laplace’s Equation	62
5.2 Spherical Harmonics	64
5.3 The Modified Spherical Harmonics Basis	65
5.4 Spherical Function Estimation with Spherical Harmonics . .	66
5.5 The Funk-Hecke Theorem	68
5.6 Properties of Spherical Harmonics Used in this Thesis	69

OVERVIEW

What is the appropriate mathematical tool to process HARDI data? This chapter covers the mathematical concepts needed to understand the thesis. In particular, what is the appropriate mathematical tool to represent discrete data that lives on the sphere? The SH basis is chosen to be the appropriate tool to represent discrete spherical functions in this thesis. The spherical harmonics are first defined as the solution to the Laplace's equation in spherical coordinates. Then, important properties of the SH basis are described and we see how the SH basis can be used to estimate spherical functions. Finally, the Funk-Hecke theorem is stated. It is an important theorem to solve integrals on the sphere.

Keywords: Laplace's equation, spherical harmonics, Funk-Hecke theorem

Organization of this chapter:

The chapter is organized as follows. We first solve the Laplace's equation to see how the spherical harmonics are defined in Section 5.1 and then review some powerful properties of the spherical harmonic basis in Section 5.2. We then describe how any spherical function can be represented with a series of spherical harmonics in Section 5.3. Next, we define and give the intuition behind the Funk-Hecke theorem in Section 5.5. Finally, Section 5.6 concludes the chapter by enumerating the important mathematical properties that will be used in this thesis.

5.1 SOLVING THE LAPLACE'S EQUATION

According to [Evans (1998)] the "*Laplace's equation is undoubtedly among the most important partial differential equation (PDE)*". It comes up in a wide variety of physical problems such as Fick's law of diffusion, Fourier's law of heat conduction and Ohm's law of electrical conduction. Here, we consider the Laplace's equation in 3-dimensions (3D). The problem is to find twice-differentiable complex-valued functions, ψ , of variables x, y, z such that

$$\frac{\partial^2 \psi}{\partial x^2} + \frac{\partial^2 \psi}{\partial y^2} + \frac{\partial^2 \psi}{\partial z^2} = 0. \quad (5.1)$$

This is the steady state form of the diffusion equation in 3D and is often written as

$$\Delta^2 \psi = 0, \quad (5.2)$$

where Δ is the Laplacian operator.

DEFINITION .1. A C^3 function (3D complex function) satisfying Eq. 5.2 is called a harmonic function.

In this thesis, we are dealing with spherical functions arising from discrete samples on the sphere, which come from HARDI acquisitions. Hence, it is natural to

seek for spherical harmonics, i.e., complex functions satisfying the Laplace's equation in spherical coordinates (r, θ, ϕ) . For the rest of the thesis, we suppose that (θ, ϕ) obey physics convention, where $\theta \in [0, \pi]$, $\phi \in [0, 2\pi)$. Laplace's equation in spherical coordinates is given by

$$\frac{1}{r^2} \frac{\partial}{\partial r} \left(r^2 \frac{\partial \psi}{\partial r} \right) + \frac{1}{r^2 \sin^2 \theta} \frac{\partial^2 \psi}{\partial \phi^2} + \frac{1}{r^2 \sin \theta} \frac{\partial}{\partial \theta} \left(\sin \theta \frac{\partial \psi}{\partial \theta} \right) = 0. \quad (5.3)$$

To solve it, we use separation of variables. We suppose the solution will have the form

$$\psi(r, \theta, \phi) = R(r)\Phi(\phi)\Theta(\theta). \quad (5.4)$$

Inserting this in Eq. 5.3 and multiplying by $r^2 \sin^2 \theta / (R\Phi\Theta)$, we have

$$\frac{\sin^2 \theta}{R} \frac{d}{dr} \left(r^2 \frac{dR}{dr} \right) + \frac{1}{\Phi} \frac{d^2 \Phi}{d\phi^2} + \frac{\sin \theta}{\Theta} \frac{d}{d\theta} \left(\sin \theta \frac{d\Theta}{d\theta} \right) = 0. \quad (5.5)$$

Separating the ϕ -dependent part first and using separation constant $-m^2$, we have

$$\frac{1}{\Phi} \frac{d^2 \Phi}{d\phi^2} = -m^2. \quad (5.6)$$

Although not necessary, the separation constant is chosen in anticipation of a periodical dependency on the azimuthal coordinate ϕ . It will make the solution easier thereafter. Now, it is easy to see that the solution to Eq. 5.6 is

$$\Phi(\phi) = A_m e^{im\phi} + B_m e^{-im\phi}, \quad (5.7)$$

where m is restricted to be an integer by the necessary periodical condition that $\Phi(\phi + 2\pi) = \Phi(\phi)$. Now, substituting Eq. 5.7 in Eq. 5.5 and dividing by $\sin^2 \theta$, we obtain

$$\frac{1}{R} \frac{d}{dr} \left(r^2 \frac{dR}{dr} \right) - \frac{m^2}{\sin^2 \theta} + \frac{1}{\Theta \sin \theta} \frac{d}{d\theta} \left(\sin \theta \frac{d\Theta}{d\theta} \right) = 0. \quad (5.8)$$

In our problem, because of the nature of the single shell HARDI acquisitions, we can suppose that all spherical values live on a sphere with constant radius. Hence, we assume that the radial part is equal to the constant $\ell(\ell + 1)$. The reason for choosing separation constant as $\ell(\ell + 1)$ is because we are anticipating the appearance of the associated Legendre differential equation [(Arfken and Weber, 1995, Ch.12)]. Hence, we have

$$\ell(\ell + 1) - \frac{m^2}{\sin^2 \theta} + \frac{1}{\Theta \sin \theta} \frac{d}{d\theta} \left(\sin \theta \frac{d\Theta}{d\theta} \right) = 0. \quad (5.9)$$

Expanding the θ -dependent part and multiplying by Θ , we obtain

$$\ell(\ell + 1) - \frac{m^2}{\sin^2 \theta} + \frac{\cos \theta}{\sin \theta} \frac{1}{\Theta} \frac{d\Theta}{d\theta} + \frac{1}{\Theta} \frac{d^2 \Theta}{d\theta^2} = 0 \quad (5.10)$$

$$\implies \frac{d^2 \Theta}{d\theta^2} + \frac{\cos \theta}{\sin \theta} \frac{d\Theta}{d\theta} + \left[\ell(\ell + 1) - \frac{m^2}{\sin^2 \theta} \right] \Theta = 0. \quad (5.11)$$

This last equation is the standard associated Legendre differential equation for $x = \cos \theta$, $\ell \geq 0$ and $|m| \leq \ell$ with solution

$$P_\ell^m(\cos \theta). \quad (5.12)$$

Therefore, with the ϕ -dependent (Eq. 5.7) solution and the θ -dependent (Eq. 5.12) solution, we can obtain the general complex solution for the constant radius Laplace's equation in spherical coordinates as

$$\psi(\theta, \phi) = \sum_{\ell=0}^{\infty} \sum_{m=-\ell}^{\ell} c_\ell^m P_\ell^m(\cos \theta) e^{im\phi} = \sum_{\ell=0}^{\infty} \sum_{m=-\ell}^{\ell} c_\ell^m Y_\ell^m(\theta, \phi), \quad (5.13)$$

where $Y_\ell^m(\theta, \phi) = P_\ell^m(\cos \theta) e^{im\phi}$ are defined as the *spherical harmonics* (SH).

Hence, spherical harmonics Y_ℓ^m satisfy the angular part of the Laplace's equation in spherical coordinates. This PDE is also called the spherical harmonic differential equation (SHDE) and is given by

$$\frac{1}{\sin \theta} \frac{\partial}{\partial \theta} \left(\sin \theta \frac{\partial Y_\ell^m}{\partial \theta} \right) + \frac{1}{\sin^2 \theta} \frac{\partial^2 Y_\ell^m}{\partial \phi^2} + \ell(\ell + 1) Y_\ell^m = 0. \quad (5.14)$$

The first two terms of this equation correspond to the Laplacian in spherical coordinates, also called the 3D Laplace-Beltrami operator Δ_b .

DEFINITION .2. *The 3D Laplace-Beltrami operator Δ_b is defined as*

$$\Delta_b = \frac{1}{\sin \theta} \frac{\partial}{\partial \theta} \left(\sin \theta \frac{\partial}{\partial \theta} \right) + \frac{1}{\sin^2 \theta} \frac{\partial^2}{\partial \phi^2}. \quad (5.15)$$

Thus, according to Eq. 5.14, the spherical harmonics satisfy the relation

$$\Delta_b Y_\ell^m = -\ell(\ell + 1) Y_\ell^m. \quad (5.16)$$

This relation will be useful in the rest of the thesis for regularization purposes.

5.2 SPHERICAL HARMONICS

DEFINITION .3. *Spherical harmonics Y_ℓ^m of order ℓ and degree m are the angular portion of Laplace's equation in spherical coordinates. They are defined as*

$$Y_\ell^m(\theta, \phi) = \sqrt{\frac{2\ell + 1}{4\pi} \frac{(\ell - m)!}{(\ell + m)!}} P_\ell^m(\cos \theta) e^{im\phi} \quad (5.17)$$

The normalization factor in Eq. 5.17 is chosen so that the spherical harmonics form an orthonormal set of functions with respect to the inner product

$$\langle Y_\ell^m(\theta, \phi), Y_{\ell'}^{m'}(\theta, \phi) \rangle = \int_{\Omega} Y_\ell^m(\theta, \phi) \bar{Y}_{\ell'}^{m'}(\theta, \phi) d\Omega \quad (5.18)$$




Figure 5.1: Real part squared $\text{Re}[Y_\ell^m(\theta, \phi)]^2$ of the spherical harmonics basis up to order 3. The color map is going from red to blue for maximal to minimal values on the sphere respectively. The 0-th order harmonic is blue because it is constant on the sphere.

where integration over Ω denotes integration over the unit sphere ($d\Omega = \sin \theta d\theta d\phi$), and $\bar{Y}_{\ell'}^{m'}$ denotes the complex conjugate of $Y_{\ell'}^{m'}$. That is,

$$\langle Y_\ell^m(\theta, \phi), Y_{\ell'}^{m'}(\theta, \phi) \rangle = \delta_{\ell\ell'} \delta_{mm'}, \quad (5.19)$$

where δ_{ij} is the Kronecker delta, i.e.

$$\delta_{ij} = \begin{cases} 1 & i = j \\ 0 & i \neq j \end{cases}$$

Finally, observe that with respect to the transformation $T : (\theta, \phi) \rightarrow (\pi - \theta, \phi + \pi)$, the spherical harmonics have the following very simple behavior,

$$Y_\ell^m(T(\theta, \phi)) = \begin{cases} Y_\ell^m(\theta, \phi), & \text{if } \ell \text{ even} \\ -Y_\ell^m(\theta, \phi), & \text{if } \ell \text{ odd} \end{cases} \quad (5.20)$$

In other words, the even order spherical harmonics are antipodally symmetric, while the odd order spherical harmonics are antipodally anti-symmetric. In order to have an idea of what spherical harmonics look like, we show in Figure 5.1 the real part squared, $\text{Re}[Y_\ell^m(\theta, \phi)]^2$, of the spherical harmonics Y_ℓ^m up to order 3.

5.3 THE MODIFIED SPHERICAL HARMONICS BASIS

The spherical harmonics Y_ℓ^m (Eq. 5.17) are a basis for complex functions on the unit sphere. Hence, any complex function defined on the sphere can be expressed as a series of spherical harmonics. This is very powerful and analogous to the Fourier

transform very often used in image processing. In that case, any image can be decomposed in a Fourier series, namely in a sum of sinusoids and cosines. In this thesis, we have physical diffusion MRI measurements that represent the average attenuation of the diffusion of water molecules. Hence, the HARDI signal is assumed to be *real* and *symmetric*. Therefore, we want to define a modified spherical harmonic basis that is also real and symmetric.

In order to impose the symmetry constraint on the expansion, we consider only spherical harmonics of even degree. As mentioned above, spherical harmonics of odd order are antipodally anti-symmetric, while spherical harmonics of even order are antipodally symmetric (Eq. 5.20). In order to impose the real-valued constraint, we appropriately choose real and imaginary parts of the spherical harmonics depending on the degree m .

For $\ell = 0, 2, 4, \dots, L$ and $m = -\ell, \dots, 0, \dots, \ell$, we define a single index j in terms of ℓ and m such that $j(\ell, m) = (\ell^2 + \ell + 2)/2 + m$. The modified basis then is

$$Y_j = \begin{cases} \sqrt{2} \cdot \text{Re}(Y_\ell^{|m|}), & \text{if } m < 0 \\ Y_\ell^m, & \text{if } m = 0 \\ \sqrt{2} \cdot (-1)^{m+1} \text{Im}(Y_\ell^m), & \text{if } m > 0. \end{cases} \quad (5.21)$$

where $\text{Re}(Y_\ell^m)$ and $\text{Im}(Y_\ell^m)$ represent the real and imaginary parts of Y_ℓ^m respectively. The basis is designed to be symmetric, real and orthonormal because of the normalization factor $\sqrt{2}$. Examples and proofs are given in [Descoteaux et al. (2005a)]. Moreover, note that there are exactly

$$R = (1/2)(L + 1)(L + 2) \quad (5.22)$$

terms in the spherical harmonic series of order L . Figure 5.2 shows the spherical harmonics in the modified basis of order 4, $L = 4$. For the rest of this thesis, $\mathbf{Y} = \{Y_1, \dots, Y_R\}$ will refer to the modified spherical harmonic basis.

5.4 DISCRETE SPHERICAL FUNCTION ESTIMATION WITH THE SPHERICAL HARMONICS

We want to use spherical harmonics to describe the HARDI signal S consisting of N discrete diffusion-weighted measurements on the sphere. As seen in Figures 1.1 and 1.2, this SH estimation step is the backbone of the entire thesis.

Since the spherical harmonics form an orthonormal basis for all functions on the unit sphere, any spherical function can be expressed by an infinite series of spherical harmonics as expressed in Eq. 5.13. High order spherical harmonics correspond to high frequency modes of the unit sphere, and thus a truncated spherical harmonic series can be effectively used to fit relatively smooth functions. Hence, the problem is to find the best coefficients of the modified SH basis (Eq. 5.21) that describe the HARDI signal S at each of the N diffusion-weighted gradient encoding directions i .




Figure 5.2: Modified spherical harmonics up to order $L = 4$. $\ell(j)$ represents the harmonic order for the j th coefficient.

Thus, the truncated smooth estimation of the HARDI signal S can be formulated as

$$S(\theta_i, \phi_i) = \sum_{j=1}^R c_j Y_j(\theta_i, \phi_i) \quad (5.23)$$

where R is the number of terms (Eq. 5.22) in the modified SH basis \mathbf{Y} of order L . Letting \mathbf{X} be the $N \times 1$ vector representing the input HARDI signal, S_i , for every encoding gradient direction, \mathbf{C} the $R \times 1$ vector of SH coefficients c_j and \mathbf{B} is the $N \times R$ matrix constructed with the discrete modified SH basis

$$\mathbf{B} = \begin{pmatrix} Y_1(\theta_1, \phi_1) & Y_2(\theta_1, \phi_1) & \cdots & Y_R(\theta_1, \phi_1) \\ \vdots & \vdots & \ddots & \vdots \\ Y_1(\theta_N, \phi_N) & Y_2(\theta_N, \phi_N) & \cdots & Y_R(\theta_N, \phi_N) \end{pmatrix}, \quad (5.24)$$

we can write the set of equations as an over-determined linear system

$$\mathbf{X} = \mathbf{B}\mathbf{C}. \quad (5.25)$$

Hence, we need to solve for the SH coefficients c_j , where

$$c_j = \int_{\Omega} S(\theta, \phi) Y_j(\theta, \phi) d\Omega. \quad (5.26)$$

Due to orthonormality of the SH basis, the coefficients of the SH series c_j can be calculated by forming the inner product of S with the spherical harmonics,

$$c_j = \langle S(\theta, \phi), Y_j(\theta, \phi) \rangle = \int_0^{2\pi} \int_0^{\pi} S(\theta, \phi) Y_j(\theta, \phi) \sin \theta d\theta d\phi \quad (5.27)$$

This idea was first used to fit the ADC obtained from HARDI data by [Frank (2002)] where S is the discrete sampling of the diffusivities. [Frank (2002)] performs the direct discretization of the integrals. This is a computationally inefficient method to obtain the coefficients; recent work by [Alexander et al. (2002)] uses a least-squares method to solve for the unknowns. The least-squares method was first proposed in the vision community for the parameterization of closed surfaces for 3D shape description by [Brechtbuhler et al. (1995)].

Here, we also use the least-squares solution for the coefficients \mathbf{C} . The least-squares solution looks for the spherical harmonic series that passes nearest to the discrete samplings on the sphere. Hence,

$$\mathbf{X} = \mathbf{BC} + \mathbf{E}, \quad (5.28)$$

where the error vector \mathbf{E} should be small. This system of overdetermined equations is solved with least-square sums over the columns of \mathbf{E} by minimizing $\|\mathbf{X} - \mathbf{BC}\|^2$ yielding

$$\boxed{\mathbf{C} = (\mathbf{B}^T \mathbf{B})^{-1} \mathbf{B}^T \mathbf{X}}. \quad (5.29)$$

The vector \mathbf{C} of spherical harmonic coefficients gives the best-fitting truncated series to the signal. The estimated signal is then simply recovered by evaluating

$$S(\theta, \phi) = \sum_{j=1}^R c_j Y_j(\theta, \phi), \quad (5.30)$$

for any (θ, ϕ) outside the discrete measurements \mathbf{X} or in the discrete case, by simple matrix multiplication $\mathbf{X} = \mathbf{BC}$. Therefore, one can interpolate over any (θ, ϕ) directions. Though this works fairly well so long as noise is small, an important contribution of this thesis will be to propose a more general fitting procedure that takes advantage of the properties of spherical harmonics and the Laplace-Beltrami operator to quantify the smoothness of spherical functions.

5.5 THE FUNK-HECKE THEOREM ---

The Funk-Hecke theorem was first published by Funk in 1916 and by Hecke in 1918 [(Andrews et al., 1999, chap.9)]. Here, we give only its 3D version. The Funk-Hecke formula is a theorem that relates the inner product of any spherical harmonic with the projection on the sphere of any continuous function $f(t)$ defined on the interval $[-1, 1]$. We write the dot product between two vectors $\mathbf{u}, \mathbf{w} \in \mathbb{R}^3$ as $\mathbf{u}^T \mathbf{w}$. The key observation is that any continuous function f on the interval $[-1, 1]$ extends to a continuous function of two variables $g(\mathbf{u}, \mathbf{w})$ on the sphere defined by $g(\mathbf{u}, \mathbf{w}) = f(\mathbf{u}^T \mathbf{w})$. With this formulation,

THEOREM .1. Funk-Hecke Theorem: *Let $f(t)$ be continuous on $[-1, 1]$ and Y_ℓ any spherical harmonic of order ℓ in \mathcal{C}^3 . Then, given a unit vector \mathbf{u}*

$$\int_{|\mathbf{w}|=1} f(\mathbf{u}^T \mathbf{w}) Y_\ell(\mathbf{w}) d\mathbf{w} = \lambda(\ell) H_\ell(\mathbf{u}), \quad (5.31)$$

where

$$\lambda(\ell) = \frac{2\pi}{P_\ell(1)} \int_{-1}^1 P_\ell(t) \lambda(\ell) = 2\pi \int_{-1}^1 P_\ell(t) f(t) dt$$

with P_ℓ the Legendre polynomial of degree ℓ .

The Funk-Hecke theorem will be very useful in this thesis to solve integrals over the sphere.

5.6 PROPERTIES OF SPHERICAL HARMONICS USED IN THIS THESIS

We now summarize the important mathematical properties of spherical harmonics that will come up often in the later chapters.

1. The modified spherical harmonic basis of Eq. 5.21 is **real** and **symmetric**. This is useful to describe the HARDI signal in a physically meaningful way.
2. The spherical harmonic basis and the modified spherical harmonic basis (Eq. 5.19) are **orthonormal** with respect to the inner product of Eq. 5.18. This will be useful to compare spherical functions and to introduce an efficient metric between spherical functions.
3. The spherical harmonics differential equation of Eq. 5.16 defines the **Laplace-Beltrami operator**. This Laplace-Beltrami operator will be useful to regularized and smooth spherical functions by simple operations of the spherical harmonic coefficients.
4. The **Funk-Hecke theorem** will be useful to solve integrals between spherical functions and spherical harmonics. This will be particularly important to develop the analytical QBI solution and solve spherical deconvolution integrals.

Note that this list is not an exhaustive list of existing spherical harmonics properties. There are more interesting properties that can be found in standard mathematical handbooks [Ferrers (1969); Arfken and Weber (1995)]. We listed the properties that will appear many times in this thesis. Note that other groups such as [Frank (2002), Alexander et al. (2002), Zhan et al. (2003), Tournier et al. (2004), Ozarslan et al. (2005a), Anderson (2005), Hess et al. (2006)] have also converged to the spherical harmonics as HARDI processing tools. This is reassuring and the spherical harmonics seem to be a natural way to decompose signals that live on the sphere, just as the Fourier transform. We will now describe how the spherical harmonic basis can be used to model the ADC, to find an analytical solution to QBI and to perform spherical deconvolution operations in a simple, efficient and robust way.

Part III

Methods

APPARENT DIFFUSION COEFFICIENT ESTIMATION AND APPLICATIONS

It would be vain to search for a rule if there were no regularity.

–Ludwig von Mises

Contents

6.1	Introduction	74
6.2	ADC Profile Estimation from HARDI	76
6.2.1	Fitting the ADC with the Spherical Harmonics	76
6.2.2	Fitting the ADC with a High Order Diffusion Tensor (HODT)	77
6.2.3	A Regularization Algorithm for HARDI Signal/ADC Estimation	78
6.2.4	From SH Coefficients to HODT Coefficients	79
6.3	High Order Anisotropy Measures from ADC Profiles	83
6.3.1	Frank and Chen <i>et al</i> Measures	83
6.3.2	Alexander <i>et al</i> Measure	84
6.3.3	Generalized Anisotropy Measure	84
6.4	Results	85
6.4.1	Optimal Regularization Parameter with the L-Curve Method	85
6.4.2	High Order Anisotropy Measures Results	89
6.4.3	Biological Phantom Results	96
6.4.4	Human Brain HARDI Data Results	96
6.5	Discussion	98
6.5.1	Contributions of this chapter:	100
6.6	Appendix A: Independent Elements of the HODT	102
6.7	Appendix B: Spherical Harmonics to the High Order Tensor	103

OVERVIEW

Given a set of discrete measurements on the sphere, how do we estimate this signal robustly? This chapter covers the HARDI estimation using spherical harmonics and regularization. As a consequence, the apparent diffusion coefficient (ADC) can be modeled and reconstructed robustly. First, we develop a closed-form regularization algorithm to estimate the HARDI signal and the ADC. Second, we develop the link between the SH basis and the high order diffusion tensor (HODT) so that it is possible to compare HARDI anisotropy measures. Finally, we see how the HARDI anisotropy measures can be used to highlight voxels containing multiple fiber distributions.

Keywords apparent diffusion coefficient (ADC), spherical harmonics (SH), high order diffusion tensor (HODT), Laplace-Beltrami operator, regularization, anisotropy measures, HARDI

Contributions of this chapter:

- The new Laplace-Beltrami regularization algorithm to estimate the HARDI signal and the ADC more robustly.
- The linear transformation between coefficients of the modified SH basis and the independent elements of the HODT.
- The comparison between the state-of-the-art high order anisotropy measures computed from the SH basis and HODT representations.
- ADC estimation on synthetic data, on a biological phantom and on a human brain dataset and ability to discriminate voxels with isotropic diffusion, single fiber and multiple fiber distributions.

Organization of this chapter:

The chapter is organized as follows. We motivate the ADC modeling problem in Section 6.1. In Section 6.2, we review the existing techniques to estimate the ADC profile from noisy HARDI data and propose a new regularization method that recovers a smoother ADC that is closer to the ADC without noise. In Section 6.3, we review the different high order anisotropy measures and in Section 6.4, we evaluate and validate our algorithm against state-of-the-art algorithms using SH and the HODT. We conclude with a discussion of the results and our contributions in Section 6.5.

6.1 INTRODUCTION

In the presence of multiple fibers, the ADC profile is oblate or planar and there is no unique principal direction, where a multiple maxima ADC profile is expected, as illustrated in Figure 6.1. Hence, clinicians and neurosurgeons are quite




Figure 6.1: ADC profile estimate from DTI fails to recover multiple fiber orientations. The maxima of the ADC profile do not agree with thin green lines corresponding to the true synthetic fiber orientations (small lines).

skeptical of tracking and segmentation methods developed on DTI data due to the well-known limitations of the DT model, as seen in Chapter 4.

In current clinical applications, people instead choose to use simple anisotropy maps computed from the ADC profile [Doug et al. (2004)] to infer white matter connectivity information. These measures are fast and easy to interpret with regions of anisotropy that clearly stand out. Many anisotropy measures exist and the most commonly used are FA (fractional anisotropy) and RA (relative anisotropy) [Basser and Pierpaoli (1996)] but again, these measures are limited in non-Gaussian diffusion areas when computed from DTI data.

In this chapter, we study the estimation of the ADC profile from HARDI data and its ability to describe complex tissue architecture. We want to design the appropriate tools to describe noisy HARDI data and explore scalar anisotropy measures computed from high order formulation. In particular, we address the problem of estimation of the ADC profile with a HODT. One proposed possibility by [Ozarslan and Mareci (2003a)] is to use direct linear regression by least-squares fitting. This can be effective but its robustness to noise is questionable as there does not appear to be any straightforward way to impose a viable smoothness maximizing criterion. We approach the problem with a SH series approximation as done in [Frank (2002); Alexander et al. (2002); Chen et al. (2004b); Zhan et al. (2003, 2004, 2006); Descoteaux et al. (2006a)].

An important contribution of our work is to propose a generalization of the standard least-squares evaluation method to include a regularization criterion. From this result, we compute the linear transformation taking the coefficients of the SH series to the independent elements of the HODT using the relation presented in [Ozarslan and Mareci (2003a)]. Therefore, any technique developed for SH formulation can be quickly and easily applied to the HODT formulation and vice versa. This bridge is very useful for comparison purposes between state-of-the-art anisotropy measures for high order models computed from SH and tensor coefficients. Published results are reproduced accurately and it is also possible to recover voxels with isotropic, single fiber anisotropic and multiple fiber anisotropic diffusion from syn-

thetic data, from a biological phantom and from a real human brain dataset.

6.2 ADC PROFILE ESTIMATION FROM HARDI

We want to estimate the HARDI signal and ADC profile. At each voxel, we have a discrete spherical function with no a priori assumption about the nature of the diffusion process within the voxel. The problem is to recover a smooth ADC close to the true ADC from the measured diffusion MRI noisy signal. Recall that the ADC was defined in Eq. 4.10 as

$$ADC := D(\mathbf{g}) = -\frac{1}{b} \ln \left(\frac{S(b, \mathbf{g})}{S_0} \right), \quad (6.1)$$

where the b -value is the diffusion-weighting factor, S_0 is the T2-weighted signal acquired without any diffusion gradients and $D(\mathbf{g})$ is the ADC along unit gradient direction $\mathbf{g} = (\sin \theta \cos \phi, \sin \theta \sin \phi, \cos \theta)^T$, where (θ, ϕ) obey physics convention ($\theta \in [0, \pi], \phi \in [0, 2\pi]$). We explore two ways to estimate this ADC profile. First, an estimation with the SH basis as described in Section 5.4 and second, with a HODT estimation.

6.2.1 Fitting the ADC with the Spherical Harmonics

We formulate the estimation of the ADC profile at each discrete sampling i as we did for the signal S in Section 5.4 of the last chapter, i.e.

$$D(\mathbf{g}(\theta_i, \phi_i)) = \sum_{j=1}^R c_j Y_j(\theta_i, \phi_i), \quad (6.2)$$

We let \mathbf{X} represent the $N \times 1$ ADC vector obtained from the raw HARDI signal S in Eq. 6.1. We can thus write the equations given above as an overdetermined linear system

$$\mathbf{X} = \mathbf{B}\mathbf{C} + \mathbf{E}, \quad (6.3)$$

where \mathbf{B} is the SH basis of Eq. 5.24, \mathbf{C} is the $R \times 1$ vector of SH coefficients and \mathbf{E} is the error vector as before. The least-squares solution to the above system Eq. 6.3 is given by

$$\mathbf{C} = (\mathbf{B}^T \mathbf{B})^{-1} \mathbf{B}^T \mathbf{X}, \quad (6.4)$$

as seen before in Section 5.4. Note that the condition number of the pseudo-inverse matrix $\mathbf{B}^T \mathbf{B}$ is always small for estimation order ℓ and sampling scheme N used in our simulations and experiments. Note also that this is an unweighted linear least-squares fit and that it can be simply extended to a weighted fit, as in [Basser et al. (1994a)], to account for the expected variation in each diffusion coefficient and distortions introduced by the logarithmic transformation of Eq. 6.1. The estimated ADC profile is thus recovered by evaluating Eq. 6.2 for any (θ, ϕ) outside the discrete measurements or in the discrete case, by simple matrix multiplication $\mathbf{X} = \mathbf{B}\mathbf{C}$.

6.2.2 Fitting the ADC with a High Order Diffusion Tensor (HODT)

In [Ozarslan and Mareci (2003a)], instead of fitting the ADC profile with a rank-2 tensor, the diffusivities are expressed in terms of a HODT. Replacing the classical rank-2 tensor \mathbf{D} by a HODT into the standard Stejskal-Tanner leads to a generalized Stejskal-Tanner equation [Ozarslan and Mareci (2003a)] given by

$$D(\mathbf{g}) = \sum_{i_1=1}^3 \sum_{i_2=1}^3 \dots \sum_{i_\ell=1}^3 T_{i_1 i_2 \dots i_\ell} g_{i_1} g_{i_2} \dots g_{i_\ell} \quad (6.5)$$

where the $T_{i_1 i_2 \dots i_\ell}$'s are the elements of HODT and the g_i 's are the components of the unit vector specifying the direction of the applied diffusion gradient. This model assumes no a priori knowledge of the ADC profile. However, one must assume that a rank- ℓ tensor is sufficient to determine the ADC profile, \mathbf{X} . Given this assumption, it is possible to simplify the problem by noting that the DT must be totally symmetric, which means that the HODT can be fully expressed in terms of its independent elements, whose vector representation is denoted by \mathbf{T}^1 . Hence, one can rewrite the expression for each ADC $D(\mathbf{g}(\theta_i, \phi_i))$ along direction $\mathbf{g}(\theta_i, \phi_i)$ in a more compact form,

$$D(\mathbf{g}(\theta_i, \phi_i)) = \sum_{k=1}^R \mu_k T_k \prod_{p=1}^{\ell} g_{k(p)}(\theta_i, \phi_i), \quad (6.6)$$

where R is the number of independent elements in the HODT (same as the number of modified SH coefficients), T_k is the k th independent element of the HODT vector representation \mathbf{T} , μ_k is the corresponding multiplicity of the element, and $g_{k(p)}$ gives the component of the gradient direction corresponding to the p th index of the k th independent element of the tensor. The complete derivation is in [Ozarslan and Mareci (2003a)] and is sketched in Appendix 6.6. There is also a rank-2 example showing the link between Eq. 6.5 and Eq. 6.6.

[Ozarslan and Mareci (2003a); Ozarslan et al. (2005b)] fit the ADC profile with a HODT (Eq. 6.6) using linear regression with processing routines written in IDL (Research Systems, Inc., Boulder, CO). Here, we have implemented a standard linear regression with least-squares fitting as described in the previous section. Letting \mathbf{X} represent the ADC profile, \mathbf{T} the vector representation of the HODT independent elements and \mathbf{R} the $N \times R$ matrix

$$\mathbf{R} = \begin{pmatrix} \mu_1 \prod_{p=1}^{\ell} g_{1(p)}(\theta_1, \phi_1) & \cdots & \mu_N \prod_{p=1}^{\ell} g_{N(p)}(\theta_1, \phi_1) \\ \vdots & \ddots & \vdots \\ \mu_1 \prod_{p=1}^{\ell} g_{1(p)}(\theta_{n_s}, \phi_{n_s}) & \cdots & \mu_N \prod_{p=1}^{\ell} g_{N(p)}(\theta_{n_s}, \phi_{n_s}) \end{pmatrix}, \quad (6.7)$$

¹For order 2, $\mathbf{T} = [T_{xx} \ T_{xy} \ T_{xz} \ T_{yy} \ T_{yz} \ T_{zz}]$, see Figures 6.2 and 6.3.

we define an error term \mathbf{E}' so that $\mathbf{X} = \mathbf{R}\mathbf{T} + \mathbf{E}'$. We seek to minimize the quantity $\mathbf{E}'^T \mathbf{E}'$ using the same technique described previously. We then obtain the expansion for the HODT

$$\mathbf{T} = (\mathbf{R}^T \mathbf{R})^{-1} \mathbf{R}^T \mathbf{X}. \quad (6.8)$$

The estimated ADC profile is thus recovered by evaluating Eq. 6.6 for any (θ, ϕ) outside the discrete measurements or in the discrete case, by simple matrix multiplication

$$\mathbf{X} = \mathbf{R}\mathbf{T}. \quad (6.9)$$

The method does not have a smoothing or regularization parameter and will therefore be more sensitive to noise.

6.2.3 A Regularization Algorithm for HARDI Signal/ADC Estimation

As seen in Chapter 4, diffusion-weighted images are quite noisy and this is even more the case at high b -values. It is thus important to seek for regularity in the estimation. Hence, instead of performing a simple least-squares minimization as in Eq. 5.29, Eq. 6.4 and Eq. 6.8, we add a regularization to the fitting procedure. Note that this can be done in the estimation step of any spherical function. In this chapter, we show the regularization on the ADC estimation but it can also be done directly on the raw HARDI signal S .

We define a measure of the deviation from smoothness E of a function f defined on the unit sphere as

$$E(f) = \int_{\Omega} (\Delta_b f)^2 d\Omega, \quad (6.10)$$

where Ω denotes integration over the unit sphere and Δ_b is the Laplace-Beltrami operator. The Laplace-Beltrami operator is a natural measure of smoothness for functions defined on the unit sphere and we know that the SH satisfy the relation $\Delta_b Y_\ell^m = -\ell(\ell + 1)Y_\ell^m$ (Eq. 5.16), as described in the SH properties of the previous chapter. Using the orthonormality property of the modified basis (Eq. 5.19), the above functional can be rewritten straightforwardly in terms of the coefficient vector \mathbf{C} as follows:

$$E(f) = \int_{\Omega} \Delta_b \left(\sum_{p=1}^R c_p Y_p \right) \Delta_b \left(\sum_{q=1}^R c_q Y_q \right) d\Omega = \sum_{j=1}^R c_j^2 \ell(j)^2 (\ell(j) + 1)^2 = \mathbf{C}^T \mathbf{L} \mathbf{C}, \quad (6.11)$$

where \mathbf{L} is simply the $R \times R$ matrix with entries $\ell(j)^2 (\ell(j) + 1)^2$ along the diagonal.² Therefore, the quantity we wish to minimize can be expressed in matrix form as

$$M(\mathbf{C}) = (\mathbf{X} - \mathbf{B}\mathbf{C})^T (\mathbf{X} - \mathbf{B}\mathbf{C}) + \lambda \mathbf{C}^T \mathbf{L} \mathbf{C}, \quad (6.12)$$

² $\ell(j)$ is the order associated with the j th element of the SH basis, i.e. for $j = 1, 2, 3, 4, 5, 6, 7, \dots$ $\ell(j) = 0, 2, 2, 2, 2, 2, 4, \dots$

where λ is a variable weighting factor on the regularization term. The coefficient vector minimizing this expression can then be determined just as in the standard unweighted linear least-squares fit ($\lambda = 0$), by setting each of the $\partial M/\partial c_j = 0$, from which we obtain the generalized expression for the desired SH series coefficient vector

$$\boxed{\mathbf{C} = (\mathbf{B}^T \mathbf{B} + \lambda \mathbf{L})^{-1} \mathbf{B}^T \mathbf{X}}. \quad (6.13)$$

The estimated ADC profile is thus recovered by evaluating Eq. 6.2 for any (θ, ϕ) outside the discrete measurements or in the discrete case, matrix multiplication $\mathbf{X} = \mathbf{B}\mathbf{C}$ as before. Note that a variation on the above derivation will hold in other geometries if a proper basis of functions is chosen. Intuitively, this approach penalizes an approximation function for having higher order terms in its modified SH series. Therefore, higher order terms will only be included in the fit if they significantly improve the overall accuracy of the approximation. This eliminates most of the high order terms due to noise while leaving those that are necessary to describe the underlying function.

It is important to point out that there are other possible choice of error function E (Eq. 6.10) to impose the regularization. A recent work by [Sakaie and Lowe (2007)] uses the gradient square of the spherical function, i.e. $(\nabla f)^2$, to regularize the spherical deconvolution reconstruction. In that case, the regularization is also analytical and can be expressed with a matrix with diagonal elements given by $\ell(j)(\ell(j) + 1)$. Hence, our regularization strength $(\ell(j)^2(\ell(j) + 1)^2)$ is the square of the one proposed by [Sakaie and Lowe (2007)]. There is also another method using Tikhonov regularization [Hess et al. (2006)] and some post-processing of the spherical functions to reduce errors involved in high frequency harmonics such as heuristic low-pass filtering methods in [Tournier et al. (2004); Tuch (2004)]. The important point here is that regularization is needed in HARDI estimation to reduce effects of noise. Which regularization method is best is part of ongoing research in the community. More details on this issue will be given in the next chapters on q-ball imaging and spherical deconvolution methods.

6.2.4 From SH Coefficients to HODT Coefficients

We now explicitly derive the correspondence between coefficients of the modified SH series and the independent elements of the HODT. [Ozarslan and Mareci (2003a)] showed the analytical relationship between the SH coefficients and the independent elements of the HODT. Conceptually, they showed that evaluating the ADC in terms of a rank- ℓ HODT is equivalent to fitting the ADC with a SH series truncated at order ℓ . In fact, in Appendix 6.7 we prove that both even order SH up to order ℓ and the rank- ℓ HODT polynomials restricted to the sphere are bases for the same function space. Therefore, it is possible to define a general linear transformation \mathbf{M} between both spaces. We express the coefficients of the modified SH series c_j in terms of the

$\ell = 2$, using standard spherical harmonics basis Y_ℓ^m :

$$\text{SH basis } \mathbf{Y} = \begin{pmatrix} \frac{1}{2\sqrt{\pi}} \\ \frac{1}{4}\sqrt{\frac{15}{2\pi}}\sin^2\theta e^{-2i\phi} \\ \frac{1}{2}\sqrt{\frac{15}{2\pi}}\sin\theta\cos\theta e^{-i\phi} \\ \frac{1}{4}\sqrt{\frac{5}{\pi}}(3\cos^2\theta - 1) \\ -\frac{1}{2}\sqrt{\frac{15}{2\pi}}\sin\theta\cos\theta e^{i\phi} \\ \frac{1}{4}\sqrt{\frac{15}{2\pi}}\sin^2\theta e^{2i\phi} \end{pmatrix} \quad \mathbf{T} = \begin{pmatrix} T_{xx} \\ T_{xy} \\ T_{xz} \\ T_{yy} \\ T_{yz} \\ T_{zz} \end{pmatrix}$$

$$\mathbf{M} = \begin{pmatrix} \sqrt{\frac{4\pi}{9}} & 0 & 0 & \sqrt{\frac{4\pi}{9}} & 0 & \sqrt{\frac{4\pi}{9}} \\ \sqrt{\frac{2\pi}{15}} & \sqrt{\frac{8\pi}{15}}i & 0 & -\sqrt{\frac{2\pi}{15}} & 0 & 0 \\ 0 & 0 & \sqrt{\frac{8\pi}{15}} & 0 & \sqrt{\frac{8\pi}{15}}i & 0 \\ -\sqrt{\frac{4\pi}{45}} & 0 & 0 & -\sqrt{\frac{4\pi}{45}} & 0 & \sqrt{\frac{16\pi}{45}} \\ 0 & 0 & -\sqrt{\frac{8\pi}{15}} & 0 & \sqrt{\frac{8\pi}{15}}i & 0 \\ \sqrt{\frac{2\pi}{15}} & -\sqrt{\frac{8\pi}{15}}i & 0 & -\sqrt{\frac{2\pi}{15}} & 0 & 0 \end{pmatrix}$$

$$\mathbf{MT} = \mathbf{C} = \begin{pmatrix} \sqrt{\frac{4\pi}{9}}(T_{xx} + T_{yy} + T_{zz}) \\ \sqrt{\frac{2\pi}{15}}(T_{xx} - T_{yy} + 2iT_{xy}) \\ \sqrt{\frac{8\pi}{15}}(T_{xz} + iT_{yz}) \\ -\sqrt{\frac{4\pi}{45}}(T_{xx} + T_{yy} - 2T_{zz}) \\ \sqrt{\frac{8\pi}{15}}(iT_{yz} - T_{xz}) \\ \sqrt{\frac{2\pi}{15}}(T_{xx} - T_{yy} - 2iT_{xy}) \end{pmatrix}$$

Figure 6.2: Illustration of transformation \mathbf{M} from independent elements of the HODT, \mathbf{T} , to the standard SH series coefficients, \mathbf{C} , for a rank-2 tensor. Using the standard SH basis, we obtain the exact same expression as presented in [Ozarslan and Mareci (2003a)].

$\ell = 2$, using modified SH basis Y_j (Eq. 5.21):

$$\text{SH basis} := \mathbf{Y} = \begin{pmatrix} 1/(2\sqrt{\pi}) \\ \sqrt{15}/(4\sqrt{\pi}) \sin^2 \theta \cos(2\phi) \\ \sqrt{15}/(2\sqrt{\pi}) \sin \theta \cos \theta \cos \phi \\ \sqrt{5}/(4\sqrt{2\pi})(3 \cos^2 \theta - 1) \\ \sqrt{15}/(2\sqrt{\pi}) \sin \theta \cos \theta \sin \phi \\ \sqrt{15}/(4\sqrt{\pi}) \sin^2 \theta \sin(2\phi) \end{pmatrix}$$

$$\mathbf{M} = \begin{pmatrix} 2\sqrt{\pi}/3 & 0 & 0 & 2\sqrt{\pi}/3 & 0 & 2\sqrt{\pi}/3 \\ 2\sqrt{\pi}/\sqrt{15} & 0 & 0 & -2\sqrt{\pi}/\sqrt{15} & 0 & 0 \\ 0 & 0 & 4\sqrt{\pi}/\sqrt{15} & 0 & 0 & 0 \\ -2\sqrt{\pi}/\sqrt{45} & 0 & 0 & -2\sqrt{\pi}/\sqrt{15} & 0 & 4\sqrt{\pi}/\sqrt{15} \\ 0 & 0 & 0 & 0 & 4\sqrt{\pi}/\sqrt{15} & 0 \\ 0 & 4\sqrt{\pi}/\sqrt{15} & 0 & 0 & 0 & 0 \end{pmatrix}$$

$$\mathbf{T} = \begin{pmatrix} T_{xx} \\ T_{xy} \\ T_{xz} \\ T_{yy} \\ T_{yz} \\ T_{zz} \end{pmatrix} \quad \mathbf{MT} = \mathbf{C} = \begin{pmatrix} 2\sqrt{\pi}/3(T_{xx} + T_{yy} + T_{zz}) \\ 2\sqrt{\pi}/\sqrt{15}(T_{xx} - T_{yy}) \\ 4\sqrt{\pi}/\sqrt{15}T_{xz} \\ -2\sqrt{\pi}/\sqrt{45}(T_{xx} + T_{yy} - 2T_{zz}) \\ 4\sqrt{\pi}/\sqrt{15}T_{yz} \\ 4\sqrt{\pi}/\sqrt{15}T_{xy} \end{pmatrix}$$

Figure 6.3: Rank-2 example of the important matrices in the algorithm. Note that coefficients of \mathbf{Y} are real. The change-of-basis matrix \mathbf{M} relating the SH coefficients \mathbf{C} and HODT independent elements \mathbf{T} are given.

independent elements of the high order tensor \mathbf{T} . From Eq. 6.2, we have the relation

$$c_j = \int_{\Omega} D(\mathbf{g}(\theta, \phi)) Y_j(\theta, \phi) d\Omega. \quad (6.14)$$

We can replace $D(\mathbf{g})$ with the high order tensor formulation of Eq. 6.6 and obtain an expression in matrix form, where c_j is the j th element of vector $\mathbf{C} = \mathbf{MT}$ and $R = (1/2)(\ell + 1)(\ell + 2)$ is the number of elements in the SH basis:

$$c_j = \sum_{k=1}^N D_k \int_{\Omega} \mu_k \prod_{p=1}^{\ell} g_{k(p)}(\theta, \phi) Y_j(\theta, \phi) d\Omega \implies \mathbf{C} = \mathbf{MT}, \text{ where} \quad (6.15)$$




Figure 6.4: Sketch of our regularization algorithm. It illustrates the main equations of the method and most importantly, it shows the bridge between SH coefficients and the independent coefficients of the HODT.

$$\mathbf{M} = \begin{pmatrix} \mu_1 \int_{\Omega} \prod_{p=1}^{\ell} g_{1(p)}(\theta, \phi) Y_1(\theta, \phi) d\Omega & \dots & \mu_N \int_{\Omega} \prod_{p=1}^{\ell} g_{N(p)}(\theta, \phi) Y_1(\theta, \phi) d\Omega \\ \vdots & \ddots & \vdots \\ \mu_1 \int_{\Omega} \prod_{p=1}^{\ell} g_{1(p)}(\theta, \phi) Y_N(\theta, \phi) d\Omega & \dots & \mu_N \int_{\Omega} \prod_{p=1}^{\ell} g_{N(p)}(\theta, \phi) Y_N(\theta, \phi) d\Omega \end{pmatrix}. \quad (6.16)$$

From Appendix 6.7, the $R \times R$ square matrix \mathbf{M} is a change-of-basis matrix, and thus is invertible. Therefore, given a vector \mathbf{C} of SH coefficients, we can use \mathbf{M}^{-1} to compute the corresponding vector of HODT coefficients.

Hence, we obtained a linear mapping between SH coefficients and HODT independent elements. When using the standard SH basis of Figure 6.2, we obtain the exact same relationship presented by [Ozarslan and Mareci (2003a)]. Figure 6.3 illustrates the transformation procedure analytically for the example of the rank-2 diffusion tensor. Note that all matrix entries are real, which captures the real-valued constraint on the problem.

Putting everything together, the linear transformation taking the ADC \mathbf{X} to the vector \mathbf{T} of HODT independent coefficients using our regularization method with is

$$\boxed{\mathbf{T} = \mathbf{M}^{-1}(\mathbf{B}^T \mathbf{B} + \lambda \mathbf{L})^{-1} \mathbf{B}^T \mathbf{X}}. \quad (6.17)$$

The technique is fast because it only involves linear transformations and more importantly, matrices \mathbf{B} (Eq. 5.24) and \mathbf{M} (Eq. 6.16) are the same for all voxels in the dataset and need only be computed once. In the end, we have a handle on both the SH series coefficients and the independent elements of the HODT. The important equations of the algorithm and the bridge between SH and HODT coefficients are clearly sketched in Figure 6.4.

6.3 HIGH ORDER ANISOTROPY MEASURES FROM ADC PROFILES

The ultimate goal when estimating and denoising the ADC profile is to use it to infer the underlying diffusion process. This is commonly done using anisotropy measures defined on the set of coefficients used to describe the ADC profile. In the literature, there are two classes of high order anisotropy measures based on the ADC profile of HARDI data. One class is defined on the coefficients of a SH series [Frank (2002); Alexander et al. (2002); Chen et al. (2004b)] and the other on the independent elements of a HODT [Ozarslan and Mareci (2003a)]. One important advantage of our algorithm is that we have a handle on both the SH and HODT coefficients and we thus can implement all these measures.

6.3.1 Frank and Chen *et al* Measures

The two measures found in [Frank (2002)] and in [Chen et al. (2004b)] papers are similar and based on SH coefficients. Frank proposed a simple Fractional Multifiber Index (FMI) that is the ratio of the sum of squared high order coefficients over order-2 coefficients. It is given by

$$\text{FMI} = \frac{\sum_{\{j:\ell \geq 4\}} |c_j|^2}{\sum_{\{j:\ell=2\}} |c_j|^2}. \quad (6.18)$$

Note however that limits and range of the FMI and other such ratios based on the SH coefficients vary depending on the input HARDI signal and cannot be known a priori.

[Chen et al. (2004b)] claim that the FMI ratio was insufficient to separate isotropic, 1-fiber and multi-fiber behavior within a voxel. They decided to introduce the variance of the ADC profile about its mean into the characterization of the underlying diffusive behavior with the following ratios:

$$R_0 = \frac{|c_0|}{\sum_{\forall j} |c_j|}, \quad R_2 = \frac{\sum_{\{j:\ell=2\}} |c_j|}{\sum_{\forall j} |c_j|}, \quad R_{multi} = \frac{\sum_{\{j:\ell \geq 4\}} |c_j|}{\sum_{\forall j} |c_j|}. \quad (6.19)$$

In their paper, they restricted the maximum number of fibers in a voxel to 2 and they used a rank-4 approximation. Hence, they only used R_0 and R_2 ratios. We generalize

their approach for more than two directions and use a rank-8 estimation. R_{multi} is a natural multi-fiber ratio extension. These anisotropy ratios can be used to formulate a 3-step algorithm to distinguish between isotropic, 1-fiber and multi-fiber diffusion. That is, large R_0 and/or small ADC variance is isotropic; if not, then large R_{multi} is multi-fiber; anything else is 1-fiber. As mentioned in [Chen et al. (2005)], there are several thresholds involved in this method that must be carefully picked by performing tests on synthetic data to obtain the best results. To overcome the previous limitations, Chen et al proposed to use cumulative residual entropy (CRE) [Rao et al. (2004)] in their recent paper [Chen et al. (2005)]. They used this new information theory measure together with the estimated diffusivities and claimed that it can classify diffusion processes with only two ratios. CRE is based on a probability measure computed from the discrete ADC profile estimates whereas we are focusing on measures defined from the coefficients of the SH or HODT vector parameterizing the ADC profile. An implementation, comparison and longer discussion can be found in [Descoteaux et al. (2005a)]. Note that another possible anisotropy measure using entropy and information theory is defined in [Ozarslan and Mareci (2003b)].

6.3.2 Alexander *et al* Measure

ANOVA is used in [Alexander et al. (2002)] to determine if truncating the series at a higher order as opposed to a lower order significantly changes the fit to the model data. The *F-test* is

$$F(M_2, M_\ell) = \frac{(n_s - p_\ell - 1)(Var(M_\ell) - Var(M_2))}{(p_\ell - p_2)E(M_\ell)} \quad (6.20)$$

where E is the mean squared error between the true ADC profile and the estimated ADC profile at the n_s sampled points and p_i is the number of free parameters in model M_i . In our case, the lower order model is M_2 (order-2 diffusion tensor) and we test it against higher order approximation models M_ℓ , for $\ell = 4, \dots, L$. In [Alexander et al. (2002); Blyth et al. (2003)], ANOVA was used to dynamically choose the order of the SH series approximation.

6.3.3 Generalized Anisotropy Measure

In [Ozarslan et al. (2005b)] the well-known FA measure computed with the DT is generalized for HARDI data fitted with HODT. Their generalized anisotropy (GA) measure is based on the generalization of the trace, generalized mean diffusivity and the variance of the normalized diffusivity of HODT. The trace of a rank-2 tensor \mathbf{T} is given by the sum of its eigenvalues, and if it is antipodally symmetric, it can be expressed as the integral over the unit hemisphere Ω of its quadratic form

$$\text{trace}(\mathbf{T}) = \frac{3}{2\pi} \int_{\Omega} \mathbf{g}^T \mathbf{T} \mathbf{g} d\mathbf{g}.$$

Thus, it is possible to generalize this trace expression for any antipodally symmetric spherical function f and obtain the generalized trace “gentr”,

$$\text{gentr}(f(\mathbf{g})) = \frac{3}{2\pi} \int_{\Omega} f(\mathbf{g}) d\mathbf{g}.$$

Inspired by the rank-2 tensor case where FA is expressed in the form

$$FA = \sqrt{\frac{1}{2} \left(3 - \frac{1}{\text{trace}(\mathbf{T}_N^2)} \right)}$$

related to the variance of the eigenvalues of the normalized tensor ($\mathbf{T}_N = \mathbf{T}/\text{trace}(\mathbf{T})$) in the case that \mathbf{T} is diagonalized, Ozarslan et al propose a generalized version using the variance of the generalized trace of the normalized ADC, $D_N(\mathbf{g})$. This generalized variance is defined as

$$V = \frac{1}{3} \left(\text{gentr}(D_N(\mathbf{g})^2) - \frac{1}{3} \right), \quad \text{where} \quad D_N(\mathbf{g}) = \frac{D(\mathbf{g})}{\text{gentr}(D(\mathbf{g}))} \quad (6.21)$$

The last step of the GA measure is to make sure it is restricted to the interval $[0, 1)$ by choosing a monotonic function that maps the interval $[0, \infty)$ to $[0, 1)$. This choice seems to be relatively “ad-hoc” and dependent on the dataset to be processed. The final generalized anisotropy (GA) measure is defined as

$$GA = 1 - \frac{1}{1 + (250V)^{e(V)}}, \quad \text{where} \quad e(V) = 1 + \frac{1}{1 + 5000V}. \quad (6.22)$$

We have skipped some important details in the GA derivation and we strongly refer the reader to the article [Ozarslan et al. (2005b)] for the complete description. Contrary to ratios/measures/algorithms proposed from spherical harmonic coefficients reviewed in this section, GA has the property of being scaled between 0 and 1 and does not assume any specific approximation order. It can be quickly and easily visualized to obtain cues on the anisotropy regions, just as FA in the DTI case.

6.4 RESULTS

We now evaluate the performance of our ADC profile fitting procedure compared with other techniques in the literature and secondly test the different HARDI anisotropy measures on both synthetic and real data. First, we want to find the optimal regularization parameter.

6.4.1 Optimal Regularization Parameter with the L-Curve Method

As mentioned earlier, the proposed algorithm penalizes the approximated SH and HODT coefficients for having higher order terms unless they significantly improve the overall accuracy of the estimation. This eliminates most contribution due to noise while leaving those that are necessary to describe the underlying ADC profile. However, obtaining this balance depends on choosing a good value for the parameter λ .

There are several numerical methods to find the optimal regularization parameter such as L-curve [Hansen (2001)] and the Generalized Cross Validation (GCV) [Wahba (1990); Sakaie and Lowe (2007)]. In this thesis, we chose the L-curve approach.

The L-curve is a plot of the norm of a regularized solution versus the norm of the corresponding residual form. The idea is to vary the values of the regularization parameter λ and to plot the points corresponding to regularization errors on the ordinate and the data perturbation errors on the abscissa. In our minimization problem given in equation $M(\mathbf{C})$ (Eq. 6.12), we record the points $(x, y) = (\|\mathbf{X} - \mathbf{BC}\|^2, \mathbf{C}^T \mathbf{LC})$ while varying λ . This gives a set of noisy points to which one can fit a *best* curve to obtain the L-curve. The optimal λ is then the point on the curve with maximum curvature. This is the location that separates the flatter and more vertical part of the graph, where the solution is dominated by perturbation errors and regularization errors respectively. In our problem, we choose a 4th order polynomial least-squares fit, which models the data well.³ We find parameters $\{a, b, c, d, e\}$ that best model the data such that $y(x) = ax^4 + bx^3 + cx^2 + dx + e$. To find the optimal λ , we seek the discrete point x where the curvature $\kappa(x)$ of the L-curve $y(x)$ is maximum. The curvature of a 2D curve expressed as $y(x)$ is given by $\kappa(x) = y''(x)/(1 + (y'(x))^2)^{3/2}$. Since $y'(x) = 4ax^3 + 3bx^2 + 2cx + d$ and $y''(x) = 12ax^2 + 6bx + 2c$, we easily get an analytic expression for the L-curves curvature and its derivative $\kappa'(x)$. In practice, we have discrete sampling of these curves and the optimal λ is then simply given by the λ associated with the discrete x value where $\kappa'(x) = 0$. Note that this is an automatic way to find the optimal λ .

L-curve Numerical Simulation

In order to find the optimal λ for our regularization algorithm, we must plot the L-curves for 1, 2, and 3 fiber distributions. It is important to note that there is a different optimal λ depending on the underlying fiber distribution. In fact, in the cases of 1 or no fibers, a HODT formulation is unnecessary and thus, it is expected to obtain a higher optimal λ for these cases than for multiple fiber diffusion. As described above, we need to plot the data errors $\|\mathbf{X} - \mathbf{BC}\|^2$ on the abscissa and the regularization term $\mathbf{C}^T \mathbf{LC}$ on the ordinate while incrementally varying $\lambda \in [0, 0.5]$. We systematically increment λ by 0.0005 for each set of simulated data. Hence, we simulate three separate tests with 1, 2, and 3-fiber distributions respectively. We decide to choose a relatively high b -value, $b = 3000 \text{ s/mm}^2$, to be able to better distinguish 2 and 3 fibers within a voxel, as noted in [Alexander and Barker (2005)]. A b -value of 1000 s/mm^2 is more standard for clinical diffusion MRI acquisitions but research-oriented acquisitions tend use higher b -values. We perform an order-8 approximation. For each L-curve plot, note that the number of fibers in the simulated ADC profile, the noise level and the approximation order ℓ are the same for all green points whereas the λ

³Before a polynomial curve can be fit, there is a necessary sorting of the x and the corresponding y and associated λ . Otherwise, the curve can potentially come back on itself.

parameter, the relative weights (p_k 's) of fibers and the orientation angle(s) and angle(s) between fibers change for every point in the graph.

L-curves Results


The L-curves for 1, 2 and 3 fiber distributions are plotted in Figure 6.5. Each curve was created separately by estimating the ADC profile while incrementally varying λ by 0.0005 from synthetic noisy ADC profiles with random fiber orientations and fiber compartments fractions. We show the set of measured points ($\|\mathbf{X}-\mathbf{BC}\|^2, \mathbf{C}^T\mathbf{LC}$), the fitted L-curve and indicate the optimal estimated λ with a circle. From the L-curves the optimal λ are 0.308, 0.006 and 0.0155 for 1, 2, and 3 fiber distributions respectively.

Numerical Simulations

One can now ask how valid these optimal λ 's really are and in particular, how well do they agree with the initial true signal without noise. Hence, in order to verify the L-curve optimal values, we test the effectiveness of recovering the original (without noise) synthetic ADC profile from noisy sparse data measurements as above. In doing so, we compare our algorithm with SH estimation without regularization of [Frank (2002); Alexander et al. (2002)] and HODT fitting of [Ozarslan and Mareci (2003a)]. Table 6.1 summarizes how we obtain each HODT estimation. For our algorithm, \mathbf{T} is obtained from Eq. 6.17. For the [Ozarslan and Mareci (2003a)], we use Eq. 6.8 to obtain \mathbf{T} and for the un-smoothed SH technique, we use Eq. 5.29 to obtain \mathbf{C} which is essentially the same as techniques proposed by [Frank (2002)] and [Alexander et al. (2002)] except that their SH basis is different. In the latter case, we need to apply our transformation matrix \mathbf{M}^{-1} (Eq. 6.16) to obtain \mathbf{T} . For all cases, we obtain the estimated ADC profiles on the unit sphere for any angle using Eq. 6.6. In the following, the errors recorded are the mean and standard deviation (std) of the error vector specified by the point-wise absolute difference between points on the estimated ADC and points on the true ADC profiles (Eq. A.2 of Appendix A.1) before adding noise)


	[Ozarslan and Mareci (2003)]	[Frank (2002)] [Alexander et al. (2002)]	Our method
Regularization	no	no	yes ($\lambda = 0.006$)
Signal \mathbf{C}	Not applicable	$(\mathbf{B}^T\mathbf{B})^{-1}\mathbf{B}^T\mathbf{X}$	$(\mathbf{B}^T\mathbf{B} + \lambda\mathbf{L})^{-1}\mathbf{B}^T\mathbf{X}$
HODT \mathbf{T}	$(\mathbf{R}^T\mathbf{R})^{-1}\mathbf{R}^T\mathbf{X}$	$\mathbf{M}^{-1}\mathbf{C}$	$\mathbf{M}^{-1}\mathbf{C}$

Table 6.1: Summary of each algorithm used to compare HODT estimation. As before, \mathbf{X} represents the measured HARDI signal, \mathbf{C} the SH representation of \mathbf{X} , \mathbf{B} is the SH basis, \mathbf{R} is the HODT matrix for the generalized Stejskal-Tanner equation and \mathbf{M} is the change of basis matrix between HODT elements and SH coefficients.



1 fiber, $\lambda = 0.308$

2 fibers, $\lambda = 0.006$



3 fibers, $\lambda = 0.0155$

Figure 6.5: L-curves for 1, 2 and 3 fibers with optimal λ circled. Points represent data perturbation errors and regularization errors ($\|\mathbf{X} - \mathbf{BC}\|^2, \mathbf{C}^T \mathbf{LC}$).

for different regularization parameter λ . We generate a single simulation with 1000 synthetic ADC profiles of 1, 2 and 3 fibers mixed together.

From Table 6.2, there are several expected observations that one can make. Firstly, from the first two columns, we note that the LR fit of a HODT (Eq. 6.8) and the SH series fit with no regularization (Eq. 6.4) are the same up to a very small numerical error. As previously discussed, this is to be expected because the two models are equivalent and there exists a direct change-of-basis transformation between the two (Eq. 6.16). Secondly, as observed with the L-curve, single fiber voxels are optimally fit with larger λ than in the multiple fiber voxels. This occurs because in the case of a single fiber, higher order estimation is unnecessary. Intuitively, a large λ in Eq. 6.13 cancels the effect of high order tensor terms so that the model is closer to a rank-2 tensor formulation. However, as soon as we have more structure, as in the cases of 2 or 3 fibers, λ of 0.006 and 0.015 are best, respectively. This is expected as we need to maintain significant contributions from higher order terms in order to accurately fit multi-fiber ADC profiles. If this is so, one can wonder why the optimal λ for 3-fiber is higher than for 2-fiber distributions. This is because as fibers are added within a voxel, the signal tends to be more isotropic due to partial volume averaging. Thus, the need for higher terms becomes less needed in the approximation. It is important to note that these best simulated λ values closely agree with the L-curve points of maximum curvature, which validates the automatic and analytic L-curve optimal λ procedure. Finally, note that the optimal value in the 2-fiber case is quite broad for different λ values which suggest that there is a range of possible λ one can choose without changing the resulting estimation. In practice, we have noted that the exact optimal λ chosen is flexible within a small range that depends on the dynamics of the signal. Nonetheless, since in all cases $\lambda = 0.006$ performs better than both LR and the SH fit with $\lambda = 0$ and the optimal value found with the L-curve in the case of 2 fibers is $\lambda = 0.006$, we choose to set $\lambda = 0.006$ for the rest of this thesis (unless specified). This avoids having to compute optimal λ at each voxel of the dataset.

Qualitatively, Figure 6.6 illustrates how we obtain a smoother ADC profile estimation that visually agrees with our quantitative results. Our regularized tensor estimation gives a ADC profile closer to the true ADC without noise. Here, we have stretched the surface mesh with respect to each ADC profile value $D(\mathbf{g}(\theta_i, \phi_i))$. The color map is red for high values and blue for low ADC values.

6.4.2 High Order Anisotropy Measures Results

We implemented the high order anisotropy measures defined in Section 6.3 to study their ability to identify the underlying synthetic diffusion process. We have used the same synthetic data generation process as before with Gaussian complex noise with standard deviation $\sigma = 1/35$ and $b = 3000$ s/mm². We create a large set of simulated profiles where the number of fibers is picked randomly between 0 and 3. Our goal is to correctly classify synthetic noisy data, generated as before, into three

		<i>1 fiber test with λ</i>						
		LR	0	0.006	0.012	0.1	0.3	0.5
mean		0.083	0.083	0.071	0.068	0.052	0.051	0.054
std		0.064	0.063	0.051	0.046	0.036	0.035	0.037
		<i>2 fiber test with λ</i>						
		LR	0	0.003	0.006	0.009	0.012	0.015
mean		0.076	0.075	0.070	0.069	0.070	0.070	0.072
std		0.052	0.052	0.043	0.041	0.040	0.041	0.042
		<i>3 fiber test with λ</i>						
		LR	0	0.006	0.009	0.012	0.015	0.1
mean		0.092	0.092	0.049	0.040	0.034	0.031	0.057
std		0.037	0.037	0.028	0.026	0.025	0.025	0.026

Table 6.2: Numerical simulation to test the behavior of our algorithm with varying regularization parameter λ . LR corresponds to the linear regression (LR) method of [Ozarslan and Mareci (2003a)], Eq. 6.8. We report the mean and standard deviation of the point-wise absolute difference between true (without noise) and estimated ADC profiles.

classes: isotropic diffusion, 1-fiber anisotropic diffusion and multi-fiber diffusion, i.e. diffusion corresponding to either 2 or more fibers.

First, we show the results of computed GA values in Table 6.3 for different order approximations. Even though the exact values are different because we use a different synthetic data generation, we are able to reproduce the same behaviors of the GA measure published by [(Ozarslan et al., 2005b, Table 4)]. In particular, for a single fiber, the measures are almost independent of the rank being used for the estimation (first row in table). However, GA measures are significantly higher for rank- $\{4,6,8\}$ than for the rank-2 model when there are multiple fibers. There is also an overlap in GA values between diffusion in 3-fiber distributions and isotropic diffusion. It is in these cases that the rank-2 model is unsatisfying. We have computed a similar table for all other anisotropy measures (not shown here). One can then attempt to find the best thresholds that separate the different measures into these three classes for all anisotropy measures. We test the robustness of the classifications with and without smoothing parameter while changing the rank of the approximation. We also analyze the behavior of the different measures by recording the classification success rate.

From Table 6.4, we see that the best results for Chen et al (Eq. 6.19) and ANOVA (Eq. 6.20) obtained from $\lambda = 0.006$ produce better classification percentages than the best outcome of FMI (Eq. 6.18). FMI is the simplest of the measures and is not able to capture as many details as the other two methods which include statistical information and a better comparison between estimation order used to model the data.




Figure 6.6: Visualizing the ADC profile fitting procedure. We have in row 1 the ADC profile for a single synthetic fiber, in row 2 the ADC profile for two orthogonal fibers and row 3 is for three orthogonal fibers. Qualitatively, the last column gives the best approximation of the true ADC, confirming the L-curve results and Table 6.2

	Our method ($\lambda = 0.006$), LR			
	$\ell = 2$	$\ell = 4$	$\ell = 6$	$\ell = 8$
1 fiber	0.92, 0.92	0.92, 0.92	0.92, 0.92	0.92, 0.92
2 fibers	0.64, 0.65	0.71, 0.73	0.72, 0.73	0.70, 0.71
3 fibers	0.09, 0.09	0.19, 0.21	0.20, 0.13	0.10, 0.09
isotropic	0.09, 0.09	0.09, 0.10	0.09, 0.09	8×10^{-4} , 0.01

Table 6.3: Mean GA measure over 10000 simulated ADC profiles. This table shows the mean GA calculated from our algorithm and Ozarslan’s LR method. We observe similar overall behaviors as the ones reported in [(Ozarslan et al., 2005b, Table 4)].

ℓ	FMI		Chen et al		ANOVA		GA	
	$\lambda = 0$	$\lambda = 0.006$	$\lambda = 0$	$\lambda = 0.006$	$\lambda = 0$	$\lambda = 0.006$	$\lambda = 0$	$\lambda = 0.006$
8	93.1%	85.1%	83.2%	96.0%	91.7%	94.8%	99.9%	99.8%
6	93.1%	84.2%	89.1%	97.0%	90.9%	94.5%	99.6%	99.8%
4	83.2%	82.2%	93.1%	97.0%	88.7%	91.1%	99.9%	100%
2			54.4%	54.5%			97.6%	97.6%

Table 6.4: Success rate for the classification of synthetic voxels into isotropic, 1-fiber and multi-fiber diffusion class using different high order anisotropy measures. For each measure, we select thresholds that best distinguish the three class. FMI and ANOVA measure make no sense in the $\ell = 2$ case.

The advantage of Chen et al’s algorithm over other measures is that it involves well-chosen SH coefficient ratios and the variance of ADC profile about its mean. Experimentally, we find it more robust to noise and less sensitive to the threshold selection. In fact, we do not have to adapt thresholds between tests when changing λ and the rank of the HODT. This is not the case with ANOVA. It is also important to note the better performance of our regularization procedure with a non-zero λ used in conjunction with the Chen et al and the ANOVA measures than with no regularization ($\lambda = 0$). This illustrates the added value brought by our regularization method. This is because the estimated ADC profile is smoother, and hence, the SH and HODT coefficients incorporate less noise. On the other hand, for FMI measure, we note that it performs better when $\lambda = 0$. This is not surprising because of the nature of the ratio and in particular, because the numerator depends only on high order terms of the SH series. With a non-zero λ , the coefficients of higher order are penalized, and therefore have lower values than without a smoothing term. The ratio is thus not designed to be used with our method. Note also the higher success rate of the GA measure compared with the other measures. Being a high order generalization of FA, it is better theoretically founded and its characterization is better than with ‘ad hoc’ ratios on SH coefficients. In our experiments, we also note that the GA values are more stable, i.e. have less variations between random trials, which makes it easier to find the best thresholds.

In order to validate the approach on a field of noisy spherical functions, we have also extended our single voxel synthetic generation to construct a noisy high angular resolution 30 x 30 slice with a 5 x 5 crossing configuration in the center Figures 6.7 and 6.8. The ADCs were estimated with an order 6 SH series and $\lambda = 0.006$. We perform one experiment with two and another with three orthogonal fibers crossing. Note that all voxels in this slice have the same mean diffusivity. We compare the shape of the DTI ellipsoids with the shape of the ADC profiles computed with an order 8 SH basis and the resulting FA and GA anisotropy measures. Although we choose only to show results of the FA measure, it is important to note that RA and FA




Figure 6.7: Two orthogonal fiber bundles crossing at 90° . The color map for the DTI ellipsoids is blue for isotropic, green for planar and red for anisotropic ellipsoids. Note that ellipsoids are flat and have no preferred principal direction in the crossing area whereas ADC profiles have multiple peaks. The last row shows the FA and GA anisotropy maps. The anisotropy is much lower in the crossing area in both cases but higher in the GA case than in the FA.

have similar properties [Ozarslan et al. (2005b)] and there is no notable differences between the two measures when comparing them against GA. The DTI ellipsoids were computed using standard least-squares fit from all discrete high angular measures. The color map is blue for isotropic, green for planar and red for anisotropic ellipsoids. While the ellipsoids are flat and have no preferred principal direction in the crossing area of Figs. 6.7 and 6.8, the ADC profiles clearly have multiple peaks reflecting multiple fibers. Although the multiple peaks do not agree with the underlying fiber orientations⁴, the high order anisotropy measures like GA are able to pick up these differences. The limitation of FA measure is even more evident on a 3-fiber distribution example in Figure 6.8. The FA measure is in the same range as the background whereas the GA map is able to distinguish the multi-fiber diffusion processes.

⁴As seen in Chapter 4, diffusion profile maxima do not agree with fiber directions for two or more underlying fibers. One must define functions in real space like the orientation distribution function (ODF) to extract fiber orientations.




Figure 6.8: Three orthogonal fiber bundles crossing at 90° where one bundle is coming out of the page. GA can distinguish the crossing area whereas the FA cannot. The last row illustrates the effect of varying the threshold in the classification process. blue := isotropic, red := 1-fiber, green := multi-fiber class respectively. Some voxels are incorrectly classified due to noise and partial volume averaging.

In the last row of Figure 6.8, we see the effects of varying the threshold separating isotropic from multi-fiber diffusion. The threshold are selected experimentally and we see that the critical threshold is in the range of the mean GA measure of the $\ell = 8$ column in Table 6.3. We see that GA values can discriminate the different diffusion processes with some errors made when there are overlap between noisy 3-fiber distributions and noisy isotropic diffusion signals. If the 3-fiber threshold is too high, some multi-fiber voxels are incorrectly classified as isotropic and vice versa when the 3-fiber threshold is too low. This is due to noise and partial volume averaging.

Finally, it is worth mentioning that the choice of noise level and b -factor surely affect the classification results as one would expect. Noisy data makes classification

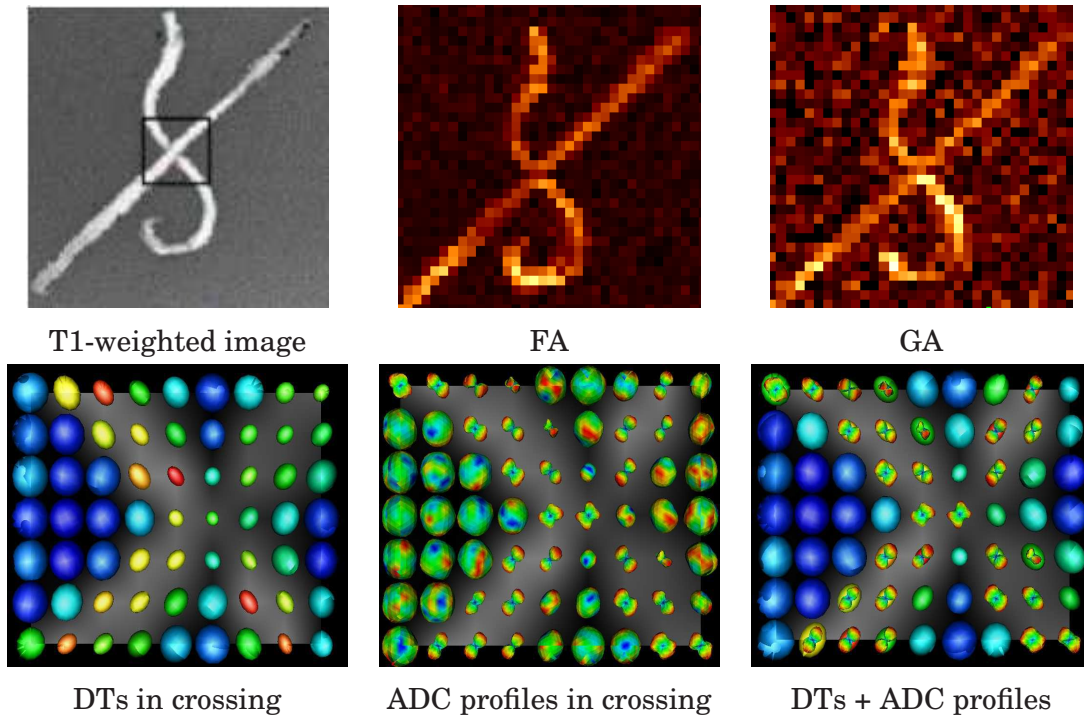


Figure 6.9: Biological phantom produced by [Campbell et al. (2005)] described in Appendix A.2. We show the baseline T1-weighted image and fields of DTI ellipsoids and estimated ADC profiles overlaid on the GA map in the crossing area. The DTI ellipsoids tend to be flat (greenish) and spherical (bluish) whereas the ADC profiles clearly have multiple peaks.

harder for all classes since major overlap appear between 1-, multi- and isotropic distributions. A lower b -value also produces less reliable results in multi-fiber voxels. This is because the ADC profiles are wider and less sharp than for high b -value.

Conclusion of Synthetic Data Experiments

There are several important messages from this section. Most importantly, we have seen that regularization is important to obtain an accurate HARDI signal estimation and a robust ADC reconstruction that agrees closely with ground truth. Then, we have seen numerical experiments that illustrate the equivalence of the SH estimation and HODT estimation, known in theory. Finally, we have seen that regularization and HARDI anisotropy measures can be exploited to discriminate isotropic voxels from voxels with single fiber distribution and multiple fiber distributions. In particular, although the multiple peaks of the ADC do not agree with the underlying fiber distribution, high order anisotropy measures can be used to detect these multiple fiber distributions.

6.4.3 Biological Phantom Results

Figure 6.9 shows the results on the biological phantom. The acquisition of this dataset is described in Appendix A.2. The baseline T1-weighted image illustrating the ground truth configurations of the fibers is seen in the top left of Figure 6.9. We compare the ellipsoid surfaces obtained from the standard DTI least-squares fit using all 90 directions against the ADC profile computed with our regularized SH estimation. The ellipsoids and ADC profiles are overlaid on the computed GA map. A rank-8 estimation with $\lambda = 0.006$ was used. As expected, the DTI ellipsoids are planar and tend to be spherical in the crossing area whereas our high order estimation of the ADC profile recover multiple peaks. Moreover, even though the GA map has slightly noisy background, it recovers both fiber bundles and is different in the crossing than in certain straight single fiber parts. However, contrary to the synthetic data characterization, the GA difference between 1- and 2-fiber regions is not clear enough to fully distinguish the two with simple thresholding.

6.4.4 Human Brain HARDI Data Results

The human brain dataset used is the CMRR dataset described in Appendix A.3.2. Based on expert knowledge of the anatomy and the FA anisotropy measures, voxels with one and two fibers in the corpus callosum and the corticospinal tract were selected to illustrate the performance of our regularization algorithm on real data. A rank-8 approximation with $\lambda = 0.006$ was used. Figure 6.10 shows some of the results.

Classification on real data is a much more difficult task than the synthetic data experiments where the noisy generated signal is more homogeneous and GA statistics are easily computed. On real data, one has to manually select regions with known crossings and regions with single fiber anisotropy to analyze the corresponding GA measure behavior. Due to noise and partial volume effects there is too much overlap between diffusion process classes. Nonetheless, we observe that the measured data ADC profiles have very similar shapes and properties as compared to our synthetic data. Even though we do not believe one can simply threshold the real data GA map to classify the diffusion process, one can make better use of the GA than FA scalar measure to look for multiple fibers crossing. Visual inspection of the peaks of the ADC profile in the small regions of interest shows agreement with known neuroanatomy (Figure 6.10) of single and multiple fiber bundles. In the left subfigure, two profiles were chosen in the corpus callosum, one with 1-fiber and the other with 2-fiber distribution evidence. The FA map illustrates the location where the 2-fiber distribution was extracted with a low FA value of 0.35. In the right sub-figure, profiles were extracted from a 1-fiber region of the corticospinal tract and a nearby 2-fiber crossing with the corpus callosum. The FA map shows the location of the 1-fiber diffusion voxel with a high FA value of 0.67.





Figure 6.10: ADC profiles selected from real data voxels in the corpus callosum and the corticospinal tract. We have manually selected voxels with one and two fibers based on knowledge of the anatomy and FA measure. The left column in each sub-figure represents the measured profile and the right column is the estimated ADC profile from our algorithm. We report the corresponding FA and GA measures for each and include the FA map with the region of interest marked with a pink overlaid cross.



Transverse, coronal and sagittal FA slice of the region of interest




Figure 6.11: Region of interest from the genu of the corpus callosum, similar as in [(Tuch, 2004, fig. 12-f)]. We show FA, GA and the thresholded classification map obtained. Green: multi-fiber, red: 1-fiber and blue: isotropic diffusion class respectively. Green voxels agree with known crossings in that area.

Finally, we reproduce a similar figure as in [(Tuch, 2004, Fig. 12-f)] from the genu of the corpus callosum in Figure 6.11. This region is known to have multiple fibers crossing. We thus observe a lower FA than that found in strong anisotropic areas and a stronger GA value in the range corresponding to 2 or 3 fibers in Table 6.3. If one locally thresholds this region of interest, the obtained classification provides some valuable information. The multi-fiber regions in red clearly stand out and we can thus identify areas of transition between single (blue) and multiple (red) fiber diffusion class. The classification of Figure 6.11 agrees with voxels containing multiple fiber crossings such as presented in [(Tuch, 2004, Fig. 12-f)].

6.5 DISCUSSION

Overall, there are several contributions in this chapter. Firstly, we proposed a new regularization algorithm to estimate a smoother ADC profile closer to the true ADC profile without noise. We chose a meaningful modified SH basis to

capture the physical constraints of the problem, defined a regularization term based on the smoothing properties of the Laplace-Beltrami operator for functions defined on the unit sphere and derived the linear transformation taking SH coefficients to HODT coefficients. We thus obtained a HODT from a regularized SH estimation of the HARDI signal. The method was tested against state-of-the-art techniques from the literature on synthetic data, on a biological phantom and on real human brain data. Our experiments verify the equivalence of the HODT and the SH formulations of the ADC profile up to a good numerical approximation. This is useful because it implies that any technique developed for the SH formulation can be quickly and easily applied to the diffusion tensor formulation and vice versa.

Moreover, the continuous smoothing parameter λ generalizing the traditional least-squares fitting algorithm for SH series is appropriate whenever there is a sparsely sampled set of noisy data on a sphere and the expected approximation function is relatively smooth. However, minimizing the squared error is not the only important characteristic of a good approximation function. In particular, if we wish to be able to extract a fiber probability distribution for the voxel from our ADC profile model, we need the approximation to be sufficient to capture details up to the desired level of resolution. For example, over a data set that contained a large proportion of voxels with only a single fiber direction, an order-2 fit would certainly minimize the squared error from the actual data averaged over the entire data set. However, the higher order voxels would be modeled very poorly, making various tasks, such as tractography, considerably more difficult. For this reason, simple truncation, which is the equivalent to fitting the data with a lower order tensor is not comparable to a procedure that retains higher-order detail, such as the proposed fit with a well-chosen regularization weight λ . Up to the resolution that we were considering, order-4 truncation was a valid choice of fitting procedure but does not have the ability to retain detail at the 6th and 8th orders shown by the smoothness maximizing criterion we propose.

Furthermore, we have used the modified SH basis of the previous chapter. This basis is physically meaningful and adapted to the HARDI modeling problem because it is real and symmetric. However, note that the physical positive constraint of the ADC profile and HARDI signal is not ensured in our derivation, i.e. the property that $D(\mathbf{g}) > 0$, for any \mathbf{g} is not ensured in theory. This is an interesting open question that is the subject of much current research in the field (see [Tournier et al. (2007); Jian and Vemuri (2007b)]). In implementing our algorithm, we found that this was rarely relevant at reasonable noise levels, as our algorithm gave everywhere positive results without being explicitly constrained to do so. However, in the standard rank-2 DTI problem, methods have been proposed to deal with data where some of the estimated tensors step out of the positive-definite matrix space [Tschumperlé and Deriche (2001); Chéfd'hotel et al. (2004); Wang et al. (2004); Arsigny et al. (2006); Lenglet et al. (2006b)]. The question of how to extend these

approaches in the space of HODT is still open. Recent works on different HODT formulations have started to appear in the literature such as [Basser and Pajevic (2007); Barmpoutis et al. (2007)]. It is part of current work to look for numerical schemes in the space of HODT that impose a positive diffusion ADC reconstruction.

In addition to the contributions mentioned above, another important part of the chapter is the careful study of the existing anisotropy measures for SH and HODT representation of the HARDI ADC profiles. We extended the study by attempting to classify the diffusion process from these anisotropy measures. This showed to be easier on synthetic data than on real data where selecting sensible thresholds is much harder. We have successfully reproduced the properties of the GA measure and clearly showed its advantages over the other scalar indices computed from SH coefficients. In addition to being the high order generalization of the popular FA measure for DTI, it is stable and has the advantage of being scaled between 0 and 1. With a fast visual inspection, this allows the detection of non-Gaussian diffusion processes in cerebral regions from real data. This can potentially be of great help to clinicians and neurosurgeons. However, even though the ADC modeling from HARDI data can describe intravoxel fiber crossings, it is important to point out that, if one is interested in the underlying fiber distribution for fiber tractography, the ADC is not the right function to be used. The ADC maxima do not agree with fiber directions for two or more underlying fibers and thus cannot be used for fiber tractography. One must use functions in real space like the diffusion PDF and diffusion ODF to extract correct fiber orientations. Therefore, it is now important to look into HARDI reconstruction algorithms that estimate functions with their maxima aligned with the underlying fiber distribution.

6.5.1 Contributions of this chapter:

In summary, the contributions of this chapter are:

- The new Laplace-Beltrami regularization algorithm to estimate the HARDI signal and the ADC more robustly.
- The linear transformation between coefficients of the modified SH basis and the independent elements of the HODT.
- The comparison between the state-of-the-art high order anisotropy measures computed from the SH basis and HODT representations.
- ADC estimation on synthetic data, on a biological phantom and on a human brain dataset and ability to discriminate voxels with isotropic diffusion, single fiber and multiple fiber distributions.

Contributions from this chapter appear in [Descoteaux et al. (2005a, 2006b,a)].

ACKNOWLEDGMENTS

We would like to thank the McConnell Brain Imaging Center (BIC) of McGill University and especially, J.S.W. Campbell, P. Savadjiev, K. Siddiqi, V. V. Rymar and B.G. Pike for the biological phantom dataset. This exchange was supported by the FFRCR grant, the CRSNG Canada graduate scholarship and INRIA International internships program. Moreover, thanks to G. Sapiro, S. Lehericy and K. Ugurbil at the Center for Magnetic Resonance Research at the University of Minnesota for the high angular human brain dataset. This part was partially supported by grants NSF-INRIA Cooperative Research.

6.6 APPENDIX A: INDEPENDENT ELEMENTS OF THE HODT

Since the HODT must be totally symmetric, the problem of determining the total number of independent elements is then equivalent to finding the number of possible combinations of two integers such that their sum is less than or equal to the rank of the tensor. This is a straightforward calculation from which the number of independent elements of a rank- ℓ tensor is

$$R = \sum_{n=1}^{\ell+1} n = \frac{1}{2}(\ell + 1)(\ell + 2). \quad (6.23)$$

As in the modified SH basis, we use letter R as well because both the number of SH coefficients in the modified basis and the number of independent elements in the HODT are the same.

In order to re-write the generalized ADC profile of Eq. 6.5 in a more compact form, the multiplicity of each independent element of the HODT is needed. Consider the group of all permutations of ℓ elements, S_ℓ . The subgroup of unique permutations is equal to S_ℓ modulo the subgroup of all permutations that do not change the indices. Note that permuting only the x 's (or y 's or z 's) does not create a distinct set of indices. With a little more algebra, the subgroup of permutations fixing the indices turns out to be the product of permutations of x 's, permutations of y 's and permutations of z 's. Thus, the multiplicity μ of a component of a rank- ℓ tensor is

$$\mu = \frac{|S_\ell|}{|S_x| \cdot |S_y| \cdot |S_z|} = \frac{\ell!}{n_x! n_y! n_z!} \quad (6.24)$$

where S_x , S_y , and S_z are respectively the permutation groups of the x , y , and z indices, and n_x , n_y , and n_z are just the number of x , y , and z indices contained in the subscript of the given independent element. With this information, much of Eq. 6.5 becomes redundant, and we can reduce it to

$$D(\mathbf{g}) = \sum_{k=1}^R \mu_k T_k \prod_{p=1}^{\ell} g_{k(p)},$$

where T_k is the k th independent element of the tensor, μ_k is its corresponding multiplicity, and $g_{k(p)}$ gives the component of the gradient direction \mathbf{g} corresponding to the p th index of the k th independent element of the tensor. To illustrate this simplification, consider the vector of independent elements of a rank-2 diffusion tensor $\mathbf{T} = (T_{xx} \ T_{xy} \ T_{xz} \ T_{yy} \ T_{yz} \ T_{zz})^T$ which has multiplicity vector $\mu = (1 \ 2 \ 2 \ 1 \ 2 \ 1)^T$. We

then have

$$\begin{aligned}
D(\mathbf{g}) &= \sum_{i_1 \in \{x,y,z\}} \sum_{i_2 \in \{x,y,z\}} T_{i_1 i_2} g_{i_1} g_{i_2} \\
&= g_x^2 T_{xx} + g_y^2 T_{yy} + g_z^2 T_{zz} + g_x g_y T_{xy} + g_y g_x T_{yx} \\
&\quad + g_x g_z T_{xz} + g_z g_x T_{zx} + g_y g_z T_{yz} + g_z g_y T_{zy} \\
&= g_x^2 T_{xx} + 2g_x g_y T_{xy} + 2g_x g_z T_{xz} + g_y^2 T_{yy} + 2g_y g_z T_{yz} + g_z^2 T_{zz} \\
&= \sum_{k=1}^6 \mu_k T_k \prod_{p=1}^2 g_{k(p)}
\end{aligned}$$

which illustrates the simplification.

6.7 APPENDIX B: RELATION BETWEEN THE SPHERICAL HARMONICS TO THE HIGH ORDER TENSOR _____

Here, we aim to prove that the M matrix (Eq. 6.16) is simply a change of basis, i.e. that for all integers $n \geq 0$, the tensor polynomials of rank $2n$ and the modified set of SH up through rank $2n$ both form a basis for the same space of functions. We note that this implies that M is invertible. Before proving this, we will make some preliminary comments. Firstly, the SH of order ℓ are defined to be the restrictions of homogeneous harmonic polynomials of order ℓ to the unit sphere [Gia (2003)]. Both the standard set of SH and the modified set we introduced in this thesis are just particular choices of basis for this space. Secondly, the tensor polynomials of rank ℓ are simply the set of unique monomials of degree ℓ , i.e. $\{x^a y^b z^c | a, b, c \geq 0, a + b + c = \ell\}$ which we also restrict to the unit sphere. Finally, both of the above sets (modified SH of order $\leq 2n$ and tensor polynomials of rank $2n$) are linearly independent and have the same number of elements. Therefore, it suffices to show that all the elements in one set can be expressed as linear combinations of elements in the other set to prove that both sets form a basis for the same space.

THEOREM .2. *For all integers $n \geq 0$, the tensor polynomials of rank $2n$, i.e. $\{x^a y^b z^c | a, b, c \geq 0, a + b + c = \ell\}$, restricted to the sphere, and the set of modified spherical harmonics of order $\leq 2n$ form a basis for the same space of functions.*

Proof: Our approach will be to show that every even order homogeneous harmonic polynomial restricted to the unit sphere of order $\leq 2n$ can be expressed as a linear combination of the tensor polynomials of rank $2n$. As discussed above, this result is equivalent to the statement of the theorem. We proceed by induction. Note that we will use \propto to denote “is proportional to”.

Base case: For $n = 0$, the set of homogeneous harmonic polynomials of order 0 restricted to the unit sphere is simply the set of constant functions, and we have that

$$\{x^a y^b z^c | a, b, c \geq 0, a + b + c = 0\} = x^0 y^0 z^0 \propto 1.$$

Therefore, any homogeneous harmonic polynomial of order 0 restricted to the unit sphere is simply a linear scaling of the unique tensor polynomial of order 0.

Inductive step: Assume that all homogeneous harmonic polynomials of order $\leq 2n - 2$ can be expressed as linear combinations of tensor polynomials of rank $2n - 2$. Since we are restricted to the sphere, we can multiply each of these formulas by $x^2 + y^2 + z^2 = 1$ to express each of the homogeneous harmonic polynomials of order $\leq 2n - 2$ as linear combinations of tensor polynomials of rank $2n$. For example, the homogeneous harmonic polynomial of order 2, xz , is itself a tensor polynomial of rank 2. However, it can also be written as a linear combination of rank 4 tensor polynomials in the following way: $xz(x^2 + y^2 + z^2) = x^3 z + xy^2 z + xz^3$. Furthermore, every homogeneous harmonic polynomial of order $2n$ can be expressed as a linear combination of tensor polynomials of rank $2n$ simply because the set of unique monomials of degree $2n$ spans the set of polynomials of degree $2n$. This completes the proof.

Here is a brief illustrative example. Below we express each of the elements of our modified SH basis of order ≤ 2 (Figure 6.3) in terms of linear combinations of tensor polynomials of rank 2. (Restricting to the unit sphere gives, as in standard spherical coordinates, the relations: $x = \cos \phi \sin \theta$, $y = \sin \phi \sin \theta$, $z = \cos \theta$ and $1 = x^2 + y^2 + z^2$).

$$\begin{aligned} Y_1 &\propto 1 && \propto x^2 + y^2 + z^2 \\ Y_2 &\propto \sin^2 \theta \cos(2\phi) \\ &\propto \sin^2 \theta (\cos^2 \phi - \sin^2 \phi) && \propto x^2 - y^2 \\ Y_3 &\propto \sin \theta \cos \theta \cos \phi && \propto xz \\ Y_4 &\propto 3 \cos^2 \theta - 1 \\ &\propto 3z^2 - (x^2 + y^2 + z^2) && \propto 2z^2 - x^2 - y^2 \\ Y_5 &\propto \sin \theta \cos \theta \sin \phi && \propto yz \\ Y_6 &\propto \sin^2 \theta \sin(2\phi) \\ &\propto \sin^2 \theta (2 \sin \phi \cos \phi) && \propto xy \end{aligned}$$

With linear combination of the Y_1 , Y_2 and Y_4 terms, one can obtain the monomials x^2 , y^2 and z^2 . Hence, the modified basis of order 2 spans the set of degree two monomials.

This result explains why fitting with the high order tensor formulation and fitting with the SH formulation are equivalent. This is seen in the similarity of the $\lambda = 0$ fit and the linear regression HODT fit in Table 6.2.

ANALYTICAL Q-BALL IMAGING

Beautiful Evidence is about the theory and practice of analytical design

– Edward Tufte

Contents

7.1 Introduction	107
7.2 Q-Ball Imaging	108
7.3 Analytical Q-Ball Imaging	109
7.4 Validation and Comparison of QBI methods	114
7.4.1 Numerical QBI Implementation	114
7.4.2 Synthetic Data Generation	114
7.4.3 Computational Complexity Analysis	115
7.4.4 Effect of Spherical Sampling Density	115
7.4.5 Robustness to Noise	115
7.4.6 ODF Shape Comparison	116
7.4.7 Fiber Detection and Angular Resolution	116
7.4.8 Power Spectrum of the Spherical Harmonics Representation	117
7.5 Results	117
7.5.1 Running Time Comparison	117
7.5.2 ODF Shape Comparison and Robustness to Noise	118
7.5.3 ODF Reconstruction for Different Sampling Schemes	120
7.5.4 Fiber Detection and Angular Resolution	121
7.5.5 Power Spectrum of the Spherical Harmonic Representation	127
7.5.6 Biological Phantom Results	129
7.5.7 Human Brain HARDI Data Results	131
7.6 Discussion	134
7.6.1 Contributions of this chapter:	136
7.7 Appendix A: Funk-Radon transform and the diffusion ODF	138
7.8 Appendix B: Rigorous Proof of the Analytical QBI Solution	140
7.9 Appendix C: Exact ODF From the Multiple-Tensor Model	141

OVERVIEW

Given our new regularized method to obtain the HARDI signal, can it be used to do q-ball imaging (QBI)? This chapter proposes a new analytical QBI solution using our regularized spherical harmonics estimation. In particular, can the Funk-Radon transform be solved in a single step with spherical harmonics? First, we develop a regularized, fast and robust analytical solution for the QBI reconstruction. To do so, a new corollary of the Funk-Hecke theorem is proved to obtain an elegant mathematical simplification of the Funk-Radon transform and to obtain the analytical QBI solution in a single step. Second, we validate and compare the analytical QBI solution with the numerical QBI solution on synthetic simulations, on a rat biological phantom and on real datasets. Finally, a discussion on its benefits over the other state-of-the-art approaches is done.

Keywords q-ball imaging (QBI), orientation distribution function (ODF), Laplace-Beltrami (LB) operator, regularization, Funk Radon transform (FRT), spherical harmonics (SH), Funk-Hecke theorem

Contributions of this chapter:

- Regularized, fast and robust analytical QBI solution.
- Proof of a new corollary of the Funk-Hecke theorem.
- Comparison of the Laplace-Beltrami regularization of the signal with the state-of-the-art numerical and analytical QBI methods using no regularization and Tikhonov regularization.
- Extensive study of the fiber detection, robustness, compression and reconstruction properties of the analytical QBI solution.
- Quantitative validation performed against ground truth from synthetic data and against real data from a biological phantom and a human brain dataset.

Organization of this chapter:

The chapter is organized as follows. We motivate the QBI problem in Section 7.1 and develop the theory of QBI in Section 7.2. Then, the analytical QBI solution is developed in Section 7.3. Next, we validate and compare QBI methods in Section 7.4 and show the fiber detection, robustness and compression properties of the ODF reconstruction in Section 7.5. Finally, we conclude with a discussion of the results and our contributions in Section 7.6.

7.1 INTRODUCTION

As already seen many times in this thesis, DTI is limited in regions of low anisotropy where there are multiple crossing fibers. Chapter 4 exposed the recent works done to generalize the existing DTI model [Basser et al. (1994a)] with new HARDI acquisition techniques such as DSI [Wedeen et al. (2000)] and single shell HARDI [Tuch et al. (1999); Tuch (2002)]. There are currently two classes of high order processing methods for these high resolution acquisition techniques. The first is based on the apparent diffusion coefficient (ADC) modeling and the second is based on the reconstruction of the diffusion probability density function (PDF) of the average spin displacement of water molecules or variants of this function. ADC modeling was seen in the previous chapter. This chapter is on the second class of HARDI processing methods. In particular, the chapter details the state-of-the-art QBI reconstruction solution and our contributions on this problem.

As seen in Chapter 4, the existing functions in the literature that approximate the diffusion PDF or variants of it are the diffusion orientation distribution function (ODF) [Tuch (2002)], the Persistent Angular Structure (PAS) [Jansons and Alexander (2003)] of the diffusion PDF, the fiber orientation distribution (FOD) [Tournier et al. (2004); Alexander (2005b)] and the diffusion orientation transform (DOT) [Ozarslan et al. (2005a)]. For all these high angular resolution functions, the important property is that their maxima agree with the underlying fiber distribution. However, these methods are based on numerical methods, lacking a straightforward regularization process and failing to take into account the useful tools for both estimation and regularization that have been developed for estimating spherical functions.

The backbone mathematical tool of this chapter is again the spherical harmonics (SH). Other groups have also converged to this tool. The SH seem to be a natural way to decompose signals that live on the sphere, just as the Fourier transform is the widespread tool used to decompose images. Hence, it is not surprising that [Anderson (2005)], [Hess et al. (2006)] and our group [Descoteaux et al. (2005b, 2006c, 2007a)] have recently developed separately and in parallel an analytical solution for the ODF reconstruction in QBI. Despite the fact that the three analytical solutions are similar, the regularized estimation part, the derivation, the experimental results and the validation phase are quite different. In this chapter, we carefully present the different solutions and put up front the differences in derivations and contributions between our proposed method and methods of [Anderson (2005); Hess et al. (2006)].

More precisely, our solution is obtained by modeling the signal with high order SH series using a Laplace-Beltrami regularization method developed for the ADC profile estimation in the previous chapter. This leads to an elegant mathematical simplification of the Funk-Radon transform (FRT) that approximates the ODF. We prove a new corollary of the Funk-Hecke theorem to obtain this simplification. We then obtain a fast algorithm for the extraction of a robust regularized ODF, which offers advan-

tages over previous numerical approaches. Another contribution is the quantitative validation on synthetic data and on real data obtained from a biological rat phantom and from a human brain dataset. Overall, the contributions are three-fold: 1) Our ODF estimation is up to 15 times faster than Tuch’s numerical method. 2) Our ODF estimation is regularized with the Laplace-Beltrami operator, which is theoretically and practically better than Tikhonov regularization [Hess et al. (2006)] as well as more robust to noise. 3) At the cost of slightly reducing angular resolution, our ODF estimation reduces errors and improves fiber detection while reducing angular error in the ODF maxima detected.

7.2 Q-BALL IMAGING

[Tuch (2002, 2004)] showed that the diffusion ODF could be estimated directly from the raw HARDI measurements on a single sphere of q-space without computing the full diffusion PDF. The basic assumption of QBI is that angular information is enough to recover fiber orientation distributions (forgetting about radial information). Hence, QBI is a modality which takes advantage of the fact that significantly less information is required to construct an angular function in real space than is required to construct a volume function, as in DSI. Specifically, QBI seeks to reconstruct the diffusion ODF, a function defined in Eq. 4.13, which is the radial projection of the diffusion PDF. Recall that it is given by

$$\Psi(\mathbf{u}) = \int_0^\infty P(\alpha\mathbf{u})d\alpha, \quad (7.1)$$

where \mathbf{u} is restricted to be a unit vector. Thus, the ODF is a function on the unit sphere describing the probability averaged over the voxel that a particle will diffuse into any given solid angle.

To compute the ODF, the QBI modality uses the Funk-Radon transform (FRT) \mathcal{G} , a transformation from the unit sphere to itself. It was illustrated in Figure 4.14. In order to find the Funk-Radon transformed value of the signal on the sphere at a given point \mathbf{u} , one needs to first find the plane through the origin with normal vector \mathbf{u} and then compute the one dimensional integral over the intersection of that plane with the function on the original sphere. Intuitively, to find the new value at an arbitrarily defined “pole”, one integrates the spherical function f over the corresponding “equator” or great circle. This can be written explicitly as

$$\mathcal{G}[f(\mathbf{w})](\mathbf{u}) = \int \delta(\mathbf{u}^T \mathbf{w})f(\mathbf{w})d\mathbf{w} \quad (7.2)$$

where \mathbf{u} and \mathbf{w} are constrained to be unit vectors. To see why this spherical transformation is close to the true diffusion ODF of Eq. 7.1, [(Tuch, 2004, Appendix A)] proves that the FRT of Eq. 7.2 can be written as

$$\mathcal{G}_{q'}[S(\mathbf{q})](\mathbf{u}) = 2\pi q' \int P(r, \theta, z)J_0(2\pi q' r)rdrd\theta dz, \quad (7.3)$$

where q' is the radius of the acquisition shell in q -space and J_0 the zeroth-order Bessel function. In Appendix 7.7, we reproduce a more detailed proof and introduce the important Fourier analysis tools that clarify the relation between the FRT and the true diffusion ODF. From this Appendix, if we assume, without loss of generality, that the evaluation direction \mathbf{u} is along the z -axis, Eq. 7.1 can also be written as

$$\Psi(\mathbf{u}) = \int P(r, \theta, z) \delta(r) \delta(\theta) r dr d\theta dz. \quad (7.4)$$

Hence, when comparing Eq. 7.3 and Eq. 7.4, it is clear that the FRT is a smoothed version of the true diffusion ODF. In fact, the higher q' (higher is the b -value), the closer the FRT approximation is to the exact ODF because the zeroth-order Bessel function J_0 gets sharper and approaches a Dirac delta function (see Appendix 7.7 and Figure 7.14). Fortunately, the values of b required to satisfy this condition are relatively small ($b = 4000 \text{ s/mm}^2$ used in [Tuch (2002)]) compared to the maximal values of b required for the DSI technique ($b = 20000 \text{ s/mm}^2$ used in [Wedeen et al. (2000)] and $b = 60000 \text{ s/mm}^2$ used in [Ozarslan and Mareci (2003a)]).

In practice, the FRT can be implemented by a matrix multiplication. However, computing this matrix involves several non-trivial numerical computations. In particular, a regridding and an interpolation of the spherical input data is needed to compute the equator points in the FRT integral of Eq. 7.2 since many points outside the actual measurements are required. This interpolation requires a good a priori basis function and many sampling directions are required to obtain a good QBI reconstruction. In fact, the implementation of [Tuch (2004)] uses a fivefold tessellated icosahedron (252 samples on the sphere) and a regridding of the signal onto equators around vertices of a fivefold tessellated dodecahedron ($48 \times 755 = 36240$ points). Hence, it is a method less intensive than DSI but that still requires non-trivial computational steps.

In this chapter, we significantly reduce the computation burden of QBI and improve its robustness to noise. We apply our technique of the previous chapter for the approximation of functions on the sphere from noisy sparse data in order to obtain a robust signal estimation with a regularization procedure. In the previous chapter, it was shown to be very useful to characterize isotropic, 1- and multi-fiber fiber distributions. We will now show that the use of the spherical harmonics basis can greatly simplify the FRT integral.

7.3 ANALYTICAL Q-BALL IMAGING ---

As mentioned in the introduction, we are proposing a similar analytical QBI solution for the ODF reconstruction as [Anderson (2005)] and [Hess et al. (2006)]. However, despite the similarity of the analytical solution, the derivations are quite different and use other properties of the spherical harmonic (SH) basis that are worth reviewing.

First, [(Anderson, 2005, Appendix B)] proposes an analytical solution using spherical harmonics. It is not the focus of the paper and this is the reason why the full potential of the solution was not studied as Hess et al did and as we do in this chapter. Nonetheless, the derivation is simple and different from that of both Hess et al and ours. The idea is that the signal on the sphere can be represented in terms of the standard polar angles (θ, ϕ) or in terms of coordinates in a rotated frame. Using the Wigner rotation matrix, there is a simple relation between SHs evaluated at (θ, ϕ) and SHs evaluated at the same physical angle but expressed in the rotated frame. This is used to evaluate and simplify the FRT by expressing the integral in a carefully chosen rotated frame.

Second, the solution of [Hess et al. (2006)] is based on the analytical solution to great circle integrals over SHs. This is based on a previous work by [Backus (1964)] for an application to geophysical data. Two important properties of the SH are exploited in this proof: 1) any rotated SH can be uniquely expressed as a linear combination of SH of the same degree and 2) SH satisfy the addition theorem. The two properties allows them to rewrite the Funk-Radon integral and solve it analytically.

In our case, we use the 3D Funk-Hecke theorem of Section 5.5 to analytically evaluate integrals of functions on the sphere. The SH formulation of Section 6.2.3 is crucial for our analytical simplification of QBI and the Laplace-Beltrami regularization gives the robustness to the solution.

Signal Approximation with the Spherical Harmonics

For the analytical solution, the signal \mathbf{S} ($S(b) / S_0$) is described in a SH series with coefficients \mathbf{C} using the regularized solution of Eq. 6.13¹. As shown in the previous chapter, the Laplace-Beltrami (LB) regularization approach is used to eliminate unnecessary higher order terms in the SH approximation. In [Hess et al. (2006)], classical Tikhonov regularization is used to obtain the c_j coefficients at high orders ($\ell > 4$) and a simple un-regularized solution for $\ell = 4$. It is a simple technique used to numerically better condition the matrix involved in the pseudo-inverse of Eq. 6.13. The solution is expressed in the same form as Eq. 6.13 but the Laplace-Beltrami matrix \mathbf{L} is replaced with the identity matrix \mathbf{I} . The regularization is thus uniform, i.e. lower and higher order coefficients are weighted in the same way, which is less desirable than a smoothing that minimizes perturbing effects occurring mainly at higher order harmonics. Moreover, the underlying assumption of Tikhonov regularization is that the space where the data lives is a flat manifold. This is a rough approximation in the ODF reconstruction problem since the data actually lives on a sphere, for which the right and appropriate tool to use is the Laplace-Beltrami operator.

Finally, other works propose some post-processing of the ODF to reduce errors involved in high frequency harmonics. In [Tournier et al. (2004)], higher order

¹The ADC vector \mathbf{X} in Eq. 6.13 is replaced with the raw HARDI signal vector $S(b)/S_0$, as in the un-regularized solution of Eq. 5.29

terms of the SH fiber ODF reconstruction are attenuated by a heuristic amount. In [Tuch (2004)], an isotropic spherical smoothing filter of a certain width is applied to the reconstructed ODFs. More recently, as mentioned in the previous chapter, [Sakaie and Lowe (2007)] proposed a gradient-based regularization which is similar in spirit as our regularization.

Funk-Radon Transform Using Spherical Harmonics

We now prove a corollary to the Funk-Hecke theorem needed to solve the Funk-Radon transform (FRT) integral. Recall that the FRT of the signal in a unit direction \mathbf{u} is the integral over the great circle perpendicular to \mathbf{u} , as stated in Eq. 7.2. Hence, replacing the signal by the SH series approximations, Eq. 5.30 can be written as

$$\begin{aligned}
\mathcal{G}[S/S_0](\mathbf{u}) &= \frac{1}{S_0} \int_{|\mathbf{w}|=1} \delta(\mathbf{u}^T \mathbf{w}) S(\mathbf{w}) d\mathbf{w} \\
&= \frac{1}{S_0} \int_{|\mathbf{w}|=1} \delta(\mathbf{u}^T \mathbf{w}) \sum_{j=1}^R c_j Y_j(\mathbf{w}) d\mathbf{w} \\
&= \frac{1}{S_0} \sum_{j=1}^R c_j \underbrace{\int_{|\mathbf{w}|=1} \delta(\mathbf{u}^T \mathbf{w}) Y_j(\mathbf{w}) d\mathbf{w}}_{I_j}
\end{aligned} \tag{7.5}$$

Note that if the Dirac delta function δ was continuous on the interval $[-1, 1]$, I_j could be directly evaluated using the Funk-Hecke formula of Eq. 5.31. However, $\delta(t)$ is discontinuous at zero. To overcome this problem, we approximate the delta function with a delta sequence $\delta_n(x)$ and take the limit as n goes to infinity. We only need the existence of such a sequence and for example, the Gaussian of decreasing variance $1/n^2$, given by $\delta_n(x) = (n/\sqrt{\pi}) \exp(-n^2 x^2)$, is a well-known delta sequence. As n goes to infinity, the variance of this Gaussian tends to zero, so that this sequence of functions satisfies the defining property of the delta sequence, i.e. that

$$\lim_{n \rightarrow \infty} \int_{-\infty}^{\infty} \delta_n(x) f(x) dx = f(0). \tag{7.6}$$

It is a straightforward corollary to show that this definition is valid over the interval $[-1, 1]$ if δ_n is continuous sequence of functions. This is the case for the Gaussian sequence for all n . Hence, we can evaluate I_j as

$$\begin{aligned}
I_j(\mathbf{u}) &= \int_{|\mathbf{w}|=1} \delta(\mathbf{u}^T \mathbf{w}) Y_j(\mathbf{w}) d\mathbf{w} \\
&= \lim_{n \rightarrow \infty} \int_{|\mathbf{w}|=1} \delta_n(\mathbf{u}^T \mathbf{w}) Y_j(\mathbf{w}) d\mathbf{w} && \text{(def'n of delta sequences in Eq. 7.6)} \\
&= 2\pi \left(\lim_{n \rightarrow \infty} \int_{-1}^1 \delta_n(t) P_{\ell(j)}(t) dt \right) Y_j(\mathbf{u}) && \text{(Funk-Hecke theorem Eq. 5.31)} \\
&= 2\pi P_{\ell(j)}(0) Y_j(\mathbf{u}),
\end{aligned} \tag{7.7}$$

where $\ell(j)$ is the order associated with the j th element of the SH basis, i.e. for $j = \{1, 2, 3, 4, 5, 6, 7, \dots\}$, $\ell(j) = \{0, 2, 2, 2, 2, 2, 4, \dots\}$ as before. The rigorous proof can be found in Appendix 7.8. The strength of this derivation is that it greatly simplifies the FRT integral in one simple step. We have thus proved the following corollary of the Funk-Hecke theorem in 3D:

COROLLARY .1. Corollary of the Funk-Hecke Theorem: *Let $\delta(t)$ be the Dirac delta function and Y_ℓ any spherical harmonic of order ℓ . Then, given a unit vector \mathbf{u}*

$$\int_{|\mathbf{w}|=1} \delta(\mathbf{u}^T \mathbf{w}) Y_\ell(\mathbf{w}) d\mathbf{w} = 2\pi P_\ell(0) Y_\ell(\mathbf{u}), \quad (7.8)$$

where $P_\ell(0)$ the Legendre polynomial of degree ℓ evaluated at 0,

$$P_\ell(0) = \begin{cases} 0 & \ell \text{ odd} \\ (-1)^{\ell/2} \frac{1 \cdot 3 \cdot 5 \cdots (\ell - 1)}{2 \cdot 4 \cdot 6 \cdots \ell} & \ell \text{ even} \end{cases} \quad (7.9)$$

Note that had we used a standard SH basis, the odd coefficients of the ODF estimation would have vanished. Therefore, Funk-Hecke theorem may be useful for anyone working with SH and seeking solutions to integrals over the sphere. Referring back to Eq. 7.5, the FRT of a function given in terms of our modified SH series in a given unit vector direction \mathbf{u} is simply given by

$$\mathcal{G}[S](\mathbf{u}) = \sum_{j=1}^R \frac{2\pi P_{\ell(j)}(0)}{S_0} c_j Y_j(\mathbf{u}). \quad (7.10)$$

Thus, the SHs are eigenfunctions of the FRT with eigenvalues depending only on the order ℓ of the SH series. When the signal S is parameterized by the vector \mathbf{C} of SH coefficients, i.e. $\mathbf{S} = \mathbf{B}\mathbf{C}$ as before, the ODF reconstruction in terms of SH coefficients, denoted by the $R \times 1$ vector \mathbf{C}' , is simply a diagonal linear transformation given by

$$\mathbf{C}' = \begin{pmatrix} \ddots & & & & \\ & 2\pi \frac{(-1)^{\ell(j)/2} 1 \cdot 3 \cdot 5 \cdots (\ell(j) - 1)}{S_0 2 \cdot 4 \cdot 6 \cdots \ell(j)} & & & \\ & & \ddots & & \\ & & & \ddots & \\ & & & & \ddots \end{pmatrix} \begin{pmatrix} \vdots \\ c_j \\ \vdots \end{pmatrix} \quad (7.11)$$

Note that if one is interested in a final estimated ODF on the sphere Ψ , it can be obtained using the simple matrix multiplication

$$\Psi = \mathbf{B}\mathbf{C}', \quad (7.12)$$

where \mathbf{B} is given by Eq. 5.24 as before. The detailed steps of the analytical QBI reconstruction algorithm are presented in Table 7.1.

Hence, by using the SH for the approximation of the signal attenuation function at a given radius in q -space, the QBI can be solved analytically, as also shown

Table 7.1: Summary of the regularized, fast and robust analytical QBI algorithm.

Input:

$\mathbf{V} : X \times Y \times Z \times N$ diffusion weighted MRI volume

$\mathbf{S}_{xyz} : N \times 1$ diffusion weighted signal vector at voxel (x, y, z)

$\Theta : 2 \times N$ matrix of gradient encoding directions in spherical coordinates, i.e. $\Theta_i = (\theta_i, \phi_i)$ for each $i \in \{1, N\}$

ℓ : order of SH basis $\implies R = (1/2)(\ell + 1)(\ell + 2)$

λ : regularization parameter

Output:

$\mathbf{C}'_{xyz} : R \times 1$ diffusion ODF vector in SH coefficients at voxel (x, y, z)

Algorithm:

$\mathbf{B} := \text{Eq. 5.24}$: Construct $N \times R$ matrix of SH basis elements using Θ and modified SH basis $\mathbf{Y} = (Y_1, \dots, Y_R)^T$ defined by Eq. 5.21.

$\ell(j)$: order associated with element \mathbf{Y}_j

For $j = \{1, 2, 3, 4, 5, 6, 7, 8, \dots\}$, $\ell(j) = \{0, 2, 2, 2, 2, 2, 4, 4, \dots\}$.

$\mathbf{L} = \begin{pmatrix} \ddots & & & & & \\ & \ell(j)^2(\ell(j) + 1)^2 & & & & \\ & & \ddots & & & \\ & & & \ddots & & \\ & & & & \ddots & \\ & & & & & \ddots \end{pmatrix} : R \times R$ LB smoothing matrix

$\mathbf{P} = \begin{pmatrix} \ddots & & & & & \\ & \frac{2\pi P_{\ell(j)}(0)}{S_0} & & & & \\ & & \ddots & & & \\ & & & \ddots & & \\ & & & & \ddots & \\ & & & & & \ddots \end{pmatrix} : R \times R$ FRT matrix ($P_{\ell(j)}(0)$ from Eq. 7.9)

$\mathbf{Tr} = \mathbf{P}(\mathbf{B}^T \mathbf{B} + \lambda \mathbf{L})^{-1} \mathbf{B}^T$: $R \times N$ signal to ODF transform matrix

For each $(x, y, z) \in \mathbf{V}$

$\mathbf{C}'_{xyz} = \mathbf{Tr} \cdot \mathbf{S}_{xyz}$: Compute SH coefficients of the ODF in one step

in [Anderson (2005); Hess et al. (2006)]. An important contribution in favor of our approach is that this solution can be obtained while imposing a well-defined regularization criterion. The accuracy of the modified SH series approximation with the Laplace-Beltrami smoothing was established in the previous chapter for sparse measurements on the sphere. The analytical ODF estimation method offers the advantages that the discrete interpolation over many equators is eliminated and the solution for all directions is obtained in a single step. We now validate our new analytical QBI method and analyze the computational complexity gain. We also point out the different and complementary results obtained in [Hess et al. (2006)].

7.4 VALIDATION AND COMPARISON OF QBI METHODS —

7.4.1 Numerical QBI Implementation

The implementation of [Tuch (2004)] involves several numerical computations such as a regridding to find points outside the actual measurements required to compute the discrete points on each great circle. Our implementation of Tuch’s numerical QBI is adapted from [Campbell et al. (2005)] QBI computation which is successfully used for multiple fiber characterization and tracking. Given N points spaced approximately uniformly on the surface of the sphere, the interpolation on each unit great circle can be done at every $(1/2)\sqrt{2\pi/N}$ radians to take full advantage of the N discrete measurements. This gives $k = \sqrt{8\pi N}$ points per equator. We also set the width parameter (angular width) of the spherical Gaussian interpolation kernel automatically to three times the angle between equator points, i.e. $\sigma = (3/2)\sqrt{2\pi/N}$.² This heuristic choice gives a good trade-off between the accuracy and stability in our experiments.

7.4.2 Synthetic Data Generation

We generate synthetic data using the multi-tensor model of Appendix A.1, which leads to an analytical computation of the exact ODF. This exact ODF is derived in Appendix 7.9. For a given b -value, noise level and encoding direction i , we generate the diffusion-weighted signal S for $N = 81$ or $N = 321$ gradient encoding directions on the hemisphere (for 3rd or 7th order tessellation of the icosahedron respectively). As in [Hess et al. (2006)] we vary SNR values between 5 and 50. This SNR range covers expected low to high quality of real in vivo HARDI data.

²In Tuch (2004), k is set to 48 and σ to 5° although it is argued that k and σ can be selected for a desired level of numerical precision.

7.4.3 Computational Complexity Analysis

We refer to Table 7.1 for the analytical QBI and to [(Tuch, 2004, Tbl.1)] for the numerical QBI. Assuming that the input diffusion MRI volume is of size $X \times Y \times Z \times N$, we have a diffusion signal vector of $N \times 1$ at each voxel, where N is the number of gradient directions taken. We let k be the number of points on each equator over which the numerical Funk-Radon integral is computed. In our technique, we have defined $R = (1/2)(\ell + 1)(\ell + 2)$ to be the number of elements in the SH basis. It is straightforward to see that the analytical ODF reconstruction is $O(XYZNR)$ because of the $O(NR)$ matrix multiplication at every voxel, while the original ODF reconstruction is $O(XYZNk)$ because of the integration of k equator points for each sampling direction N at every voxel. Therefore, the difference in computational complexity between the two methods can be explained by the difference between R and k , where R is generally smaller than k . For example, at orders $\ell = 4, 6, 8$, for our method $R = 15, 28, 45$ respectively; in [Tuch (2004)] k is set to be 48.

7.4.4 Effect of Spherical Sampling Density

To study the effect of sampling density, we vary the number of gradient encoding directions along which we generate the signal and we also vary the SNR between 30, 20 and 10. We fix $b = 3000 \text{ s/mm}^2$, estimation order $\ell = 8$, $\lambda = 0.006$ and generate signal for fibers crossing at 90° with equal volume fractions.

7.4.5 Robustness to Noise

We perform two simulations to evaluate the ODF estimation. In the first simulation, SNR was fixed to 35 while varying the b -value between 500 and 12000 s/mm^2 for estimation order $\ell = 4$ and $\ell = 8$. In the second simulation, we fix the b -value to $b = 3000 \text{ s/mm}^2$ and estimation order to $\ell = 8$ while varying the SNR between 5 and 50. In both tests, we randomly choose the number of fibers n per voxel between 1, 2 and 3 as described in Appendix A.1. The optimal regularization λ parameter can be obtained from the L-curve numerical method [Hansen (2001)] or generalized cross validation [Wahba (1990); Sakaie and Lowe (2007)]. To avoid having to compute the optimal λ for each HARDI profile at every iteration, we set $\lambda = 0.006$ for our reconstruction, a value shown to provide good separation of 1-fiber from 2-fiber distributions over a range of SNR and b -values in the previous chapter. We apply the transformation given in Eq. 7.11 to obtain the estimated SH coefficients of the ODF and use Eq. 7.12 to obtain the corresponding discrete function on the sphere for the N sampling directions.

7.4.6 ODF Shape Comparison

Letting Ψ represent the exact ODF and Ψ' the estimated ODF, we compute the average Euclidean squared error between Ψ and Ψ' over all tests and N samplings of the sphere, i.e.

$$[\Psi, \Psi'] = \frac{1}{N} \sum_{i=1}^N (\Psi'(\theta_i, \phi_i) - \Psi(\theta_i, \phi_i))^2. \quad (7.13)$$

This method is used to compare ODF shapes throughout the *Results* section, including experiments using synthetic data, data from the biological phantom and human brain data. However, other appropriate distances such as Kullback-Leibler divergence [Tuch (2004)] could also be considered.

7.4.7 Fiber Detection and Angular Resolution

To evaluate fiber detection differences between Laplace-Beltrami and Tikhonov regularization, we test on synthetic HARDI data generated using a 3rd order tessellation of the icosahedron that gives 81 samplings on the hemisphere, a SNR of 10 and 2 orthogonal fibers. We vary estimation order ℓ and use two b -values of 3000 s/mm² and 1000 s/mm². 1) We generate 1000 such HARDI data separately, 2) we estimate ODFs with/without Laplace-Beltrami and Tikhonov regularization, 3) we count the number of times we correctly detect 2 ODF maxima and 4) we report the percentage and average angular error \pm standard deviation in degrees over all 1000 trials.

We also perform a numerical experiment to evaluate angular resolution limitations of the ODF reconstruction with/without Laplace-Beltrami and Tikhonov regularization. We generate noise-free synthetic HARDI profiles for 2 fibers, for b -values of 3000 s/mm² and 1000 s/mm² and for spherical sampling densities $N = 81$ and $N = 321$. Then, we vary the crossing angle between fibers to determine the critical angle at which only a single maxima starts to be detected instead of two. We report this critical angle as the angular resolution of the estimation.

Maxima Extraction To extract the fiber ODF maxima, it is generally assumed that they are simply given by the local maxima of the normalized ODF ($[0,1]$), where the function exceeds a certain threshold (here, we use 0.5). In practice, we project the ODF onto the sphere tessellated with a fine mesh. We use a 16th order tessellation of the icosahedron, which gives 1281 sample directions on the sphere.³ Then, we implement a finite difference method on the mesh. If a mesh point is above all its neighbors and if this point has a normalized ODF value above 0.5, we keep the mesh point direction as a maxima. This thresholding avoids selecting small peaks that may appear due to noise. Other more complex methods exist to extract the maxima such as the method presented in [Hlawitschka and Scheuermann (2005)] or spherical Newton’s method [Tournier et al. (2004)], Powell’s method [Jansons and Alexander

³The angle between each spherical grid point is approximately 4°.

(2003)] or sequential quadratic programming [Sakaie and Lowe (2007)]. However, one must be careful with these more complex methods to avoid getting trapped in small local maxima. We prefer our simple finite difference scheme because it is very efficient and robust in practice.

7.4.8 Power Spectrum of the Spherical Harmonics Representation

[Hess et al. (2006)] also analyze the power spectrum of the SH representation up to order $\ell = 10$ for signal with 1, 2 and 3 fibers crossing. The contribution of each angular frequency of the power spectrum is computed with

$$E_\ell = \frac{1}{E_T} \sum_{\{j:\ell(j)=\ell\}} |c'_j|^2, \quad (7.14)$$

where c'_j are the coefficients of describing the ODF (Eq. 7.11), $\ell(j)$ is the order associated with the j th coefficient as before and E_T denotes the total power of the 10th-order harmonic representation for all coefficients of order $\ell \geq 2$. The 0th-order harmonic is excluded as it is constant and does not contribute to the anisotropy of the ODF. We have reproduced the power spectrum study for different b -values and we study the difference between the power spectrum of the regularized ($\lambda = 0.006$) and un-regularized ($\lambda = 0$) analytical QBI reconstructions. The power spectrum was computed and averaged for 1000 simulated noisy HARDI signal of 1 fiber, 2 fibers and 3 fibers crossing respectively with SNR of 35 and sampling density of $N = 81$. As before, the separation angle for crossings was random between 45° and 90° and fibers had random fiber fractions as described in Appendix A.1.

7.5 RESULTS

We show five contributions of the regularized analytical QBI method; 1) it is up to 15 times faster than Tuch’s numerical method, 2) it requires fewer DWI measurements than numerical QBI to perform a stable ODF reconstruction, 3) it is robust to noise, 4) it improves accuracy in ODF maxima detection at the cost of slightly reducing angular resolution and 5) it recovers fiber crossings from synthetic data, from a biological phantom and from real human brain data.

7.5.1 Running Time Comparison

Table 7.2 shows that analytical QBI is up to 15 times faster than Tuch’s numerical QBI in practice. Computation was performed on a Dell single processor, 3.4 GHz, 2 GB RAM machine. Given N samples on the sphere and a SH basis of order ℓ , the theoretical speed-up factor for order $\ell = 4, 6, 8$ is approximately 3, 2, 1 respectively. However, in practice the running time of the analytical QBI is nearly 15 times faster. The factor of 5 gap between the theoretical and experimental speed-up is because of the constant time operation hidden and not accounted for in the “big O” analysis


dataset size (voxels)	Phantom (A.2) 35 x 17 x 35 x 90	Brain (A.3.1) 128 x 128 x 63 x 99	Brain (A.3.3) 93 x 116 x 93 x 61
Numerical QBI	0:13.59	13:27.12	5:43.86
Analytical QBI			
$\ell = 4$	0:01.27	0:40.42	0:27.07
$\ell = 6$	0:01.61	0:53.06	0:35.76
$\ell = 8$	0:02.09	1:49.37	0:58.38

Table 7.2: Analytical QBI is up to 15 times faster than Tuch’s numerical QBI in practice. Computation experiments are performed on a Dell single processor, 3.4 GHz, 2 GB RAM machine.


of the running time in theory. This is mainly due to the interpolation kernel width of Tuch’s approach which is not considered in the complexity analysis and which adds a constant number of operations (4 to 7 in practice) at every equator point. The reading/writing of 4D volumes with 4th dimension N instead of R is also slower. Thus, as mentioned in [Hess et al. (2006)], there are interesting potential data compression applications offered by the SH basis, since only a few harmonic coefficients need to be stored per voxel.

7.5.2 ODF Shape Comparison and Robustness to Noise

Figure 7.1 shows that the Laplace-Beltrami regularization reduces ODF estimation errors. First, Figure 7.1a shows that the estimation is precise. As expected, we observe that for optimal b -value, the error is less than 1%. Note also that the best results are not for the highest b -values because in this case, the diffusion-weighted signal decreases sharply in SNR. This is because the signal intensity of the Gaussian-like profile of the signal varies more rapidly with ADC. At the extremity of very high b -values, most of the signal is lost. One would need to choose a much higher λ to prevent the regularized curve from approaching the un-regularized curve. It is also expected that for low b -values, the accuracy of the estimation is reduced. This is mostly due to the Bessel function averaging effect explained in Section 7.2 and Appendix 7.7. The signal is too smooth and there is very small contrast. Regularization is unnecessary in these cases. It is also interesting to note that a lower estimation order ($\ell = 4$) can outperform the regularized and un-regularized estimation of order $\ell = 8$ for low b -values ($b < 1500$ s/mm²). the higher angular resolution afforded by higher estimation order is not possible. Finally, lowest error is observed in a plateau of relatively high b -values between 2000 and 6000 s/mm², which agrees with reported results in the literature [Tuch (2002); Tournier et al. (2004)]. How to choose the optimal b -value for a particular HARDI acquisition is still an open question but in our particular synthetic experiment, we find the smallest error occurring for b -values between $b =$



(a)



(b)

Figure 7.1: Laplace-Beltrami regularization reduces ODF estimation errors. Plotted is the point-wise Euclidean squared error $[f, f']$ of Eq. 7.13 between the ground truth f and estimated normalized ODFs f' . (a) In the first simulation, the SNR was fixed to 35 while varying the b -value between 500 and 12000 s/mm^2 for combinations of regularization parameter and estimation order of $(\lambda = 0.006, \ell = 8)$ in blue, $(\lambda = 0, \ell = 8)$ in red and $(\lambda = 0, \ell = 4)$ in black. (b) In the second simulation, we fix $b = 3000 s/mm^2$ and estimation order $\ell = 8$ while varying the SNR between 5 and 50.

4000 and 5000 s/mm².

Last, Figure 7.1b shows that the error decreases when the noise level decreases, from more than 12% for a noisy signal (small SNR) to less than 1% for high quality data (large SNR). It is also important to compare the green curve (numerical QBI) with the analytical QBI blue/red curves with/without regularization that we have added in this test to report robustness results of the different approaches. We note that for high quality data, numerical QBI and analytical QBI with/without regularization are almost identical whereas for noisy data, the analytical QBI with regularization $\lambda = 0.006$ performs best, while numerical QBI is better than analytical QBI without regularization ($\lambda = 0$). Finally, as in Figure 7.1a, we have analyzed the behavior of the un-regularized estimation of lower order $\ell = 4$. The un-regularized estimation errors of order 4 overlap the regularized estimation errors of order 8 for low quality data with SNR less than 20 and for high quality data with SNR larger than 20, the un-regularized estimation errors of order 4 overlap the un-regularized estimation errors of order 8. Because of these overlaps and for clarity of Figure 7.1b, we do not plot the curve.

Finally, it is worth mentioning that Tikhonov regularization produces curves that overlap with the $\lambda = 0$ curves in the two experiments. This is why it is not necessary to illustrate it in Figure 7.1. As pointed in the last section, this is expected because Tikhonov regularization does not alter the shape of the ODF. Therefore, the Laplace-Beltrami regularization weight reduces the effect of noise while allowing the use of a high order ℓ in the estimation.

7.5.3 ODF Reconstruction for Different Sampling Schemes

Figures 7.2 and 7.3 show the effect of varying the number of DWI measurements (sampling density) while varying noise level. Figure 7.2 was taken from [(Hess et al., 2006, Fig.3)] and Figure 7.3 extends the study for more sampling density schemes. The signal generation used here is also slightly different from that used in [Hess et al. (2006)]. Overall, we see that the reconstructions behave similarly for high SNR values and also behave similarly for high sampling density schemes. However, as noted in Hess et al. (2006), the analytical ODF reconstruction is more stable for low SNR and low number of DWI directions. Tuch's numerical QBI as well as the un-regularized analytical QBI ($\lambda = 0$) are more affected by noise in the SNR 10 column and when the number of directions is less than 100. The reconstructions are slightly more noisy. Note also that for very low sampling density $N = 21$, Tuch's numerical reconstruction is very smooth and maxima less clear than for the analytical reconstructions. Overall, reconstructions of the ODF with Laplace-Beltrami regularization ($\lambda = 0.006$) more stable over a large range of sampling density schemes and SNR values.

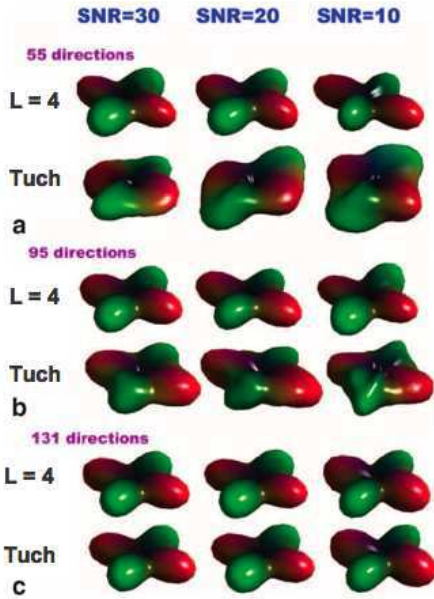


Figure 7.2: ODF reconstruction for different spherical sampling schemes and noise level. Figure taken from [(Hess et al., 2006, Fig.3)]. The analytical QBI solution requires fewer DWI measurements to obtain a stable reconstruction at low SNR. The signal was generated with slightly different parameters, namely $b = 3000 \text{ s/mm}^2$, 90° crossing fibers with equal volume fractions and FAs of 0.7 for each tensor profile.

7.5.4 Fiber Detection and Angular Resolution

In this section, we show four results: 1) Laplace-Beltrami regularized ODFs decrease small perturbations due to noise that can create false maxima, 2) Laplace-Beltrami regularized ODFs improve the detection of crossing fibers while reducing angular error as calculated from the maxima of the ODF and 3) ODF maxima agree with the known underlying fiber configurations under different signal parameters. 4) Angular resolution of the QBI method depends on regularization technique, estimation order ℓ , acquisition b -value and spherical sampling density N .

First, Figure 7.4 shows that Laplace-Beltrami regularized ODFs remove small perturbations due to noise that can create false maxima and errors in detected maxima in ODFs estimated without regularization and with Tikhonov regularization (last two rows of Figure 7.4). In particular, in this test, there are 3 maxima detected for Tikhonov regularization at $\ell = 8$ and for no regularization at $\ell = 8, 10$ and up to 4 maxima (one is less obvious and comes out of the page) for Tikhonov regularization at order $\ell = 10$. Hence, not only there is a danger of over-modeling the data when using a high order ($\ell > 6$) with/without Tikhonov regularization but also, there is an angular error made on the detected maxima at $\ell = 6$. One can potentially tweak the threshold in the fiber detection method to remove some of the spurious maxima but cannot correct for the angular error made. The Tikhonov regularization results

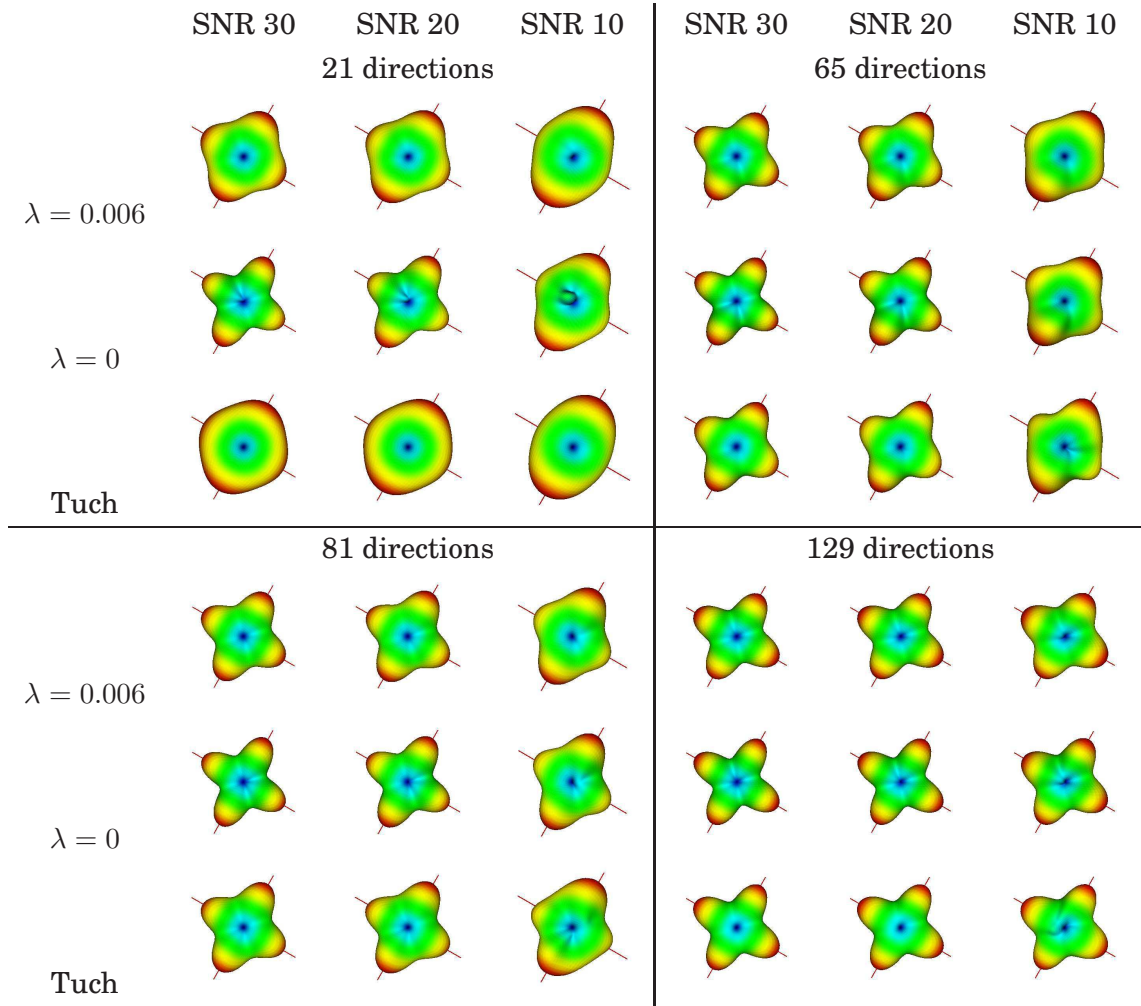


Figure 7.3: At low SNR, the regularized analytical QBI solution requires fewer DWI measurements to obtain a stable reconstruction. For high SNR and with sufficient DWI measurements the un-regularized ($\lambda = 0$) and regularized analytical ($\lambda = 0.006$) QBI as well as the numerical QBI of (Tuch (2004)) are stable and behave similarly. The signal was generated with $b = 3000 \text{ s/mm}^2$, 90° crossing fibers with equal volume fraction and profiles as described in Appendix A.1. The ODF reconstruction was done for order $\ell = 4$. Sampling density $N = 21, 65, 81$ and 129 correspond to a 2nd, 3rd, 3rd and 4th order tessellation of the icosahedron, tetrahedron, icosahedron and octahedron respectively.

shown in Figure 7.4 were obtained with parameter $\lambda = 2$.

Second, Table 7.3 confirms the observation made in Figure 7.4. Table 7.3 shows that Laplace-Beltrami regularized ODFs improve the detection of crossing fibers while reducing angular error as calculated from the maxima of the ODF, as seen in the example of $\ell = 6$ of Figure 7.4. For these tests, we recorded the percentage of correct 2-fiber ODF maxima detected and noted the average angular error \pm standard

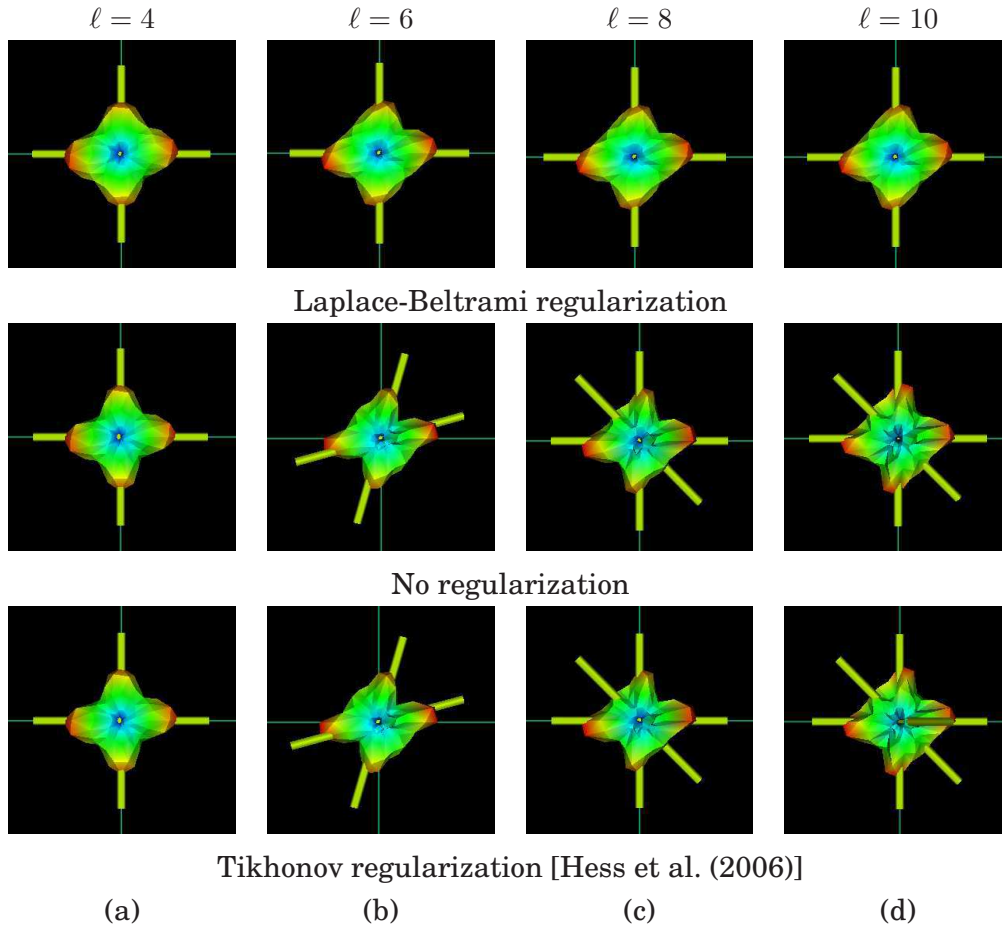


Figure 7.4: ODF regularization decreases small perturbations due to noise that can create false maxima and errors on the detected maxima. Thin lines are the true underlying fiber directions and thicker tubes are the detected ODF maxima. Simulations are done with harmonic order between $\ell = 4, 6, 8, 10$ without regularization ($\lambda = 0$), with Tikhonov regularization ($\lambda = 2$) and with our Laplace-Beltrami regularization ($\lambda = 0.006$). ODFs are generated with $N = 81$, $\text{SNR} = 10$, $b = 3000 \text{ s/mm}^2$ and orthogonal fibers crossing.

deviation made in degrees. In the case there were more than 2 maxima detected, the error was estimated on the two closest ODF maxima to ground truth. With Laplace-Beltrami regularization the detection is nearly perfect at $b = 3000 \text{ s/mm}^2$ and above 88% at $b = 1000 \text{ s/mm}^2$ for all orders whereas the detection dramatically decreases for high order estimations $\ell = 6, 8, 10$ for Tikhonov regularization and without regularization. Table 7.3 shows also that Laplace-Beltrami smoothing reduces the average angular error as calculated from the ODF maxima detected and their ground truth. Overall, orthogonal fibers are detected accurately by all methods for $\ell = 4$ even at the lower b-value. The approximation is smooth enough that effects due to noise are reduced. However, Laplace-Beltrami regularization is necessary to obtain good results

b -value s/mm ²	order ℓ	LB ($\lambda = 0.006$) $\lambda = 0.006$	without LB $\lambda = 0$	Tikhonov $\lambda = 2$
3000	4	99.9%, $2.1 \pm 5.4^\circ$	99.6%, $1.6 \pm 4.7^\circ$	99.6%, $1.6 \pm 4.7^\circ$
	6	99.6%, $2.8 \pm 6.1^\circ$	95.8%, $4.4 \pm 7.2^\circ$	95.9%, $4.2 \pm 7.0^\circ$
	8	99.4%, $2.5 \pm 5.8^\circ$	62.9%, $4.6 \pm 7.4^\circ$	63.1%, $4.5 \pm 7.2^\circ$
	10	99.6%, $2.6 \pm 5.8^\circ$	31.5%, $6.0 \pm 7.5^\circ$	31.1%, $6.6 \pm 7.7^\circ$
1000	4	96.2%, $8.6 \pm 10.6^\circ$	96.1%, $7.1 \pm 8.9^\circ$	96.2%, $7.0 \pm 8.8^\circ$
	6	90.3%, $10.4 \pm 10.8^\circ$	69.4%, $11.9 \pm 10.1^\circ$	71.0%, $11.1 \pm 9.7^\circ$
	8	88.5%, $10.8 \pm 11.4^\circ$	23.4%, $11.3 \pm 10.9^\circ$	24.3%, $11.1 \pm 10.5^\circ$
	10	88.0%, $10.8 \pm 11.3^\circ$	4.5%, $12.6 \pm 10.7^\circ$	4.4%, $12.8 \pm 10.2^\circ$

Table 7.3: Laplace-Beltrami regularization improves the percentage of detecting crossing fibers while reducing angular error as calculated from the maxima of the ODF. Percentage reflects if ODF were correctly detected with 2 maxima. We report success %, average angular error in degrees \pm standard deviation in degrees for the ODF estimated with Laplace-Beltrami (LB) regularization, with Tikhonov regularization and without regularization. The simulations are HARDI signal of 1000 separate pairs of orthogonal fibers, with a SNR of 10 and spherical sampling density of $N = 81$, while varying estimation order ℓ and the b -value. For $\ell > 4$, Laplace-Beltrami regularization is necessary to obtain good fiber detection.

for higher order ($\ell > 4$).

The behavior of ODFs in Figure 7.4 and better performance in Table 7.3 due to the Laplace-Beltrami regularization in the ODF estimation are expected. High-order modeling error due to noise is avoided while minimally altering the lower order coefficients involved in the description of the ODF. A regularization approach like that of [Sakaie and Lowe (2007)] based on the gradient squared would most likely also be able to reproduce this behavior. However, the Tikhonov regularization used in [Hess et al. (2006)] is an approach mainly used to improve the numerical conditioning of the matrices. However, it is not designed to smooth the spherical functions as it perturbs the diagonal elements uniformly which has the effect of adding λ to every eigenvalue. This does not change the overall shape of the ODF and does not eliminate spurious peaks. Moreover, in our experiments, we found that both Hess et al SH basis and our basis are well conditioned. That is, the ratio of the largest over the smallest eigenvalue of the $\mathbf{B}^T \mathbf{B}$ matrix involved in the least-square expression of Eq. 6.13 remains small, even when varying density N for $N = 30, 66, 81, 99, 130, 321$ samplings on the hemisphere.

Next, Figure 7.5 qualitatively shows that we are able to reliably recover the underlying fiber population for imaging parameters similar to clinical applications. The example is generated with b -value of 1500 s/mm², with sampling density $N = 81$, with noise level of SNR = 15 and with a 2-fiber population crossing at 90° and 60° respec-

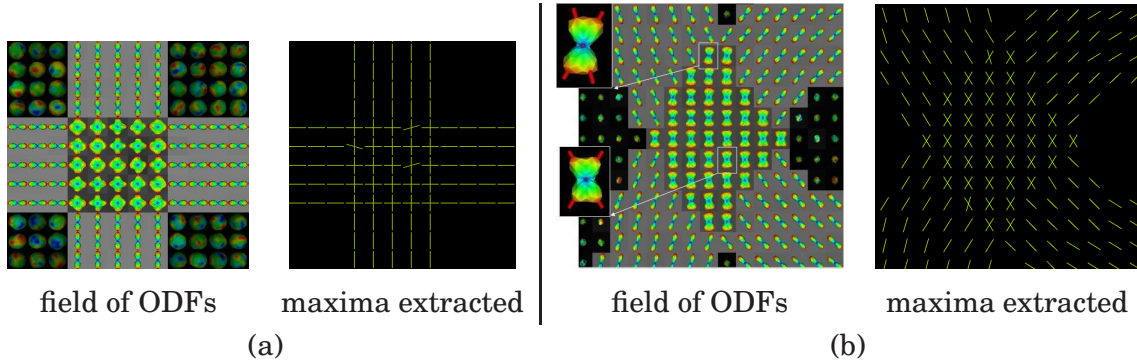


Figure 7.5: ODFs recover multiple fiber crossing in synthetic data generated with $b = 1500 \text{ s/mm}^2$, $\text{SNR} = 15$. (a) 90 degree crossing and (b) 60 degree crossing. An order 6 estimation with $\lambda = 0.006$ was used.

tively. Other examples for b -values ($500, 1500, 3000 \text{ s/mm}^2$) and other SNR ($5, 15, 35$) can be found in [Descoteaux et al. (2005b)]. Note that the detected ODF maxima qualitatively agree with the underlying fibers with some small angular error due to noise level, even though a relatively low b -value is used. The average angular error between ground truth directions and detected maxima was calculated to be approximately 10° for both datasets. Simulation with a 3rd order tessellation gives a 16° difference between each ODF reconstruction point. The error is thus less than an angular sampling unit and separating both fiber configurations is done with the same quantitative precision. As one would expect, the smaller the angle between fibers, the harder it is to distinguish them.

Last, Table 7.4 illustrates the angular resolution of the analytical QBI technique with respect to the SH order ℓ and the regularization method used in the ODF estimation as well as the b -value and spherical sampling density used to generate the synthetic signal. The critical separation angle is reported, i.e. the angle between 2 fibers under which only a single ODF maximum starts to be detected. Five expected observations can be made from Table 7.4. 1) It is harder to distinguish crossing fibers for lower b -values. In fact, there is approximately a 15° gain in angular resolution when going from b -value 1000 s/mm^2 to b -value 5000 s/mm^2 and higher. 2) For $\ell > 4$ there is an improvement in angular resolution of roughly 3 to 10° for all methods because higher order estimation include higher order frequencies in the approximation of the signal. The largest angular resolution improvement occurs for high b -value and high spherical sampling density ($N = 321$). 3) Tikhonov and no regularization have a better angular resolution than using Laplace-Beltrami regularization. In the Laplace-Beltrami case, the critical angle remains reasonable and about 0 to 5° higher. This decrease in angular resolution is due to smoothing of the higher frequency information in Laplace-Beltrami regularization. Hence, there is a trade-off between accuracy of the fiber detection and angular resolution. 4) There is an improvement of

(a) $N = 81$												
ℓ	$b = 10000 \text{ s/mm}^2$			$b = 5000 \text{ s/mm}^2$			$b = 3000 \text{ s/mm}^2$			$b = 1000 \text{ s/mm}^2$		
	LB	TK	$\lambda = 0$	LB	TK	$\lambda = 0$	LB	TK	$\lambda = 0$	LB	TK	$\lambda = 0$
4	57°	55°	55°	59°	56°	56°	63°	60°	60°	75°	71°	71°
6	53°	49°	49°	55°	52°	53°	59°	53°	54°	74°	68°	68°
8	52°	46°	48°	55°	50°	50°	58°	53°	53°	74°	68°	68°
10	52°	46°	47°	55°	49°	49°	58°	52°	53°	74°	68°	68°

(b) $N = 321$												
4	54°	54°	54°	56°	55°	55°	60°	59°	59°	72°	71°	70°
6	44°	42°	43°	47°	45°	46°	52°	50°	50°	69°	67°	67°
8	39°	35°	36°	44°	40°	41°	50°	46°	47°	69°	66°	67°
10	36°	31°	31°	43°	38°	39°	50°	46°	46°	69°	66°	67°

Table 7.4: Angular resolution limitations of the ODF reconstruction with/without Laplace-Beltrami (LB) ($\lambda = 0.006$) and Tikhonov (TK) [Hess et al. (2006)] regularization. Noise-free synthetic HARDI profiles are generated for 2 fibers, for b -values of 10000, 5000, 3000 and 1000 s/mm^2 and for spherical sampling density using $N = 81$ and $N = 321$. The crossing angle between fibers is varied to report the critical angle under which only a single maxima starts to be detected instead of two.

a few degrees in angular resolution when increasing spherical sampling density used to generate the synthetic data. 5) For higher b -values and higher sampling density, there is an important angular resolution improvement of about 10° when comparing the Laplace-Beltrami regularized estimation of order $\ell > 4$ with the Tikhonov and $\lambda = 0$ estimations at order $\ell = 4$. This increase in angular resolution is even more apparent at higher b -values.

In a different angular resolution study, [Hess et al. (2006)] have studied the theoretical relationship between the approximation order ℓ and angular resolution, as calculated from the full width half max of the main lobe of the spherical point spread function (PSF). In particular, for $\ell = 4, 6, 8, 10$, the angular resolution is approximately $65^\circ, 45^\circ, 35^\circ, 30^\circ$ respectively. This is described in [(Hess et al., 2006, Fig.1)]. This study mainly looks at the impact of SH estimation order without considering the HARDI signal parameters or the regularization technique used in the estimation. In practice, we observe critical angles that are higher because there are many other parameters than order ℓ that may influence angular resolution of the solution. In this section, not only have we studied the SH order ℓ but we have also studied the regularization method and λ parameter as well as the b -value and the spherical sampling density N in synthetic simulations.

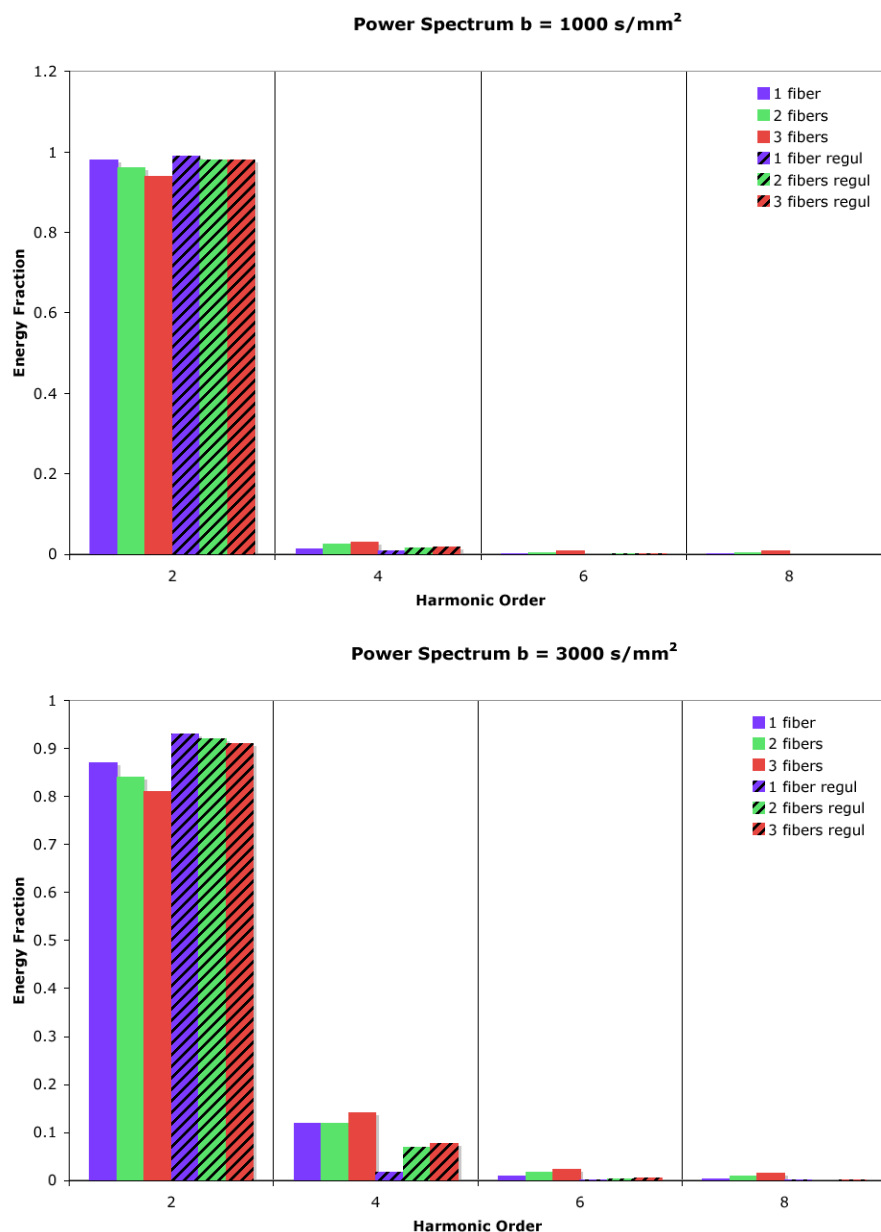


Figure 7.6: Power spectrum (Eq. 7.14) of the spherical harmonic representation at low b -values. ODFs were estimated with regularization $\lambda = 0.006$ (indicated with *regul* in the legend of the figure and dashed texture on the chart bars) and without regularization for $\ell = 10$ from $N = 81$ spherical sampling density and SNR 35 for different b -value. 1000 HARDI profiles were simulated for 1 fiber (blue), 2 fibers (green) and 3 fibers (red). Note that for low b -values, most of the SH information is contained in the SH orders $\ell \leq 4$.

7.5.5 Power Spectrum of the Spherical Harmonic Representation

Figures 7.6 and 7.7 are the power spectrum graphs summarizing the fraction of the power that each harmonic order has in the ODF reconstruction for different b -values.




Figure 7.7: Power spectrum (Eq. 7.14) of the spherical harmonic representation at high b -values. ODFs were estimated with/without regularization as in the previous figure. Note that for high b -values, high order $\ell \geq 6$ SH information is more important than at low b -values. The effect of noise is also seen in the un-regularized reconstructions where there is an abnormal amount of energy in the high SH orders.

Figure 7.6 shows the power spectrum of the ODF SH representation at low b -values. We note that most of the SH information is contained in the SH orders $\ell \leq 4$. In fact, at $b = 1000 \text{ s/mm}^2$, more than 95% and 98% of the energy is contained in order $\ell \leq 4$ without/with regularization respectively. At $b = 3000 \text{ s/mm}^2$ the energy is slightly lowered by more or less 6 to 10%. Hence, most of the SH information is contained in

harmonics of order 4 and less for QBI reconstructions from HARDI signal at relatively low b -values ($b \leq 3000$ s/mm²).

Figure 7.7 shows the power spectrum of the ODF spherical harmonic representation at high b -values. Note that for high b -values, high order ($\ell \geq 6$) SH information is more important than at low b -values. We note strong contributions in the SH order of $\ell = 6$ and $\ell = 8$. As expected, energy is even higher for $\ell = 6$ and $\ell = 8$ when the signal has a multiple fiber configuration, where higher frequency information is needed to describe the crossings. Another important effect seen is the strong effect of noise in the un-regularized reconstructions (solid chart bars). There is an abnormal amount of energy in the high SH orders of $\ell = 6$ and $\ell = 8$. This effect is dramatic at $b = 8000$ s/mm² where energy in SH coefficients of order 8 is higher than energy of coefficients of order 6. This is due to the over-modeling of the estimation and the fact that high order coefficients are perturbed by noise. In the regularized case (dashed chart bars), this noise effect is not seen and as expected, the energy decreases as SH order increases.

Conclusion of Synthetic Data Experiments

There are several important messages from this section. Overall, we showed that there is a trade-off between angular resolution and accuracy of the fiber detection. We have seen that the Laplace-Beltrami regularization reduces errors in fiber detection while keeping a reasonable angular resolution, and in the case of high b -value acquisitions with high number of sampling directions on the sphere, the Laplace-Beltrami regularized solution for high orders ($\ell > 4$) has better angular resolution than the unregularized solution at order $\ell = 4$. The Laplace-Beltrami regularization also allows for more stable ODF reconstruction over a large range of sampling scheme density and SNR values. Moreover, although still an open question, we find that the smallest estimation errors in our simulations occur for b -values between 4000 s/mm² and 5000 s/mm². The best angular resolution is also observed for high b -values of 5000 s/mm² and larger where a 15° gain in angular resolution is observed when compared to a $b = 1000$ s/mm² simulation. Finally, we have also seen that the analytical QBI solution is up to 15 times faster than the standard numerical QBI solution. These are advantages in favor of our analytical QBI solution.

7.5.6 Biological Phantom Results

Figure 7.8 shows that ODFs recover multiple fiber crossing in the rat biological phantom of Appendix A.2. Note that the ODFs have multiple peaks that agree with the known underlying fiber population, which we have emphasized with the sub-figure showing only ODF maxima. This is not the case when looking at the DTI ellipsoids which are flat and sphere-like in the crossing. Moreover, ODF shapes shown in the second row of Figure 7.8 are qualitatively nearly identical. In fact, when comput-

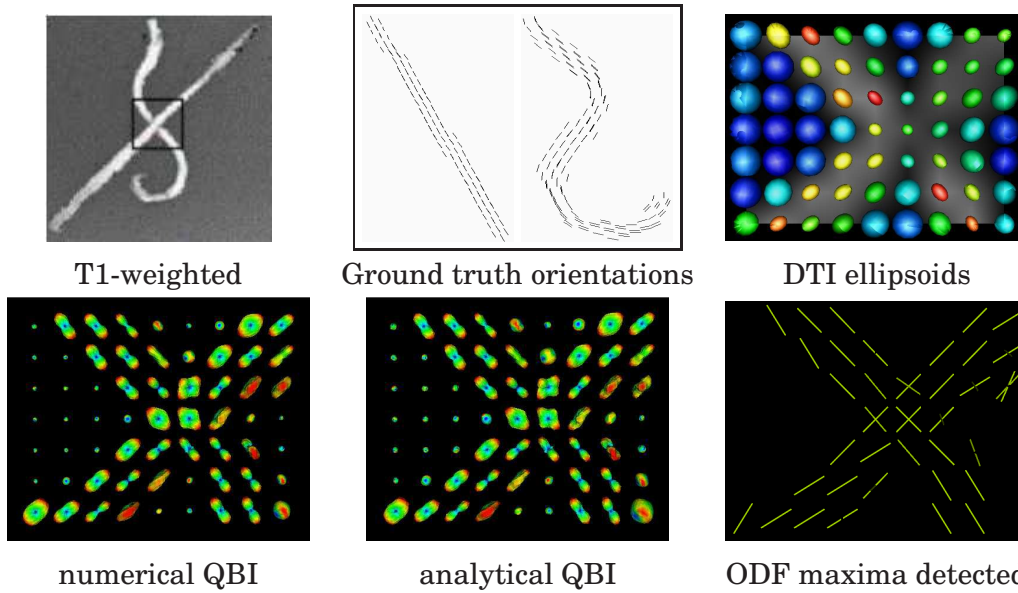


Figure 7.8: ODFs recover multiple fiber crossing in the rat biological phantom. Ground truth orientations were obtained from [Savadjiev et al. (2006)]. An order 6 estimation with $\lambda = 0.006$ was used.

ing the mean and standard deviation of the average Euclidean squared difference (Eq. 7.13) between the ODFs at every voxel of the volume, we obtain a $0.55 \pm 0.17\%$ difference, demonstrating the strong agreement between the methods. That is, ODFs computed from the analytical and numerical QBI [Tuch (2004)] are more than 99% in agreement.

Table 7.5 shows that analytical and numerical QBI [Tuch (2004)] methods yield essentially the same results on the rat phantom while reducing errors obtained from DTI. To perform a quantitative evaluation of the ODF maxima, we used the “ground truth” orientations from [Savadjiev et al. (2006)] (illustrated in Figure 7.8). The orientations are in fact more like a silver standard (a gold standard does not exist in a biological phantom) as they were determined by extracting the centerlines of each super-sampled rat cord and then smoothly extended to the center of the boundaries. For each ODF dataset and for the DTI ellipsoids, the maxima are extracted. Then, at each voxel, the smallest angular difference between the available maximum(a) and ground truth orientation(s) is recorded. The median and mean \pm standard deviation (std) orientation errors in degrees are shown in Table 7.5. The analytical and numerical QBI methods yield essentially the same results while significantly reducing the errors obtained from the DTI profiles. It is important to note that due to the discrete sampling of the sphere, even perfect ODF data will be expected to have some error with respect to the ground truth orientations. This minimal expected error is related to the solid angle subtended by one facet of the sphere tessellation induced by the sampling. For example, it can be found to equal 7.2° for a uniform sampling of the

	Comparison with ground truth directions	
	median	mean \pm std
analytical QBI	12.20°	15.94 \pm 15.32°
numerical QBI	12.19°	15.94 \pm 15.40°
DTI	15.2°	19.4 \pm 16.2°

Table 7.5: Analytical and numerical QBI methods yield essentially the same results on the rat phantom while reducing errors obtained from DTI. Ground truth orientations and comparison are done as in [Savadjiev et al. (2006)].

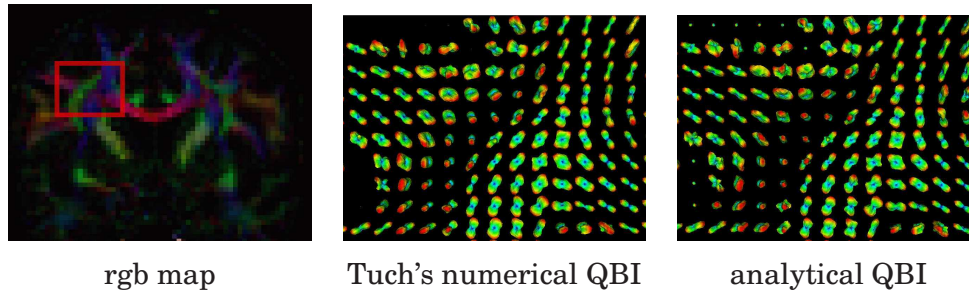


Figure 7.9: Overall shapes of the ODFs are nearly the same for the analytical and numerical QBI. The region of interest shows the cortical spinal tract (cst) and corpus callosum (in the plane) and the cst and longitudinal superior fibers (coming out of the plane) cross.

hemisphere by 100 directions [Savadjiev et al. (2006)].

7.5.7 Human Brain HARDI Data Results

As for the rat biological phantom, if we compare ODFs on the human brain dataset of Appendix A.3.1, the overall shapes of the ODFs are nearly the same for the analytical and numerical QBI [Tuch (2004)] methods. We record a small mean and standard deviation of the average Euclidean squared difference (Eq. 7.13) between ODFs from the two methods of $0.68 \pm 0.23\%$. This is seen qualitatively in Figure 7.9. Here, $b = 3000 \text{ s/mm}^2$ with isotropic voxels of 2 mm.

Crossings in the Max Planck Institute Dataset

Qualitatively, Figure 7.10 shows that ODFs recover multiple fiber crossing in the human brain dataset of Appendix A.3.3 where DTI profiles are limited. Here, $b = 1000 \text{ s/mm}^2$ with isotropic voxels of 1.7 mm. Diffusion tensors (DTs) were estimated with a classical least-squares method. DTs and ODFs are overlaid on the classical FA anisotropy measure and its high order generalization GFA [Tuch (2004)] respectively. We zoom on the ODFs in two ROI. Again, ODF maxima agree with our knowledge of




Figure 7.10: ODFs recovering multiple fiber crossing in a ROI similar to the previous figure where DTI profiles are limited. The ODFs and diffusion tensors are overlaid on the GFA and FA measure respectively and we zoom on the ODFs of two smaller ROI. The ROI(a) shows crossing fibers between the cortical spinal tract (cst) and superior longitudinal fibers (coming out of the plane) and the ROI(b) shows crossing between the corpus callosum (in the plane) and the cst. ROI(a) is tilted to see the fiber crossing better. An order 6 estimation with $\lambda = 0.006$ was used.

the crossings between the cortical spinal tract (cst) and superior longitudinal fibers (coming out of the plane) in ROI (a) and crossings between the cst and corpus callosum (in the plane) in ROI (b). Figure 7.10 also emphasizes the limitations of DTI and the ability of the ODF to recover multiple fiber orientations that are nearly orthogonal even from a dataset with relatively low b -value of 1000 s/mm^2 .

Crossings in the CMRR Dataset

We illustrate the analytical ODF estimation on three regions of interests (ROI) in the human brain dataset of Appendix A.3.2. Here, $b = 1000 \text{ s/mm}^2$ with isotropic voxels




Figure 7.11: Axial slice showing intersection between the genu of the corpus callosum, the capsule fibers and the superior frontal gyrus fibers such as in Tuch (2004). An order 6 estimation with $\lambda = 0.006$ is used.

of 3 mm. Figures 7.11, 7.12 and 7.13 were selected in an axial, coronal and sagittal slice respectively. We highlight the ROI on the FA map and compute the least-squares DTI ellipsoids in that ROI. We adjust the opacity of each ODF surface with respect to the underlying GFA measure.

The first ROI in Figure 7.11 is in an axial slice showing the intersection between the genu of the corpus callosum, the capsule fibers and the superior frontal gyrus fibers. This slice is taken towards the front of the head and corresponds to the ROI of [(Tuch, 2004, Fig.11)]. Note that we pick up the main fiber bundles as well as multiple fiber voxels where we detect evidence for multiple directions.

The second ROI in Figure 7.12 is a coronal slice of the intersection between transcallosal projections of the body of the corpus callosum, the corona radiata and the superior longitudinal fasciculus. This region is very similar to the ROIs in the centrum semiovale of [Tuch et al. (2005)] and [Tournier et al. (2004)]. We see two important crossings. First, the corona radiata crossing the transcallosal fibers projecting to the precentral gyrus and secondly, we have high-lighted the crossings between transcallosal projections and superior longitudinal fasciculus. The superior longitudinal fibers are harder to see as they come out of the page. This is also the case for the corpus callosum body coming out of the page as well in the lower left corner.

Finally, we show a sagittal slice in Figure 7.13 taken in the corona radiata showing diverging and crossing fibers. We detect multiple fibers towards the top right of the brain. Some crossings are due to diverging and splitting fibers in the corona radiata. There are also crossings between the corona radiata and the superior longitudinal

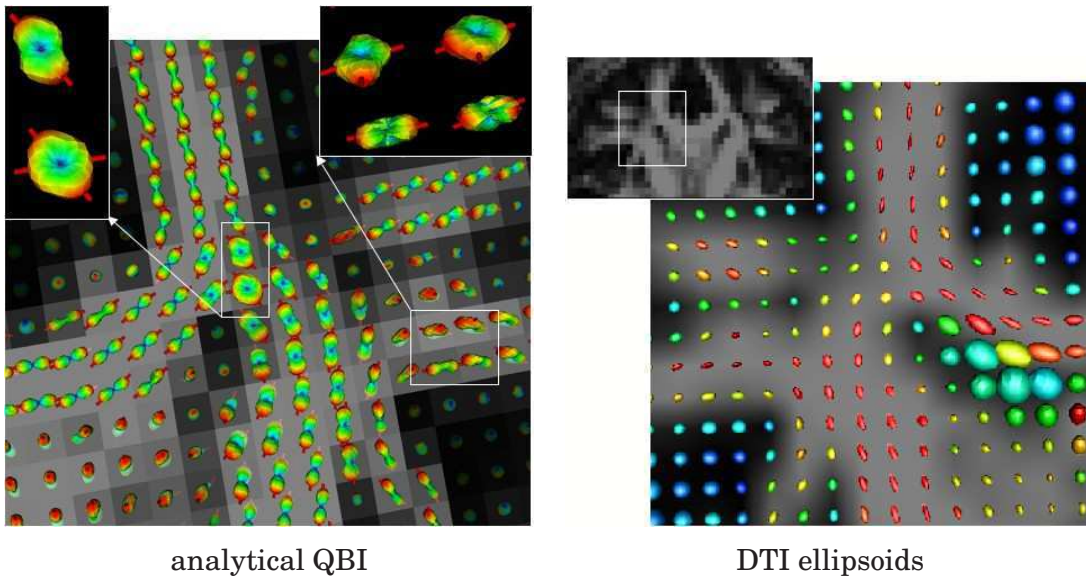


Figure 7.12: Coronal slice in the centrum semiovale, where the intersections between the corpus callosum commissural fibers and the corona radiata and superior longitudinal fasciculus are seen. The zoomed ROI on the right is slightly tilted to better see the directions coming out of the plane. An order 6 estimation with $\lambda = 0.006$ is used.




Figure 7.13: Sagittal slice showing the corona radiata diverging fibers and crossings with the superior longitudinal fasciculus. An order 6 estimation with $\lambda = 0.006$ is used.

fasciculus (left-right) fibers.

7.6 DISCUSSION

In this chapter, we have developed a regularized, fast and robust analyti-

cal solution for the ODF estimation problem in QBI. We have derived the analytical solution for QBI that agrees with previous solutions [Anderson (2005); Hess et al. (2006)] recently developed in parallel in the literature. As already pointed out in the chapter, [Hess et al. (2006)] paper is very close in spirit to our approach but both complement each other. Each makes separate contributions that support the strength of the analytical solution using SH. Hess et al have shown interesting properties of the analytical solution such as the theoretical angular resolution limitation with respect to the harmonic order ℓ of the basis, the better performance of the analytical solution for low sampling density schemes of q-ball ODF reconstruction compared to Tuch's numerical technique and the data compression potential of the approach. Moreover, they have studied the power spectrum of the SH decomposition of the ODF. We have reported these results and successfully reproduced and extended them for different signal generation and different b -value, sampling density scheme N and SNR.

Further, we have made additional contributions: 1) we have included the Laplace-Beltrami regularization in the signal-fitting step, 2) we have proved a novel corollary to the Funk-Hecke theorem to derive the analytical solution for ODF reconstruction, 3) we have done a complexity analysis with a practical running time experiment that shows an increase in speed by a value of 15 over numerical QBI [Tuch (2004)], 4) we have shown the advantages of Laplace-Beltrami regularization theoretically and experimentally and 5) we have performed a careful validation of both analytical and numerical QBI [Tuch (2004)] techniques on noisy synthetic data and real data where ground truth is known.

The new analytical solution comes from the fitting of the signal with SH, which allowed us to impose the Laplace-Beltrami regularization criterion. The Funk-Hecke corollary was proved using a delta sequence so that the Funk-Hecke formula could be used to solve the FRT integral and obtain a simple regularized expression for the ODF reconstruction. Without this derivation, the FRT can only be computed with a more complicated numerical scheme. This solution eliminates the discrete numerical integration step over each equator needed in Tuch's numerical QBI implementation which speeds up computation by a factor of up to 15, while solutions stay in close agreement. As imaging techniques are improved and the number of gradient directions are increased, this can potentially be an important speed-up value even with the optimal numerical QBI implementation. It can also have potential data compression applications.

We have also shown the better performance and robustness of the ODF reconstruction in the presence of noisy synthetic data. Theoretically, we have argued that the Laplace-Beltrami is the right regularization criterion to use. Experimentally, at the cost of slightly reducing angular resolution, we have shown that it introduces less noise in the high order SH coefficients describing the ODF and thus better describes voxels with multiple fibers, especially in the presence of high noise level. Therefore, it is possible to use a high order approximation while limiting the over-modeling of

perturbations due to noise. This is not the case without regularization or when using Tikhonov regularization.

Determining the angular resolution of the analytical QBI method is a difficult problem that depends on the b -value and spherical sampling density N used in the QBI acquisition, and regularization parameter λ and estimation order ℓ used in the ODF reconstruction. Some answers to this angular resolution problem have been proposed in the literature. [Tuch (2004)] reported the angular limitation of the QBI protocol due to the intrinsic Bessel function smoothing of the FRT whereas [Hess et al. (2006)] studied the theoretical angular limitation arising from the choice of SH order ℓ . In general, it would be of great interest to know the achievable angular resolution given the specific sampling N , b -value and SNR in the diffusion-weighted data and given the order ℓ and regularization λ used in the ODF reconstruction.

Overall, we illustrated the power of the SH representation, which have interesting post-processing applications. First, Gaussian/Laplacian smoothing on the sphere by the Laplace-Beltrami operator extension is easy to compute and allows one to impose a regularization criterion on the solution. Second, derivatives and integrals on the sphere have analytical expressions. This could allow for an automatic maxima extraction on the ODFs. The SH basis will also allow easy ODF sharpening operations to attempt to transform the diffusion ODF into a sharp fiber ODF. This will be the topic of the next chapter. Finally, we will also integrate this SH analytical ODF reconstruction in high order tracking and segmentation algorithms in the following chapters.

7.6.1 Contributions of this chapter:

In summary, the contributions of this chapter are:

- Regularized, fast and robust analytical QBI solution.
- Proof of a new corollary of the Funk-Hecke theorem.
- Comparison of the Laplace-Beltrami regularization of the signal with the state-of-the-art numerical and analytical QBI methods using no regularization and Tikhonov regularization.
- Extensive study of the fiber detection, robustness, compression and reconstruction properties of the analytical QBI solution.
- Quantitative validation performed against ground truth from synthetic data and against real data from a biological phantom and a human brain dataset.

Contributions from this chapter appear in [Descoteaux et al. (2005b, 2006c, 2007d,a)].

ACKNOWLEDGMENTS

Thanks to J.S.W. Campbell and P. Savadjiev for the numerical QBI implementation and for the ground truth computation and comparison on the biological phantom. Moreover, thanks to A. Anwander, Tim Wetzel and T. Knösche of the Max Planck Institute for Human Cognitive and Brain Sciences, Leipzig, Germany for the human brain dataset. Finally, thanks to G. Sapiro of the University of Minnesota and Kamil Ugurbill and Stéphane Lehericy of the Center for Magnetic Resonance Research in Minneapolis, USA.

7.7 APPENDIX A: FUNK-RADON TRANSFORM APPROXIMATES THE DIFFUSION ODF

The proof in this section is more detailed and slightly different than the one published in [(Tuch, 2004, Appendix A)]. Let \mathcal{F}_{2D} , \mathcal{F}_{3D} and \mathcal{G} represent the 2D Fourier, 3D Fourier and Funk-Radon (FRT) transforms respectively. We want to prove that the FRT of Eq. 7.2 approximates the diffusion ODF of Eq. 7.1, i.e. that $\mathcal{G}[S(\mathbf{q})](u) \approx \Psi(u)$. For this, we first state the important Fourier analysis tools required for the proof.

1. **Parseval-Plancherel Theorem** relates the integral of two functions over real space to the integral of their Fourier transforms over q-space as

$$\int_{-\infty}^{\infty} f(x)\bar{g}(x)dx = \int_{-\infty}^{\infty} F(k)\bar{G}(k)dk, \quad (7.15)$$

where $F(k)$ and $\bar{G}(k)$ are the Fourier transforms of $f(x)$ and $\bar{g}(x)$ respectively and \bar{g} is the complex conjugate of g .

2. The **Central Slice Theorem** states that the 2D Fourier transform of the projection of a function $f(x)$ onto the plane defined by the vector \mathbf{u} is the same as the intersection of that plane with the 3D Fourier transform of $f(x)$. Letting \mathbf{u} be a unit normal vector defining the projection plane and \mathbf{x} a point on that plane, we can define the projection \mathcal{L} of a 3D function $f(\mathbf{x})$ as

$$\mathcal{L}[f(\mathbf{x})](\mathbf{u}) = \int_{-\infty}^{\infty} f(\mathbf{x} + \alpha\mathbf{u})d\alpha. \quad (7.16)$$

Define now the intersection of a 3D function $f(\mathbf{x})$ with the plane defined by the normal vector \mathbf{u} to it as $\mathcal{I}[f(\mathbf{x})](\mathbf{u}) = f(\mathbf{x})\delta(\mathbf{x}^T\mathbf{u})$, where δ is the Dirac delta function. Hence, we can formally write the central slice theorem as

$$\mathcal{F}_{2D}[\mathcal{L}[f(\mathbf{x})](\mathbf{u})] = \mathcal{I}[\mathcal{F}_{3D}[f(\mathbf{x})]](\mathbf{u}) \quad \text{or} \quad \mathcal{L}[f(\mathbf{x})](\mathbf{u}) = \mathcal{F}_{2D}[\mathcal{I}[\mathcal{F}_{3D}[f(\mathbf{x})]]](\mathbf{u}) \quad (7.17)$$

if f is antipodally symmetric.

3. The **Hankel transform** \mathcal{H} is a special case of the 2D Fourier transform of a function f when this function has no angular dependence, i.e. $f(x, y) = f(r)$ when written in polar coordinates. Letting $x = r \cos \theta$, $y = r \sin \theta$, $u = q \cos \phi$,

$v = q \sin \phi$, we can derive the Hankel transform as

$$\begin{aligned}
\mathcal{F}_{2D}[f(x)] &= \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} f(x, y) \exp(-2\pi i(ux + vy)) dx dy \\
&= \int_0^{\infty} \int_0^{2\pi} f(r) \exp(-2\pi iqr(\cos \theta \cos \phi + \sin \theta \sin \phi)) r dr d\theta \\
&= \int_0^{\infty} \int_0^{2\pi} f(r) \exp(-2\pi iqr \cos(\theta - \phi)) r dr d\theta \\
&= \int_0^{\infty} f(r) \left(\int_0^{2\pi} \exp(-2\pi iqr \cos \zeta) d\zeta \right) r dr \\
&= 2\pi \int_0^{\infty} f(r) J_0(2\pi qr) r dr \quad \text{since } J_0(z) = \frac{1}{\pi} \int_0^{\pi} \exp(iz \cos \zeta) d\zeta \\
&= \mathcal{H}[f(x, y)]
\end{aligned} \tag{7.18}$$

We can now derive the relation between the diffusion ODF Ψ and Funk-Radon transform \mathcal{G} of the measured signal $S(\mathbf{q})$. To do so, we express both functions in cylindrical coordinates (r, θ, z) . Without loss of generality, the point of the unit sphere \mathbf{u} where we want to evaluate the diffusion ODF is along the z -axis. Thus, we can express the diffusion ODF along \mathbf{u} as the projection of diffusion PDF onto the xy -plane with value taken at the origin.

$$\begin{aligned}
\Psi(\mathbf{u}) &= \int_0^{\infty} P(\alpha \mathbf{u}) d\alpha \\
&= \mathcal{L}[P(0)](\mathbf{u}) \quad \text{(using Eq.7.16)} \\
&= \int_{-\infty}^{\infty} P(0, 0, z) dz \\
&= \int_{-\infty}^{\infty} \int_0^{2\pi} \int_0^{\infty} P(r, \theta, z) \delta(r) \delta(\theta) r dr d\theta dz
\end{aligned} \tag{7.19}$$

Now, the raw signal on a single shell of q -space of radius q' is given by $S(\mathbf{q})$. We can expand the FRT of this signal, $\mathcal{G}_{q'}[S(\mathbf{q})](\mathbf{u})$ and show that it approximates the diffusion ODF. First, since \mathbf{u} is along the z -axis, the FRT in the direction \mathbf{u} is integral over the great circle in the xy -plane, i.e.

$$\begin{aligned}
\mathcal{G}_{q'}[S(\mathbf{q})](\mathbf{u}) &= \int_0^{2\pi} S(q', q\theta, 0) dq\theta \\
&= \int_0^{2\pi} \int_0^{\infty} S(q', q\theta, 0) \delta(q_r - q') q_r dq_r dq\theta \\
&= \int_0^{2\pi} \int_0^{\infty} \underbrace{\mathcal{F}_{2D}[S(q', q\theta, 0)]}_{I_1} \cdot \underbrace{\mathcal{F}_{2D}[\delta(q_r - q')]}_{I_2} r dr d\theta
\end{aligned} \tag{7.20}$$

using Parseval-Plancherel theorem (Eq. 7.15). We solve for I_1 using the central slice theorem (Eq. 7.17) through the xy -plane defined by \mathbf{u} as

$$I_1 = \mathcal{F}_{2D}[S(q', q\theta, 0)] = \mathcal{F}_{2D}[\mathcal{L}[\mathcal{F}_{3D}[P(r, \theta, z)]]](\mathbf{u}) = \mathcal{L}[P(r, \theta, z)](\mathbf{u}) = \int_{-\infty}^{\infty} P(r, \theta, z) dz.$$

Noting that $\delta(q_r - q')$ is independent of θ , we can use the Hankel transform (Eq. 7.18) to evaluate I_2 ,

$$I_2 = \mathcal{H}[\delta(q_r - q')] = 2\pi \int_0^{\infty} \delta(q_r - q') J_0(2\pi q_r r) q_r dq_r = 2\pi q' J_0(2\pi q' r).$$




Figure 7.14: Zeroth-order Bessel function $y = J_0(2\pi zx)$ approaches a Dirac delta function as z increases.

Therefore, inserting I_1 and I_2 into Eq. 7.20, we obtain

$$\mathcal{G}_{q'}[S(\mathbf{q})](\mathbf{u}) = 2\pi q' \int_{-\infty}^{\infty} \int_0^{2\pi} \int_0^{\infty} P(r, \theta, z) J_0(2\pi q' r) r dr d\theta dz. \quad (7.21)$$

Comparing Eq. 7.19 and Eq. 7.21, we note that the approximation depends on how close the zeroth-order Bessel function resembles the Dirac delta function. The width of J_0 is inversely proportional to q' . Hence, the larger b , the narrower the Bessel function, the closer the FRT approximation is to the exact diffusion ODF. This is clear from the shapes of the graph in Fig. 7.14.

7.8 APPENDIX B: RIGOROUS PROOF OF THE ANALYTICAL QBI SOLUTION

The delta sequence definition is

$$\lim_{n \rightarrow \infty} \int_{-\infty}^{\infty} \delta_n(x) f(x) dx = \int_{-\infty}^{\infty} \delta(x) f(x) dx = f(0).$$

It is a straightforward corollary to show that for any continuous function f on the interval $[-1, 1]$, the relation

$$\lim_{n \rightarrow \infty} \int_{-1}^1 \delta_n(x) f(x) dx = \int_{-1}^1 \delta(x) f(x) dx = f(0), \quad (7.22)$$

also holds true.

In order to do a rigorous proof of the step between the first and second line of Eq. 7.7, we need to do a change of variable. The integral along the great circle orthogonal to vector the \mathbf{u} , $I_j(\mathbf{u})$, can be rewritten without lost of generality as the integral along the xy -plane equator, where \mathbf{u} is assumed to be along the z -axis. Letting (θ, ϕ) be the natural spherical coordinates defined relatively to vector \mathbf{u} , we can write $I_j(\mathbf{u})$

as

$$\begin{aligned}
I_j(\mathbf{u}) &= \int_{|\mathbf{w}|=1} \delta(\mathbf{u}^T \mathbf{w}) Y_j(\mathbf{w}) d\mathbf{w} \\
&= \int_0^{2\pi} Y_j\left(\frac{\pi}{2}, \phi\right) d\phi && \text{(equator integral on the } xy\text{-plane)} \\
&= \int_{-1}^1 \delta(\zeta) \left(\int_0^{2\pi} Y_j(\arccos(\zeta), \phi) d\phi \right) d\zeta && (\arccos(0) = \pi/2) \\
&= \lim_{n \rightarrow \infty} \int_{-1}^1 \delta_n(\zeta) \left(\int_0^{2\pi} Y_j(\arccos(\zeta), \phi) d\phi \right) d\zeta && \text{(delta sequence corollary Eq. 7.22)} \\
&= \lim_{n \rightarrow \infty} \int_0^{2\pi} \int_0^\pi \delta_n(\cos \theta) Y_j(\theta, \phi) \sin \theta d\theta d\phi && \text{(change of variable } \theta = \arccos(\zeta)\text{)} \\
&= \lim_{n \rightarrow \infty} \int_{|\mathbf{w}|=1} \delta_n(\mathbf{u}^T \mathbf{w}) Y_j(\mathbf{w}) d\mathbf{w} && (\mathbf{u}^T \mathbf{w} = \cos \theta) \\
&= 2\pi \left(\lim_{n \rightarrow \infty} \int_{-1}^1 \delta_n(t) P_{\ell(j)}(t) dt \right) Y_j(\mathbf{u}) && \text{(Funk-Hecke theorem Eq. 5.31)} \\
&= 2\pi P_{\ell(j)}(0) Y_j(\mathbf{u}), && \text{(delta sequence corollary Eq. 7.22)}
\end{aligned}$$

This is the detailed proofs using the delta sequence.

7.9 APPENDIX C: EXACT ODF FROM THE MULTIPLE-TENSOR MODEL

Assuming signal $S(\mathbf{u})$ is generated from the multi-tensor model for n fibers with relative weight p_k ,

$$S(\mathbf{u}) = \sum_{k=1}^n p_k e^{-b\mathbf{u}^T \mathbf{D}_k \mathbf{u}},$$

we want to derive the exact ODF for this fiber distribution. Recalling that the signal is the Fourier transform \mathcal{F} of the underlying probability density function of water molecules $P(\mathbf{r})$, we first need to compute the inverse Fourier transform \mathcal{F}^{-1} of $S(\mathbf{u})$ to obtain the corresponding PDF. We consider a single tensor and then, by linearity, we obtain the result for the general multiple tensors. We will need the following inverse Fourier transform expression

$$\mathcal{F}^{-1}[\exp(-bk^2)](x) = \frac{1}{2\pi} \sqrt{\frac{\pi}{b}} \exp\left(\frac{-x^2}{4b}\right). \quad (7.23)$$

We assume a symmetric diffusion tensor \mathbf{D} with diagonal entry $[e_2, e_2, e_1]$, $S_0 = 1$ and signal formation given by $S(\mathbf{u}) = \exp(-b\mathbf{u}^T \mathbf{D} \mathbf{u})$ is⁴

$$\begin{aligned}
P(\mathbf{r}) &= \mathcal{F}^{-1}[\exp(-b\mathbf{u}^T \mathbf{D} \mathbf{u})] \\
&= \mathcal{F}^{-1}[\exp(-b(e_2 u_x^2 + e_2 u_y^2 + e_1 u_z^2))] \\
&= \mathcal{F}^{-1}[\exp(-be_2 u_x^2)] \mathcal{F}^{-1}[\exp(-be_2 u_y^2)] \mathcal{F}^{-1}[\exp(-be_1 u_z^2)] \\
&= \frac{1}{2\pi} \sqrt{\frac{\pi}{be_2}} \exp\left(\frac{-x^2}{4be_2}\right) \frac{1}{2\pi} \sqrt{\frac{\pi}{be_2}} \exp\left(\frac{-y^2}{4be_2}\right) \frac{1}{2\pi} \sqrt{\frac{\pi}{be_1}} \exp\left(\frac{-z^2}{4be_1}\right) \quad (\text{Eq. 7.23}) \\
&= \frac{1}{8\pi^3} \sqrt{\frac{\pi^3}{b^3 e_2^2 e_1}} \exp\left(-\frac{1}{4b} \left(\frac{x^2}{e_2} + \frac{y^2}{e_2} + \frac{z^2}{e_1}\right)\right) \\
&= \frac{1}{\sqrt{(4\pi b)^3 |\mathbf{D}|}} \exp\left(\frac{-\mathbf{r}^T \mathbf{D}^{-1} \mathbf{r}}{4b}\right)
\end{aligned}$$

Now, we can obtain the exact ODF by integrating the radial direction of the PDF (Eq 7.1). We will also need the Gaussian integral formula $\int_0^\infty \exp(-r^2 C) dr = 1/2\sqrt{\pi/C}$. Hence, for one fiber we have,

$$\begin{aligned}
\Psi_1(\mathbf{u}) &= \int_0^\infty P(r\mathbf{u}) dr = \frac{1}{\sqrt{(4\pi b)^3 |\mathbf{D}|}} \int_0^\infty \exp\left(\frac{-r^2 \mathbf{u}^T \mathbf{D}^{-1} \mathbf{u}}{4b}\right) dr \\
&= \sqrt{\frac{\pi b}{(4\pi b)^3 |\mathbf{D}|}} \sqrt{\frac{1}{\mathbf{u}^T \mathbf{D}^{-1} \mathbf{u}}} = \frac{1}{Z} \sqrt{\frac{1}{\mathbf{u}^T \mathbf{D}^{-1} \mathbf{u}}}
\end{aligned} \quad (7.24)$$

where Z is the normalization factor for which the integral over the sphere of the ODF is 1, i.e.

$$Z = \int_{|\mathbf{u}|=1} (\mathbf{u}^T \mathbf{D}^{-1} \mathbf{u})^{-1/2} d\mathbf{u} = \sqrt{\frac{(4\pi b)^3 |\mathbf{D}|}{\pi b}} = 8\pi b \sqrt{|\mathbf{D}|} \quad (7.25)$$

Thus, for n fibers, we have

$$\Psi(\mathbf{u}) = \sum_{k=1}^n \frac{p_k}{\bar{Z}} \sqrt{\frac{\pi b}{\mathbf{u}^T \mathbf{D}_k^{-1} \mathbf{u}}}, \quad (7.26)$$

where \bar{Z} is again a normalization factor that assures that the ODF still has the properties of a PDF.

⁴Note that if we use the standard signal formation in \mathbf{q} -space, $S(\mathbf{q}) = \exp(-\tau \mathbf{q}^T \mathbf{D} \mathbf{q})$ and $b = \tau |\mathbf{q}|^2$, we obtain the same expression with τ instead of b (as in [(Tuch, 2004, Eq.2)]).

SEGMENTATION IN HIGH ANGULAR RESOLUTION DIFFUSION MRI

“Go with the flow”

Contents

8.1 Introduction	145
8.2 Background on Segmentation	145
8.3 Statistical Surface Evolution	148
8.3.1 Distances between ODFs	148
8.3.2 Segmentation by Surface Evolution	150
8.4 Segmentation Results & Discussion	153
8.4.1 Synthetic Simulation Results	153
8.4.2 Biological Phantom Results	155
8.4.3 Human Brain HARDI Data Results	156
8.4.4 Multi-Subject Study on a public HARDI Database	159
8.5 Discussion	162
8.5.1 Contributions of this chapter:	164

OVERVIEW

Now that we have robust q-ball reconstructions, can we segment this image of q-ball diffusion ODF into coherent fiber bundles? This chapter covers a new statistical region-based active contour algorithm to segment HARDI data. In particular, how can the segmentation problem be formulated and solved efficiently on HARDI data? What is gained by the HARDI segmentation with respect to the DTI segmentation? Is it possible to validate the segmentation results and make the segmentation automatic? First, a new method to segment HARDI data is developed using our fast, robust and analytical QBI solution. Then, a region-based statistical surface evolution is done on this image of diffusion ODF to efficiently find coherent white matter fiber bundles. Next, our method is shown to be appropriate to propagate through regions of fiber crossings and our results outperform state-of-the-art DTI segmentation methods. Finally, results obtained on synthetic data, on a biological phantom, on real datasets and on all 13 subjects of a public HARDI database show that our method is reproducible and brings a strong added value to DW-MRI segmentation.

Keywords: region-based segmentation, level set framework, Euclidean and Riemannian distances, DTI, HARDI, QBI, diffusion ODF

Contributions of this chapter:

- Definition of an efficient and simple distance measure between spherical functions based on their SH representation.
- New statistical region-based segmentation framework for HARDI data segmentation.
- Efficient and semi-automatic implementation using the level set method.
- Validation and reproducibility of the results on a biological phantom and the public HARDI database.

Organization of this chapter:

The chapter is organized as follows. We motivate the HARDI segmentation problem in Section 8.1. Then, we review the existing methods in DTI and HARDI segmentation in Section 8.2. Next, Section 8.3 describes our new statistical region-based segmentation method implemented using the level set method. Results on complex synthetic HARDI datasets, on a biological phantom and on real datasets are presented in Section 8.4. Finally, we conclude with a discussion of the results and our contributions in Section 8.5.

8.1 INTRODUCTION

In this chapter, we are interested in recovering larger bundles or groups of fiber tracts instead of recovering fibers one by one¹ like done in tractography. We would like to segment white matter fiber bundles in which diffusion properties are similar and ultimately compare their features to those in other ROI in the same subject or on multiple subjects. The goal is thus to find global coherence that exists among white matter fiber tracts belonging to the same fiber bundle. Most existing segmentation methods are based on DTI [Zhukov et al. (2003); Wiegell et al. (2003); Feddern et al. (2003); Rousson et al. (2004); Wang and Vemuri (2004); Jonasson et al. (2005a); Wang and Vemuri (2005); Lenglet et al. (2006a)] and are inherently limited by the DT model. As expected, these segmentation methods are most often blocked in regions of fiber crossings where DTs are oblate and isotropic. Hence, it is natural to extend existing DT-based methods and to design new segmentation methods based on HARDI data to aid the inference of crossing, branching and kissing fiber configurations. New methods have thus started to appear to segment bundles from fields of ODFs [Jonasson et al. (2005b); Hagmann et al. (2006a); McGraw et al. (2006b); Jonasson et al. (2007)].

In this chapter, we answer the following three questions: 1) How can the segmentation problem be formulated and solved efficiently on a field of ODFs? 2) What is gained by the ODF with respect to the DT? 3) Is it possible to validate the segmentation results? To do so, we propose an efficient region-based level set approach using a regularized and robust SH representation of the ODF presented in Chapter 7. We first show that a better local modeling of fiber crossings improves segmentation results globally. Then, we show that our ODF segmentation is more accurate than the state-of-the-art DTI segmentation based on the Euclidean and Riemannian distances in regions of complex fiber configurations from synthetic data, from a biological phantom and from real data. The ODF better captures statistics in crossing areas and is thus able to flow through complex fiber regions without leaking in the whole white matter. Finally, we show that our q-ball segmentation is reproducible by segmenting automatically the corpus callosum (CC) and the cortico spinal tract (CST) of the 13 subjects of a public HARDI database [Poupon et al. (2006)].

8.2 BACKGROUND ON SEGMENTATION

DT-based segmentation The existing DT-based segmentation techniques in the literature are [Zhukov et al. (2003); Wiegell et al. (2003); Feddern et al. (2003); Rousson et al. (2004); Wang and Vemuri (2004, 2005); Jonasson et al. (2005a);

¹Note that in tractography we are not actually tracking individual fibers at the micrometer scale but actually tracking bundles of fibers at the millimeter scale corresponding to the size of the imaging voxel.

Lenglet et al. (2006a)]. They can be grouped into four categories: 1) scalar-based method [Zhukov et al. (2003)], 2) Euclidean distance methods based on the full DT [Wiegell et al. (2003); Feddern et al. (2003); Rousson et al. (2004); Wang and Vemuri (2004)], 3) Geometric-based measure of dissimilarity between DTs [Jonasson et al. (2005a)] and 4) more complex DT metric-based approaches such as the Kullback-Leibler divergence metric [Wang and Vemuri (2005)] and Riemannian metric [Lenglet et al. (2006a)].

In [Zhukov et al. (2003)], they define an invariant anisotropy measure in order to drive the evolution of a level set and isolate strongly anisotropic or strongly isotropic regions of the brain, such as the CC and the ventricles. However, the reduction of the full tensor to a single scalar value can result in a relatively low discriminate capability and potentially yields the segmentation of mixed structures. Alternatively, [Wiegell et al. (2003); Feddern et al. (2003, 2004); Rousson et al. (2004); Wang and Vemuri (2004, 2005); Jonasson et al. (2005a); Lenglet et al. (2006a)] use or propose different measures of dissimilarity between DTs. In [Wiegell et al. (2003); Wang and Vemuri (2004); Rousson et al. (2004)], the authors use the Frobenius norm of the difference of tensors (i.e. the Euclidean distance). A k -means algorithm based on the Euclidean distance with a spatial coherence constraint and an active contour model with a regularity term were respectively used by the last two methods ([Wang and Vemuri (2004)], [Wiegell et al. (2003)]) to perform the segmentation of different cerebral structures such as the thalamus nuclei or the corpus callosum. The first method of [Rousson et al. (2004)] also used the Euclidean distance between DTs in a region-based surface propagation inspired by general works on image segmentation [Chan and Vese (2001); Paragios and Deriche (2002)]. In [Wang and Vemuri (2004)], a generalization of the region-based active contours to matrix-valued images is proposed for two-dimensional images, also using the Euclidean distance between DTs. In [Feddern et al. (2003, 2004)], partial differential equations based on mean curvature motion, self-snakes and geodesic active contours models are extended to two-dimensional and three-dimensional tensor-valued images by generalizing the notion of structure tensor to matrix-valued data. This method still relies on the Euclidean metric between tensors. The authors apply this framework to the regularization and segmentation of diffusion tensor images. The other methods [Wang and Vemuri (2005); Jonasson et al. (2005a); Lenglet et al. (2006a)] do not use the Euclidean distance and propose different similarity measures between DTs. In [Jonasson et al. (2005a)], the authors introduce a new geometric measure of dissimilarity by computing the normalized tensor scalar product of two tensors, which can be interpreted as a measure of overlap. Finally, the methods presented in [Wang and Vemuri (2005)] and [Lenglet et al. (2006a)] rely on the symmetrized Kullback-Leibler divergence and Riemannian geodesic metric to derive an affine invariant dissimilarity measure between DTs.

In this chapter, our statistical surface evolution is inspired by the original meth-

ods of [Chan and Vese (2001); Paragios and Deriche (2002)] and recent DTI extension of [Rousson et al. (2004); Lenglet et al. (2006a)]. We use a similar region-based surface propagation but the surface evolution is based on a Euclidean distance between q-ball ODFs described in a SH basis representation.

HARDI-based segmentation To our knowledge there are only two existing HARDI-based segmentation algorithms, namely [Jonasson et al. (2005a); Hagmann et al. (2006a); Jonasson et al. (2007)] and [McGraw et al. (2006b)]. In [Hagmann et al. (2006a); Jonasson et al. (2007)], the segmentation is developed in a non-Euclidean 5-dimensional (5D) position-orientation space, $(x, y, z, \theta, \phi) \in \mathbb{R}^3 \times \mathbb{S}^2$. using the ODF map reconstructed from diffusion spectrum imaging (DSI), a procedure computationally more expensive than QBI. The extension from 3D to 5D space leads to huge 5D matrices, creating large problems with data handling and storage, which only allows one to segment small parts of cerebral structures. [Jonasson et al. (2007)] and [Hagmann et al. (2006a)] use the same algorithm but [Jonasson et al. (2007)] is implemented with a level set framework whereas [Hagmann et al. (2006a)] is implemented with a hidden Markov random field framework. In [McGraw et al. (2006b)], the main contribution is to model the ODF with a mixture of von Mises-Fisher distributions and use the associated metric in a hidden Markov measure field segmentation scheme. Both the ODF modeling and the segmentation technique are different from our proposed method. Moreover, results are promising but only preliminary on synthetic 2D simulations and on a rat spinal cord dataset. Hence, the first approach suffers from computational efficiency and the second is still preliminary and lacks real data results. Furthermore, both HARDI-based methods do not show the gain in using a high-order spherical function such as the ODF in comparison with the simpler DT model. Therefore, the motivations for our new HARDI segmentation algorithm are the three questions previously asked: 1) How can the segmentation problem be formulated and solved efficiently on a field of diffusion ODFs? 2) What is gained by the ODF with respect to the DT? 3) Is it possible to validate the segmentation results on human brain datasets?

In this chapter, we propose an efficient level set statistical surface evolution based on the ODFs computed from QBI that we thoroughly test on synthetic and real data with complex fiber configurations. The level set implementation is simple and efficient and the statistical evolution is robust with automatic convergence. Moreover, the ODF reconstruction from QBI has several advantages that make it a good choice of input image. First, samples are only taken on a single sphere in q-space and thus, the imaging time is much smaller than that of the DSI despite significantly higher angular resolution measurements and good signal to noise ratio. QBI also has the advantage of being model-independent, robust, fast to compute and represented in a compact SH basis, as seen in Chapter 7.

8.3 STATISTICAL SURFACE EVOLUTION

Given a HARDI signal with N discrete measurements on the sphere, we transform it into a spherical harmonic (SH) series of order ℓ with R coefficients ($R \ll N$) describing the diffusion ODF. Then, we look for the most probable partition of this ODF image of SHs using our new region-based statistical surface evolution. We now want to find a global coherence in this field of q-ball ODFs represented in the SH basis. We denote the image of ODFs by $\mathcal{F} : \mathbb{R}^3 \mapsto \mathbb{R}^R$ so that for all $\mathbf{x} \in \mathbb{R}^3$, $\mathcal{F}(\mathbf{x})$ is the ODF of order ℓ at voxel \mathbf{x} represented by a vector of R real SH coefficients, $\mathcal{F}(\mathbf{x}) := \{f_1, \dots, f_R\} \in \mathbb{R}^R$. Recall that $R = (1/2)(\ell + 1)(\ell + 2)$. Now, the question is what is a good metric to compare ODFs?

8.3.1 Distances between ODFs

We want to capture similarities and dissimilarities between two ODFs, i.e two spherical functions $\Psi, \Psi' \in \Omega$ that can be represented by real SH vectors $f, f' \in \mathbb{R}^R$, as shown in the previous chapters. Since the ODFs come from real physical diffusion measurements they are bounded and form an open subset of the space of real-valued \mathcal{L}^2 spherical functions with an inner product \langle, \rangle defined as

$$\langle \Psi, \Psi' \rangle = \int_{\Omega} \Psi(\theta, \phi) \cdot \Psi'(\theta, \phi) d\Omega = \int_{\Omega} \left(\sum_{i=1}^R f_i Y_i(\theta, \phi) \sum_{j=1}^R f'_j Y_j(\theta, \phi) \right) d\Omega, \quad (8.1)$$

where Ω denotes integration over the unit sphere. As before, because of the orthonormality of the SH basis, the cross terms cancel and the expression is simply

$$\langle \Psi, \Psi' \rangle = \sum_{j=1}^R f_j \cdot f'_j. \quad (8.2)$$

Therefore, the induced \mathcal{L}^2 norm $\|\Psi\| = \sqrt{\langle \Psi, \Psi \rangle}$ giving us the distance metric between two ODFs is

$$\|\Psi - \Psi'\| = \sqrt{\sum_{j=1}^R (f_j - f'_j)^2}. \quad (8.3)$$

How does this metric behave when comparing two spherical functions?

To illustrate the metric behavior, we generate synthetic profiles with $b = 1000$ s/mm², FA = 0.7, $N = 81$ directions and vary the azimuth angle θ of the ODF at each pixel (x, y) according to the angle with the center pixel $(x_{\text{center}}, y_{\text{center}})$, i.e. $\theta = \arctan((y - y_{\text{center}})/(x - x_{\text{center}}))$. This gives a distribution of ODFs with varying azimuth angle dispersion, as illustrated in Figure 8.1. At every pixel, the brighter the underlying color, the more distant from the central single maximum ODF is the ODF and conversely, the darker the underlying color, the closer it is. The distance range for each figure is normalized differently to give better idea how the distances between crossing ODFs and single maximum ODFs vary. The metric behaves as




Figure 8.1: Illustration of the distance measure between a single maximum ODF placed at the center and single maximum ODFs on the left. On the right, illustration of the distance measure between a single maximum ODF placed at the center and 90° crossing ODFs with varying angular dispersion. At every pixel, the brighter the underlying color, the more distant from the central single maximum ODF is the ODF and conversely, the darker the underlying color, the closer it is.

one would expect. First, for single maxima ODFs, the distance between two ODFs is smaller when they are aligned and increases as their angular difference increases. Second, for ODFs with crossing configurations and two maxima, the distance between a single maximum ODF aligned with one of the two maxima is smaller than if the crossing ODF is not aligned with the single maximum ODF. Finally, we note that the distance between two single maximum ODFs with angular difference less than about 60° to 75° is smaller than the distance between a crossing ODF and a single maximum ODF aligned with one of the maxima of the crossing ODF. That is, it is easier to connect two single maximum ODF even if they are separated by angles of 0-75° than to connect a crossing ODF with a single maximum ODF. This is sketched in Figure 8.2. This property turns out to have the desired effect of facilitating the flow through fanning and branching configurations as well as allowing to flow through pure crossing configurations when the maxima are aligned.

The Euclidean distance was also used successfully for DTI segmentation in [Rousson et al. (2004); Lenglet et al. (2006a)] and DTI registration [Zhang et al. (2004)] even though more appropriate metrics exist such as the symmetrized Kullback-Leibler distance [Wang and Vemuri (2005); Lenglet et al. (2006a)] and Riemannian geodesic distance [Lenglet et al. (2006a)]. Similarly, one can think of choosing another metric to compare ODFs. For instance, since the ODF can be viewed




Figure 8.2: Sketch of distance between crossing and single peaked ODF. The distance between two single maxima ODFs at angles difference between 0 and 75° is smaller than the distance between a crossing ODF with a single maximum ODF that has its peak aligned with one of the crossing maxima.

as a probability distribution function (PDF) of fiber orientations, one can use the Kullback-Leibler distance between two PDFs. However, in that case the problem quickly blows up computationally because one needs to use all N discrete data on the sphere instead of the R SH coefficients ($R \ll N$). For example, one needs to process $N = 200$ values instead of $R = 15$ SH coefficients.

8.3.2 Segmentation by Surface Evolution

Inspired by general works on image segmentation by [Chan and Vese (2001); Paragios and Deriche (2002)], we search for the optimal partition boundary \mathcal{B} in two regions \mathcal{F}_1 and \mathcal{F}_2 of the image \mathcal{F} . We maximize the *a posteriori* frame partition probability $p(\mathcal{B}|\mathcal{F})$ of obtaining the desired segmentation for the observed image of ODFs \mathcal{F} . The major difference in our approach is that we use order-4 ODFs, with $R = 15$ real coefficients whereas in the DT Euclidean flow of [Rousson et al. (2004)], DTs represented by 6D vectors are used as input to the region-based segmentation.

We use the level set framework [Dervieux and Thomasset (1979, 1981); Osher and Sethian (1988)] to represent the optimal partition boundary \mathcal{B} as the zero-crossing of the level set function ϕ . ϕ is defined as the usual signed distance function, i.e.

$$\phi(\mathbf{x}) = \begin{cases} \mathcal{D}(\mathbf{x}, \mathcal{B}) & \text{if } \mathbf{x} \in \mathcal{F}_1 \\ -\mathcal{D}(\mathbf{x}, \mathcal{B}) & \text{if } \mathbf{x} \in \mathcal{F}_2 \\ 0 & \text{if } \mathbf{x} \in \mathcal{B}, \end{cases} \quad (8.4)$$

where $\mathcal{D}(\mathbf{x}, \mathcal{B})$ is the Euclidean distance between \mathbf{x} and \mathcal{B} . Hence, the optimal partition is obtained by maximizing $p(\phi|\mathcal{F}) \propto p(\mathcal{F}|\phi)p(\phi)$ using Bayes' rule. Assuming




Figure 8.3: Histograms of the SH coefficients representation of the ODF $\{f_1, \dots, f_{15}\}$ in the corpus callosum (CC). The CC was manually segmented from subject 7 of the public HARDI database [Poupon et al. (2006)].

ODFs to be independent within each region, \mathcal{F}_1 and \mathcal{F}_2 , of the image \mathcal{F} , we have

$$p(\mathcal{F}|\phi) = \prod_{\mathbf{x} \in \mathcal{F}_1} p_1(\mathcal{F}(\mathbf{x})) \cdot \prod_{\mathbf{x} \in \mathcal{F}_2} p_2(\mathcal{F}(\mathbf{x})). \quad (8.5)$$

At this point, the main and debatable assumption is that p_1 and p_2 are Gaussians, which means that each SH coefficient of the ODFs follow a Gaussian distribution in the different partitions of the q-ball ODF image. Figure 8.3 and Figure 8.4 show that this Gaussian assumption is reasonable. In fact, we see that the histograms for most of the R coefficients of the ODF are “bell-shaped”. Histograms were computed from all voxels in a manual segmentation of the CC and CST in a subject of the HARDI database [Poupon et al. (2006)] (histograms are similar for all subjects of the database). Hence, we consider a parametric representation with a R -dimensional Gaussian. Letting $\bar{\mathbf{F}}_r \in \mathbb{R}^R$ be the mean SH ODF vector in region $r = 1, 2$ and $\mathbf{\Lambda}_r$ be the $R \times R$ covariance matrix of the ODF vectors in region r , the likelihood of the ODF $\mathcal{F}(\mathbf{x})$ to be part of region r is defined as

$$p_r(\mathcal{F}(\mathbf{x})|\bar{\mathbf{F}}_r, \mathbf{\Lambda}_r) = \frac{1}{(2\pi)^3 |\mathbf{\Lambda}_r|^{1/2}} \exp\left(-\frac{1}{2}(\mathcal{F}(\mathbf{x}) - \bar{\mathbf{F}}_r)^T \mathbf{\Lambda}_r^{-1} (\mathcal{F}(\mathbf{x}) - \bar{\mathbf{F}}_r)\right). \quad (8.6)$$

Concerning the term $p(\phi)$, it expresses the probability that the evolving surface represents the structure of interest and it can be used to introduce a prior shape knowledge [Lenglet et al. (2006a)]. Here, we want to favor structures with smaller surface. Letting $|\mathcal{B}|$ represent the number of points on the boundary, i.e., points on the zero level set, we can introduce the shape prior in the term $p(\phi)$ so that it is proportional to $\exp(-\nu|\mathcal{B}|)$ and favors smooth surfaces. This regularization term can be expressed with ϕ by introducing the Dirac function δ [Zhao et al. (1996); Rousson




Figure 8.4: Histograms of the SH coefficients representation of the ODF $\{f_1, \dots, f_{15}\}$ in the cortico spinal tract (CST). The CST was manually segmented from subject 7 of the public HARDI database [Poupon et al. (2006)].

(2004)] and thus, we have

$$p(\phi) = \exp \left(-\nu \int_{\mathcal{F}} \delta(\phi) |\nabla \phi(\mathbf{x})| dx \right). \quad (8.7)$$

The optimal segmentation is then obtained by maximizing $p(\mathcal{F}|\phi)p(\phi)$ or by minimizing of the negative logarithms of Eqs. (8.6, 8.7). Therefore, the final energy minimization is

$$E(\phi, p_1, p_2) = - \int_{\mathcal{F}_1} \log p_1(\mathcal{F}(\mathbf{x})|\bar{\mathbf{F}}_1, \Lambda_1) dx - \int_{\mathcal{F}_2} \log p_2(\mathcal{F}(\mathbf{x})|\bar{\mathbf{F}}_2, \Lambda_2) dx + \nu \int_{\mathcal{F}} \delta(\phi) |\nabla \phi| dx. \quad (8.8)$$

Then, for given statistical parameters $\bar{\mathbf{F}}_1, \bar{\mathbf{F}}_2, \Lambda_1, \Lambda_2$ and regularization ν of the evolving surface, the Euler-Lagrange [Arfken and Weber (1995)] equation can be computed to derive the implicit surface evolution

$$\begin{aligned} \frac{\partial \phi}{\partial t} &= \delta(\phi) \left(\nu \operatorname{div} \frac{\nabla \phi}{|\nabla \phi|} + \log \frac{p_1(\mathcal{F}(x)|\bar{\mathbf{F}}_1, \Lambda_1)}{p_2(\mathcal{F}(x)|\bar{\mathbf{F}}_2, \Lambda_2)} \right) \\ &= \delta(\phi) \left(\nu \operatorname{div} \frac{\nabla \phi}{|\nabla \phi|} + \frac{1}{2} \log \frac{|\Lambda_2|}{|\Lambda_1|} \right. \\ &\quad \left. + \frac{1}{2} \left(-(\mathcal{F}(x) - \bar{\mathbf{F}}_1)^T \Lambda_1^{-1} (\mathcal{F}(x) - \bar{\mathbf{F}}_1) + (\mathcal{F}(x) - \bar{\mathbf{F}}_2)^T \Lambda_2^{-1} (\mathcal{F}(x) - \bar{\mathbf{F}}_2) \right) \right). \end{aligned} \quad (8.9)$$

The statistics can be updated after each iteration of the ODF flow, as described in [Lenglet et al. (2006a)]. More details on this level set optimization can be found in [Chan and Vese (2001); Vese and Chan (2002); Rousson (2004); Cremers et al. (2007)]. The flow formulation is flexible and one can easily replace the input image vector of ODFs \mathcal{F} with the standard DT coefficient vector, as in [Rousson et al.

(2004)] or any vector representation of a spherical function. We now compare the Euclidean DT [Rousson et al. (2004)] version of the flow and more complex Riemannian DT [Lenglet et al. (2006a)] version of the flow with our proposed ODF flow.

8.4 SEGMENTATION RESULTS & DISCUSSION

ODF-based segmentations can deal with regions of complex fiber configurations. We first show that the ODF flow is able to propagate through regions of complex fiber crossings better than the DT-based flow using the Euclidean [Rousson et al. (2004)] and Riemannian distances [Lenglet et al. (2006a)]. We show comparison results on synthetic data, on a biological phantom and on real datasets. We also show that the corpus callosum and the corticospinal tract of the HARDI database [Poupon et al. (2006)] can be segmented automatically for all subjects and that there exist a variability in these structures across the subjects.

Visualization In this chapter, as always, DTs and ODFs are visualized as spherical functions stretched with respect to the surface values. However, the surfaces are colored according to the FA or GFA value, with color map going from red to blue for anisotropic to isotropic profiles. The more isotropic is the spherical function, the bluer it is and the more anisotropic is the spherical function, the redder it is. Greenish spherical functions are often locations where there are crossing fiber configurations. We use this same color map for both the DTI and ODF segmentations to allow easy visual comparisons between DTs and ODFs.

8.4.1 Synthetic Simulation Results

First, Figure 8.5 shows that initialization has a strong influence on the final segmented surface. If the initialization contains strictly anisotropic DTs/ODFs, the final surface is not able to pass through the fiber crossing area, as seen in Figure 8.5(a,b). Similarly, the final surface is trapped in the crossing area when initializing strictly in the 2-fiber region, as seen in Figure 8.5(c). This is because the statistics of the initial region have a large difference with the rest of the DTs/ODFs and hence, the evolving surface is blocked from connecting to the rest of the structure. However, if the initialization contains a mixture of both single fiber and 2-fiber DTs/ODFs, the DT flow propagates through the crossing region to connect to the similar anisotropic DTs on the other side of the crossing and the second fiber is completely ignored, as seen in Figure 8.5(d,e). The DTs in the crossing are oblate and there is no information on the second orientation. In contrary, there is information about the second orientation in the ODF flow and the surface evolution, seen in Figure 8.6, finds the whole 2-fiber structure as coherent.

Figure 8.7 and Figure 8.8 show more complex fiber configurations with a branching example and the “U”-fibers example. In the Riemannian DT flow, we see that the




Figure 8.5: Segmentation of the 90° crossing example. In (a-e), from left to right, the initialization used overlaid on the FA map, the DT Riemannian [Lenglet et al. (2006a)] segmentation and the ODF flow segmentation.




Figure 8.6: Evolution in time of the ODF flow on the 90° crossing. The flow is able to propagate through the fiber crossing area and recover the whole 2-fiber structure.

surface remains trapped in the regions of the initial seeding for all initializations. In contrary, in the ODF case, when the flow is initialized in the bottom and middle part of the branch, the whole branching structure is recovered because the ODF contains a broader range of orientations in its statistics. For all synthetic data experiments, DT Riemannian and DT Euclidean flows produced nearly identical qualitative segmentations and we decided only to show results for the DT Riemannian case. A factor of $\nu = 2$ for the smoothness of the surface was used for all flow evolutions.




Figure 8.7: Segmentation on a synthetic branching example. In (a-d), from left to right, the initialization used overlaid on the FA map, the DT Riemannian [Lenglet et al. (2006a)] segmentation and the ODF flow segmentation. The ODF flow is able to propagate through the fiber crossing area and recover the whole structure.

8.4.2 Biological Phantom Results

Figure 8.9 and Figure 8.10 show that the DT flow with the Euclidean distance is unable to segment the structure. The surface leaks outside the cords and the surface diverges because the isotropic DTs in the fibers and isotropic DTs in the background have a mean diffusivity in a similar range. Hence, the Euclidean distance cannot make the difference between the two regions. There is no convergence of the flow. This result was also obtained in [Lenglet et al. (2006a)]. However, our new ODF flow segments the whole structure quite easily. The segmentation also agrees with similar results published using the DT Riemannian flow [(Lenglet et al., 2006a, Figs.12-13)]. This similarity between the ODF flow and DT Riemannian flow can be explained because the crossing region is very small (roughly $3 \times 3 \times 2$ voxel region) and thus, the overall coherence and statistics of the desired structure are not dramatically affected by this crossing. Although the overall shape of the segmentations look the same in Figure 8.9, the actual zero level set of the segmenting surface in DT Riemannian and




Figure 8.8: Segmentation on a synthetic ‘U’-fiber example. In (a-d), from left to right, the initialization used overlaid on the FA map, the DT Riemannian [Lenglet et al. (2006a)] segmentation and the ODF flow segmentation. The ODF flow is able to propagate through the fiber crossing area and recover the whole structure.

ODF case are slightly different for some voxels at the border the structure, as seen in Figure 8.10. A factor of $\nu = 2$ for the smoothness of the surface was used for all flow evolutions.

8.4.3 Human Brain HARDI Data Results

We have segmented two large and well-known fiber bundles, the CC and the CST, on a human brain with b -value 1000 s/mm^2 and 60 sampling directions (MPI dataset of Appendix A.3.3) and on the public HARDI database (Appendix A.3.4) with a higher b -value 3000 s/mm^2 and 200 sampling directions. Our ODF segmentation on real datasets recovers more structure than other published results on the CC and CST [Zhukov et al. (2003); Rousson et al. (2004); Rousson (2004); Jonasson (2006); Lenglet et al. (2006a)]. First, we reproduced the results from [Lenglet et al. (2006a)] with the DT-based flows using both the Euclidean and Riemannian distances in Figure 8.11. Figure 8.11 and Figure 8.12 also show that in the DT Euclidean flow, the

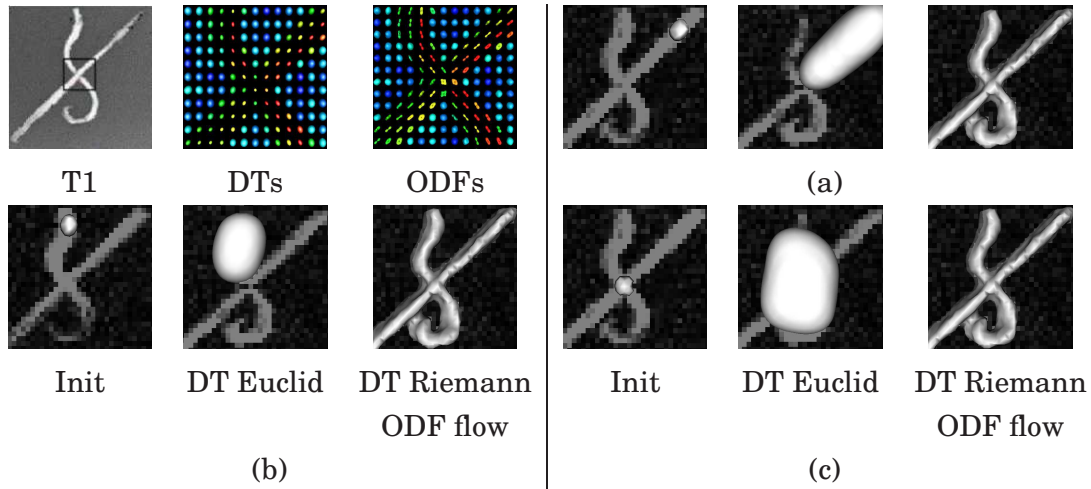


Figure 8.9: Segmentation on a biological phantom. In (a-c), from left to right, the initialization used overlaid on the FA map, the DT Euclidean [Rousson et al. (2004)] flow at $t = 40$ where we see the surface starting to leak outside the fiber structure and the segmentation of the ODF flow.




Figure 8.10: Evolution in time of the zero level set for the DT Euclidean [Rousson et al. (2004)] flow (green), DT Riemannian [Lenglet et al. (2006a)] flow (red) and our ODF flow (blue). The contour is placed over the inverted FA map.

evolving surface stops near complex crossing area where oblate and isotropic DTs (yellow-greenish DTs) block the flow. The DT Riemannian is able to connect more voxels than the DT Euclidean by slightly evolving into the crossing area. It is interesting to note, in the CST example (Figure 8.12), how the flow evolves as to go around complex fiber crossings, i.e. the evolution leaves the expected straight inferior-superior direction to go and pick up the projections of the CC overlapping with the CST and going to project to the superior cortex of the brain. However, in the CST, the flow is still unable to recover the branching fiber structure projecting to the cortex. The ODF flow recovers that branching structure to the different sulci and also recovers more of the posterior parts of the splenium of the CC.

We also observe that the ODF flow evolves more easily through the crossing area and does so after fewer iterations of the evolution. Convergence is obtained after 100




Figure 8.11: ODF flow segmentations can propagate through crossing regions and get a more accurate segmentation of the corpus callosum (CC). DT-based results from [Lenglet et al. (2006a)] are reproduced. The DT-based segmentations are overlaid on an axial slice with DTs and the ODF flow is overlaid on the same slice with the ODFs. The last row shows the evolution in time of the ODF flow with convergence after 100 iterations.

iterations in the CC and more rapidly after 60 iterations in the CST. It is important to point out that this convergence is obtained automatically without having to heuristically stop the surface evolution at a certain time t . At some point in the evolution, the background and foreground statistics stabilize, i.e. the mean and covariance matrix in the two regions become coherent and convergence is attained. Only a few voxels oscillate in and out of the regions \mathcal{F}_1 and \mathcal{F}_2 on the boundary of the evolving structure.² However, as it is often the case in surface evolution algorithms with a boundary term and/or a smoothness term we have to choose a certain value of ν . Here, a factor of

²This is seen in the surface evolutions of the CC and CST on the author's web page: <http://www-sop.inria.fr/odyssee/team/Maxime.Descoteaux/pages/seg.html>




Figure 8.12: ODF flow can propagate through crossing regions and get a more accurate segmentation of the cortico-spinal tract (CST). DT-based results from [Lenglet et al. (2006a)] are reproduced. The last row shows the evolution in time of the ODF flow with convergence after 60 iterations.

$\nu = 5$ was used for all flow evolutions shown. In practice, we observe that for ν too large the evolution stops prematurely and for ν too small the evolution leaks and connects the whole white matter. Fortunately, results are not extremely sensitive to ν and in our experiments, convergence is obtained for ν between 2 and 10 for both synthetic and real data. One still needs to try a few values of ν to obtain the best qualitative results.

8.4.4 Multi-Subject Study on a public HARDI Database

Figure 8.13 and Figure 8.14 show the ODF segmentations for the CC and CST of all subjects of the public HARDI database [Poupon et al. (2006)] (Appendix A.3.4). The same parameters were used for all 13 subjects from a single voxel in the medial part




Figure 8.13: Automatic segmentation of the corpus callosum (CC) using the ODF flow on the 13 subjects of the HARDI database [Poupon et al. (2006)] from a single seed point in the middle of the CC. Overall CCs are similar and we observe some variability across subjects.

of the CC and CST selected (manually) to initialize the flow. A factor of $\nu = 10$ for the smoothness of the surfaces was used for all flow evolutions on the database. It is thus possible to segment the datasets automatically and results are reproducible across many subjects.

Convergence was always obtained automatically for all subjects. Depending on the subject, 80 to 120 iterations of the flow were needed. An iteration takes roughly 1 second on a Dell single processor, 3.4 GHz, 2 GB RAM machine when $R = 15$. It is thus quite fast to obtain the segmentation for all subjects.

In Figure 8.13, we have segmented the complete CC for most subjects with the longer posterior parts of the splenium and the full genu, as in Figure 8.11. In the CST example of Figure 8.14, we have also obtained segmentation results that are in most cases as complete as the segmentation in Figure 8.12. For the CST example, we overlaid the segmented surface on the GFA to clearly see the intended white matter structure to be segmented and also to note the white matter structure differences across subjects. For all subjects, the GFA slice was always the 53rd coronal slice in voxel space.

In both the CC and CST, we note that some evolutions prematurely stopped near the crossing areas for some subjects (e.g. subjects 1, 11 and 13 for the CC and subjects 6 and 7 for the CST). One may think that playing with the segmentation parameters such as initialization and the smoothness ν of the surface might improve and change




Figure 8.14: Automatic segmentation of the cortico-spinal tract (CST) using the ODF flow on the 13 subjects of the HARDI database [Poupon et al. (2006)] from a single seed point in the middle of the CST. For this example, the segmented surface is overlaid on the 53rd coronal slice of the GFA map for all subjects to clearly see the multi-subject variability of the white matter. Overall CSTs are similar and we observe an important variability across subjects.

the results. In fact, in our experiments, this was not the case. A careful visual inspection of the QBI data for these subjects in regions where the flow stops shows two things: 1) Some subjects have “unusual” and “less complete” structures anatomically than others. For instance, subject 13 does not have a curving genu of the CC as all the other subject have. Hence, the CC recovered agrees with the anatomy of subject 13. 2) The ODFs are very isotropic in regions where the flow stops prematurely. The spherical shape of the ODFs are almost as isotropic as the DTs. We believe this is because the HARDI signal suffers from major partial signal averaging because of noise due to the high b -value acquisitions and multiple fibers crossing (up to three) in those crossing areas. Hence, there is no coherence between the crossing area and the rest of the desired CC/CST structure to segment. It is thus not surprising that the surface evolution stops. Why does it occur for these subjects and not for the others and why does it not occur for the dataset presented in Figures 8.11 and 8.12? Is it better to have a higher spatial resolution with less angular sampling of the sphere and smaller b -value or a higher angular resolution sampling density with a larger b -value at the cost of larger voxel size and lower SNR?

Overall, most CC and CST structures are similar and a significant variability across subject is observed. Hence, it is now important to better understand why these differences occur and propose ways to quantify this multi-subject variability.

8.5 DISCUSSION

We have presented a unified statistical surface evolution framework for the segmentation of HARDI data and in particular, for the segmentation of ODF images reconstructed from QBI. The proposed method combines state-of-the-art HARDI reconstruction and state-of-the-art region-based surface evolution. To answer questions from the introduction: 1) The segmentation problem on ODF images can be formulated efficiently with the level sets evolving to partition similar ODF based on their SH representation. This allows a computationally feasible similarity measure between ODFs based on the Euclidean distance between their SH representation. Note that this segmentation framework is general and could be used with any input vector of coefficients, for instance coefficients from other HARDI reconstruction methods such as persistent angular structure (PAS), spherical deconvolution (SD), diffusion orientation transform (DOT) or other high order spherical functions. 2) The ODF flow is able to deal with complex fiber configurations such as crossing and branching fibers better than DT-based segmentation using the Euclidean and Riemannian distances. 3) It is possible to validate the segmentation results. In particular, we obtained sets of globally coherent ODFs agreeing with well-known real data cerebral anatomical structures as well as with synthetic and biological phantom datasets where the ground truth was known. Moreover, compared to DTI segmentation, the ODF flow produces more complete segmentations of fiber bundles with crossings.

One might ask why we have not done the segmentation on a field of fiber ODFs, as for the tractography chapter. In our experience, a smoother ODF representation is more appropriate when looking for global statistics of ODFs in white matter. In practice, we find that the best results are obtained when using the diffusion ODF and not the fiber ODF. The segmentation applied on a field of fiber ODFs produce similar successful segmentations on synthetic datasets and on the biological phantom. However, on the real datasets, the fiber ODF flow is unstable and it is hard to find good convergence. We obtain more complete and accurate segmentations and obtain them with more stable convergence across subjects with the diffusion ODF flow. We find that the sharper peaks of the fiber ODF and the small perturbations around the origin of the fiber ODFs make the statistics between fiber ODFs within the same bundle further apart than if we use the diffusion ODF. Therefore, it seems that a smoother ODF representation is more appropriate when looking for a global coherence of white matter tracts.

There is room for improvement in our proposed segmentation algorithm. First, we believe that the search for good distance measures between spherical functions and orientation distributions is an important open problem in this field. Candidate met-

rics have been proposed in the literature such as the Kullback-Leibler distance but, as mentioned before, this measure works on the discrete spherical values and cannot be applied directly on the SH representation. Is it possible to propose a better distance measure between ODFs? In our opinion, the best distance measure should be computed directly on the SH coefficients for efficiency and should capture geometrical and shape properties of the ODF. It might also be interesting to investigate measures that take into account the diffusion properties of the original HARDI signal.

Another improvement to the algorithm could be made around the Gaussian distribution assumption of the SH coefficients in each region of the ODF image to be segmented. Is there a more appropriate parametric distribution that better describes the variability of the ODF in their SH representation? Is it possible to use a non-parametric probability distribution? The choice of a parametric family of functions to approximate distributions may be seen as a limitation to model region statistics. Contrary to parametric models, non-parametric distribution estimations can approximate any type of distribution for a sufficiently large dataset. In the scalar-valued and vector-valued images, people have estimated non-parametric distributions (see [Rousson (2004)] and references therein) of intensities to perform segmentation. These non-parametric methods have had some success but are much more involved computationally. Nonetheless, it might be worthwhile exploring this possibility.

Finally, a last improvement of the algorithm could be made regarding the initialization step. As for any image segmentation algorithm, initialization is an important problem. As seen in the result section, the initialization is crucial and will influence the final segmentation surface. Is it possible to find a more automatic and optimal initialization? This question is even more important if we intend to do multiple subject studies. Since we have an a priori knowledge about the white matter structure we are looking for, there is no reason why this cannot be done more automatically with atlases and appropriate image analysis tools. A first attempt to be able to do such a reproducible and automatic seeding is proposed in [Clayden et al. (2006, 2007)], where results from multiple subject probabilistic tracking are used to find consistent and reproducible segmentation of tracts.

Overall, the segmentation algorithm proposed in this chapter gives promising results. We have showed the reproducibility of the surface evolution on real datasets with different b -values and also on the 13 subjects from the public HARDI database. It is now important to focus on more difficult and small fiber bundles to test the limits of the algorithm. Nonetheless, it is possible to imagine performing a multi-subject study with segmented fiber bundles to quantify certain diffusion properties and attempt to follow the evolution of white matter diseases such multiple sclerosis, Parkinson, Alzheimer, etc...

8.5.1 Contributions of this chapter:

- Definition of an efficient and simple distance measure between spherical functions based on their SH representation.
- New statistical region-based segmentation framework for HARDI data segmentation.
- Efficient and semi-automatic implementation using the level set method.
- Validation and reproducibility of the results on a biological phantom and the public HARDI database.

Parts of the contributions from this chapter appear in [Descoteaux and Deriche (2007b,a)].

ACKNOWLEDGMENTS

We would like to thank Mikael Rousson and Christophe Lenglet for their previous work and implementation of the DTI segmentations algorithm with the Euclidean, Kullback-Leibler divergence and Riemannian distances. Thanks also to Cyril Poupon and Jean-François Mangin and colleagues from NeuroSpin / CEA for the HARDI database of 13 subjects.

IMPROVING Q-BALL IMAGING: FROM DIFFUSION ODF TO FIBER ODF ESTIMATE

The sharp thorn often produces delicate roses
—Ovid

Contents

9.1 Introduction	167
9.2 Spherical Deconvolution Using Spherical Harmonics	167
9.2.1 Spherical Deconvolution of the q-Ball diffusion ODF	168
9.2.2 Spherical Deconvolution of the raw HARDI Signal	170
9.3 Evaluation of the Spherical Deconvolution Methods	178
9.3.1 Synthetic Data Experiment	178
9.3.2 Real Data Experiment	179
9.4 Results	179
9.4.1 Effect of Varying Signal Generation Parameters	180
9.4.2 Fiber Detection and Angular Resolution	184
9.4.3 Robustness to Noise	186
9.4.4 Real Data Spherical Deconvolution Reconstructions	188
9.5 Discussion	194
9.5.1 Contributions of this chapter:	197
9.6 Appendix A: Diffusion ODF Kernel for Sharpening	198
9.7 Appendix B: Coefficients of A in the FORECAST Solution	199
9.8 Appendix C: Relation between the fODF and FORECAST	200

OVERVIEW

Can we improve angular resolution of QBI? This chapter covers the problem of reconstructing a fiber ODF and not only a diffusion ODF. In particular, can the diffusion ODF be transformed into a sharper fiber ODF? Can this sharper fiber ODF be obtained from QBI? Can it be related to the raw HARDI signal and other spherical deconvolution methods that also estimate a fiber ODF? First, we develop a new spherical deconvolution operation that transforms the q-ball diffusion ODF into a sharper fiber ODF. This deconvolution transformation improves angular resolution and fiber detection of QBI by approximately 15° . Then, our new fiber ODF is shown to behave similarly to the fiber ODF estimated from the filtered, regularized and constrained spherical deconvolution methods of [Tournier et al. (2004, 2007); Sakaie and Lowe (2007)] as well as the fiber ODF estimated with FORECAST [Anderson (2005)]. Next, we extensively compare these methods quantitatively by varying the harmonic order ℓ , the b -value, the sampling density N , the noise level (SNR), the separation angle between fibers and the volume fraction of each fiber. We also evaluate the robustness to noise and the angular resolution of each method. Finally, all fiber ODFs are reconstructed and compared on the biological phantom and the real datasets.

Keywords diffusion ODF, fiber ODF, spherical deconvolution (SD), spherical harmonics (SH), Funk-Hecke theorem

Contributions of this chapter:

- New deconvolution sharpening operation transforming the q-ball diffusion ODF into a sharper fiber ODF.
- Derivation, implementation and comparison of the state-of-the-art spherical deconvolution methods based on the spherical harmonic basis.
- Extensive validation study of our new fiber ODF from numerical simulations, on the biological phantom and on the real human brain datasets.

Organization of this chapter:

The chapter is organized as follows. We motivate the spherical deconvolution of QBI data in Section 9.1. We develop the analytical spherical deconvolution operation that transforms the q-ball diffusion ODF into a fiber ODF in Section 9.2.1 and we describe the relationship and our implementation of the other state-of-the-art spherical deconvolution methods in Section 9.2.2. Then, Section 9.3 describes the different numerical simulations used to compare and validate the fiber ODF reconstruction methods. Results of the comparison between fiber ODF methods are presented in Section 9.4 on the synthetic simulations, on the biological rat phantom and on the human brain datasets. We conclude with a discussion of the results and our contributions in Section 9.5.

9.1 INTRODUCTION

A current open problem that we have not discussed so far in the thesis is the relation between the measured diffusion ODF and the underlying fiber distribution, the fiber ODF. This relation between diffusion and fiber ODF is still unknown today [Tuch (2002); Perrin et al. (2005)] even though we know that the relation between the diffusion of the water molecule and the underlying fiber distribution can depend on the physics of diffusion, the cell membrane permability, the free diffusion coefficients, the axonal packing, the distribution of axonal diameters and the degree of myelination in the underlying fiber bundles. As described in Chapter 4, spherical deconvolution methods simplify the problem by assuming simple mixture models to reconstruct an estimation of the true fiber ODF, also called a fiber orientation density (FOD) [Tournier et al. (2004); Alexander (2005b); Anderson (2005); Ramirez-Manzanares et al. (2007); Dell’Acqua et al. (2007); Kaden et al. (2007); Jian and Vemuri (2007b)]. These spherical deconvolution methods work directly on the raw HARDI signal, without the need of the diffusion ODF. However, anybody working with QBI and diffusion ODFs such as [Tuch (2004); Khachaturian et al. (2007); Perrin et al. (2005); Hagmann et al. (2006a); Bergmann et al. (2006); Jonasson et al. (2007); Haroon and Parker (2007); Fonteijn et al. (2007)] has the problem of dealing with smooth ODFs that have a large diffusion part outside the principal fiber directions (see Figure 9.2). Hence, it is important to investigate SD operations that can work on the diffusion ODF and also to study its relation with the other spherical deconvolution methods working directly on the HARDI signal.

In this chapter, we propose a new ODF sharpening deconvolution operation that has the desired effect of transforming the q-ball diffusion ODF into a sharp fiber ODF. Again, the solution is based on the spherical harmonic representation of the HARDI signal. Hence, we extensively compare our new fiber ODF estimation with the other spherical deconvolution methods based on spherical harmonics. We implement the fiber ODF estimation from the classical filtered SD (fSD) [Tournier et al. (2004)] and also compare against the more recent fiber orientation estimated using continuous axially symmetric tensors (FORECAST) [Anderson (2005)], the Tikhonov regularized SD [Sakaie and Lowe (2007) and the very recent Constrained SD (CSD) and super-resolution CSD (super-CSD) of [Tournier et al. (2007)]. These methods are all based on the spherical harmonics representation. We show that our fiber ODF is a valid choice to obtain a stable fiber ODF estimation. We also show that our new fiber ODF behaves similarly to the other SD methods studied.

9.2 SPHERICAL DECONVOLUTION USING SPHERICAL HARMONICS

In this section, we describe our new method to perform spherical deconvolution




Figure 9.1: Spherical deconvolution illustration as in [(Tournier et al., 2004, Fig.1)]. The measured HARDI signal S is the convolution of the single fiber response function R and the fiber distribution.

of the q-ball diffusion ODF. We also describe our implementation of the FORECAST [Anderson (2005)], the fSD [Tournier et al. (2004)], the Tikhonov regularized SD [Sakaie and Lowe (2007)] and the CSD and super-CSD [Tournier et al. (2007)] methods. All these methods are similar in spirit to our approach and are all based on the spherical harmonic representation of the HARDI signal.

9.2.1 Spherical Deconvolution of the q-Ball diffusion ODF

We describe the sharpening operation that aims to transform the smooth diffusion ODF into a sharp fiber ODF. The sharpening operation is a simple linear transformation of the SH coefficients describing the diffusion ODF. Our approach is inspired by the original spherical deconvolution approach proposed in [Tournier et al. (2004)] and uses a diffusion model for the single fiber similar to that of [Anderson (2005)]. In SD methods, the measured HARDI signal S is expressed as the convolution on the unit sphere of the fiber response function R with the fiber distribution, as illustrated in Figure 9.1. Assuming a particular fiber response function representing the diffusion signal attenuation that would be measured for a single fiber, the deconvolution of the signal with the response function gives an estimation of the fiber distribution. Here, we call this fiber distribution estimation the fiber ODF.

The starting point of our deconvolution method is not the measured signal but the estimated q-ball diffusion ODF described in Chapter 7. We want to deconvolve the estimated smooth diffusion ODF by the diffusion ODF for a single fiber in order to obtain a sharpened fiber ODF. The procedure is sketched in Figure 9.2. We assume that the estimated diffusion ODF, Ψ , is formed by the convolution between the single fiber diffusion ODF kernel, R' , and the true fiber ODF F . Note that we use R' not to confuse with the single fiber response function R used in Figure 9.1. Hence, the deconvolution of the diffusion ODF by the diffusion ODF kernel can recover an estimation of the true fiber ODF. In this chapter, we make use of the term *fiber ODF* for the estimation of the true fiber ODF. We abbreviate our new fiber ODF with fODF and use Ψ_{sharp} in the equations.

The convolution on the sphere between the single fiber diffusion kernel R' and

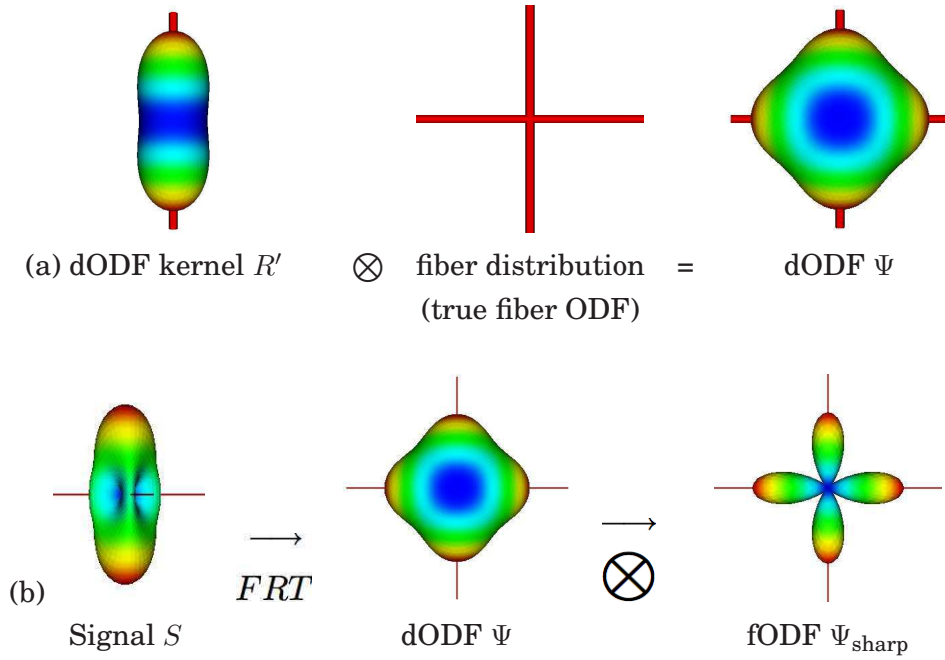


Figure 9.2: Sketch of the convolution/deconvolution. In (a), the convolution between the dODF kernel and true fODF produces a smooth dODF. In (b), we show a sketch of the deconvolution sharpening. The Funk-Radon Transform (FRT) of the simulated HARDI signal on the sphere produces a smooth dODF. This dODF is transformed into a sharp fODF by the deconvolution with the dODF kernel of (a). The simulated HARDI signal was generated with noise-free orthogonal fibers crossing each with a FA = 0.7, with b -value of 3000 s/mm², spherical sampling density $N = 60$ and an order $\ell = 6$ reconstruction.

Ψ_{sharp} (Figure 9.2a) can be written as

$$\Psi(\mathbf{u}) = \int_{|\mathbf{w}|=1} R'(\mathbf{u} \cdot \mathbf{w}) \Psi_{\text{sharp}}(\mathbf{w}) d\mathbf{w}. \quad (9.1)$$

In order to solve this integral, we first replace Ψ and Ψ_{sharp} with their respective SH estimation of order ℓ , $\Psi = \sum_j c'_j Y_j(\mathbf{u})$ and $\Psi_{\text{sharp}} = \sum_j f_j Y_j(\mathbf{u})$ and obtain

$$\sum_{j=1}^R c'_j Y_j(\mathbf{u}) = \sum_{j=1}^R f_j \int_{|\mathbf{w}|=1} R'(\mathbf{u} \cdot \mathbf{w}) Y_j(\mathbf{w}) d\mathbf{w}. \quad (9.2)$$

At this point, we use the Funk-Hecke theorem stated in Section 5.5 to solve the convolution integral between R' and the spherical harmonic Y_j over the sphere. We then obtain

$$f_j = \frac{c'_j}{r'_j}, \quad \text{where} \quad r'_j = 2\pi \int_{-1}^1 P_{\ell(j)}(t) R'(t) dt, \quad (9.3)$$

where coefficients r'_j come from the Funk-Hecke formula.

The main consideration is thus the creation of a viable single fiber diffusion ODF kernel R' . As done in [Anderson (2005)], we assume an axially symmetric tensor

model of diffusion for a single fiber. Hence, we choose a prolate profile with eigenvalues $[\lambda_2, \lambda_2, \lambda_1]$ such that $\lambda_1 \gg \lambda_2$ to represent this tensor. Then, using the analytical relation between diffusion tensor and diffusion ODF, we obtain the diffusion ODF kernel R' as

$$R'(t) = \frac{1}{Z'} \frac{1}{\sqrt{(\lambda_2/\lambda_1 - 1)t^2 + 1}}, \quad (9.4)$$

where Z' is a normalization constant so that the integral over all directions is 1. In the case of an axially symmetric tensor model, $Z' = 8\pi b\sqrt{\lambda_2^2\lambda_1}$. Details can be found in Appendix 9.6. In the end, the SH coefficients of the new fiber ODF can be reconstructed with

$$f_j = \frac{4b\sqrt{\lambda_2\lambda_1}}{A'_{\ell(j)}(1 - \lambda_2/\lambda_1)} c'_j, \quad (9.5)$$

where $A'_{\ell(j)}(\alpha) = \int_{-1}^1 (1 - \alpha t^2)^{-1/2} P_{\ell(j)}(t) dt$ are given analytically in Appendix 9.7 up to order 8.

Therefore, the final sharp fiber ODF expression is computed with

$$\Psi_{\text{sharp}}(\mathbf{u}) = \sum_{j=1}^R f_j Y_j(\mathbf{u}), \quad (9.6)$$

for any direction \mathbf{u} .

9.2.2 Spherical Deconvolution of the raw HARDI Signal

We now describe in detail the SD methods working directly on the raw HARDI signal. The convolution on the sphere of Figure 9.1 between the single fiber response function R and signal S can be written in two different ways. First, the original formulation of [Tournier et al. (2004)] is

$$\mathbf{S}^n = \mathbf{R}^n \mathbf{F}^n, \quad (9.7)$$

where \mathbf{S}^n and \mathbf{F}^n are vectors of length $(2n + 1)$ of spherical harmonic coefficients representing the signal and fiber ODF respectively and \mathbf{R}^n is a $(2n + 1) \times (2n + 1)$ matrix representing the rotational harmonic decomposition of the axially symmetric response function R . Rotational harmonics are defined and given in [(Healy et al., 1998, Eq. 2.1)].

The second formulation is the approach by [Anderson (2005)]. The SD can be written as

$$S(\mathbf{u}) = \int_{|\mathbf{w}|=1} R(\mathbf{u} \cdot \mathbf{w}) F(\mathbf{w}) d\mathbf{w}, \quad (9.8)$$

where F is the fiber ODF and R is assumed to be the response function arising from an axially symmetric diffusion tensor model. We now describe the implementation of these two methods and relate them to our sharp fiber ODF described in the previous section. Even though the second formulation appeared after the original SD formulation of Tournier et al, we start with the second because it is closest to our method described in the previous section.

FORECAST Spherical Deconvolution

FORECAST [Anderson (2005)] is similar in spirit to our deconvolution method. It also reconstructs an estimation of the true fiber ODF from a linear transformation of the SH coefficients describing the raw HARDI signal. As now seen many times in this thesis, the signal S and fiber ODF F can be expressed in SH series of order ℓ as $S(\theta, \phi) = \sum_j c_j Y_j(\theta, \phi)$ and $F(\theta, \phi) = \sum_j f_j Y_j(\theta, \phi)$. Hence, the integral of Eq. 9.8 becomes

$$\sum_{j=1}^R c_j Y_j(\mathbf{u}) = \sum_{j=1}^R f_j \int_{|\mathbf{w}|=1} R(\mathbf{u} \cdot \mathbf{w}) Y_j(\mathbf{w}) d\mathbf{w}. \quad (9.9)$$

At this point, we use a different derivation from [Anderson (2005)]. Our derivation is equivalent but uses the Funk-Hecke theorem to solve the SD integral, as done in our solution of the previous section. We then obtain

$$f_j = \frac{c_j}{r_j}, \quad \text{where} \quad r_j = 2\pi \int_{-1}^1 P_{\ell(j)}(t) R(t) dt, \quad (9.10)$$

where coefficients r_j come from the Funk-Hecke formula. R is chosen to be an axially symmetric diffusion tensor \mathbf{D} with eigenvalues $[\lambda_2, \lambda_2, \lambda_1]$ with $\lambda_1 \gg \lambda_2$, which can be written as

$$R(t) = S_0 e^{-b(\lambda_2 + t^2(\lambda_1 - \lambda_2))}. \quad (9.11)$$

This is the standard signal attenuation formula of Appendix A.1, which we have re-written for an axially symmetric fiber along the z -axis using similar simplifications as in Appendix 9.6 (see also [Anderson (2005)]).

Therefore, the SH coefficients of the fiber ODF from the FORECAST reconstruction are given by

$$f_j = \frac{e^{b\lambda_2}}{2\pi S_0 A_{\ell(j)}(b(\lambda_1 - \lambda_2))} c_j := \text{FORECAST}_j, \quad (9.12)$$

where $A_{\ell(j)}(a) = \int_{-1}^1 e^{-at^2} P_{\ell(j)}(t) dt$ are given analytically in Appendix 9.7 up to order 8.

Recall the coefficients of our sharp fiber ODF in Eq. 9.5. They are written in terms of the q-ball diffusion ODF SH coefficients c'_j . From the last chapter, we know that these coefficients can be written in terms of the HARDI signal S/S_0 represented in SH coefficients c_j , namely $c'_j = (2\pi/S_0) P_{\ell(j)}(0) c_j$ (Eq. 7.10). Hence, Eq. 9.5 can be re-written as

$$f_j = \frac{8\pi b \sqrt{\lambda_2 \lambda_1} P_{\ell(j)}(0)}{S_0 A'_{\ell(j)}(1 - \lambda_2/\lambda_1)} c_j := \text{fODF}_j. \quad (9.13)$$

This expression is similar to the FORECAST solution of Eq. 9.12 but they are not equal. However, in Appendix 9.8 we show that they have the same behavior. It turns out that the normalized spherical function projected back to the sphere from fODF_j and FORECAST_j are the same up to small numerical errors propagating in the high order harmonics (see Appendix 9.8 and Figures 9.17 and 9.18) on a noise-free example.

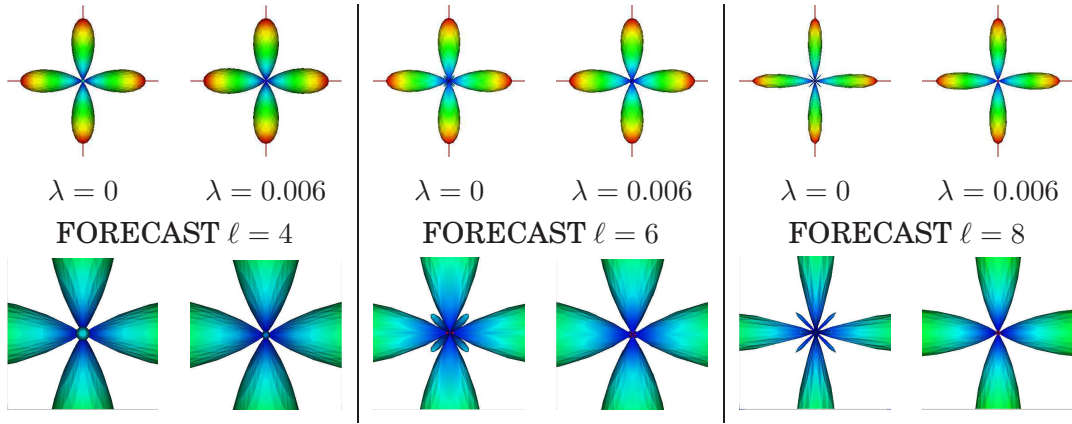


Figure 9.3: Effect of varying estimation order ℓ on the FORECAST reconstruction with and without ($\lambda = 0$) Laplace-Beltrami regularization incorporated in the signal estimation. Signal is noise-free and generated with two orthogonal fibers each with a FA = 0.7, $b = 3000 \text{ s/mm}^2$ and $N = 60$ data points.

Implementation: In addition to the axially symmetric tensor assumption, the original FORECAST [Anderson (2005)] also assumes uniform diffusivity in the white matter tissue, i.e. it assumes that the mean diffusivity ($\bar{e} = (\lambda_1 + 2\lambda_2)/3$) is constant in all the white matter. Hence, the FORECAST solution can be reformulated only in terms eigenvalue λ_2 and \bar{e} , which are estimated from the real data. Moreover, the SH coefficients representing the signal, the c_j coefficients, are estimated using a simple least-squares method, as described in Section 5.4 and Eq. 5.29.

In the rest of the chapter, we implement an extended and improved version of FORECAST based on previous work presented in this thesis. First, we do not assume uniform diffusivity in the white matter and estimate λ_1 and λ_2 directly from our real datasets from voxels where we have high confidence that they contain only a single fiber population (typically voxels of the corpus callosum). Second, we use a Laplace-Beltrami regularization incorporated in the least-squares solution to solve for the c_j coefficients representing the signal. This regularization of Eq. 6.13 has showed success in our ADC and q-ball diffusion ODF estimation in the earlier chapters. Here, it also improves FORECAST reconstructions. To show this, Figure 9.3 shows a noise-free example with and without the regularization in the estimation of c_j coefficients. As expected, spurious peaks and lobes around the center of the crossing area are eliminated with the Laplace-Beltrami λ -regularization. Perturbations of the fiber ODF arising numerical instabilities of the deconvolution occurring in the high order SH coefficients are reduced. Therefore, we are consistent with our fiber ODF implementation and previous methods proposed in this thesis.

Filtered Spherical Deconvolution

The original formulation of SD was proposed in [Tournier et al. (2004)]. Here, the f_j coefficients are the unknowns. We need to estimate the rotational harmonic matrices for each even order n . Due to the axial symmetry assumption of the response function, only the rotational harmonic coefficients along the diagonal are non-zero. Hence, each \mathbf{R}^n collapse down to a single real scalar constant, which can be evaluated using the rotational harmonic expression given in [Healy et al. (1998)]. As pointed out in [(Tournier et al., 2007, Appendix A)], the actual rotational harmonics do not need to be computed explicitly. If we imagine that the fiber ODF is a spherical Dirac function aligned with the z -axis, then the signal S measured has to be identical to the response function R , i.e. if the f_j coefficients correspond to a Dirac delta function, then the c_j coefficients representing the signal must correspond to the response function. In this case, as both functions are axially symmetric, all coefficients corresponding to $m \neq 0$ coefficients are zero. Therefore, if the Dirac delta function δ and response function R are expressed in a SH series as $\delta(\theta, \phi) = \sum_j \delta_j Y_j^0(\theta, \phi)$ and $R(\theta, \phi) = \sum_j r_j Y_j(\theta, \phi)$, the rotational harmonic coefficients can be computed by simple scalar division of the n th order, $m = 0$ SH coefficients of R by the corresponding coefficient for the Dirac delta function. Note that there is also a simple expansion for the SH coefficients of the spherical Dirac function [Bulow (2004)]. The δ_j 's are given by

$$\delta_j = \sqrt{\frac{2\ell+1}{4\pi}} Y_{\ell(j)}^0. \quad (9.14)$$

Hence, we can solve for the fiber ODF without the need to compute the actual rotational harmonics. The SD Eq. 9.7 becomes

$$f_j = R_{jj}^{-1} c_j, \quad (9.15)$$

where $R_{jj} = r_j/\delta_j$ are the diagonal entries of matrix \mathbf{R} . Finally, the reconstructed fiber ODF on the sphere is computed as before with

$$F(\mathbf{u}) = \sum_{j=1}^R f_j Y_j(\mathbf{u}), \quad (9.16)$$

for any direction \mathbf{u} and f_j given in Eq. 9.15.

Implementation: We implement the fSD [Tournier et al. (2004)] because it reduces the effects of noise and it has been extensively studied in the literature. Noise can introduce spurious peaks and negative values on the sphere. So, we implement the fSD using a low-pass filter $[1, 1, 1, 0.8, 0.1]$ multiplying each coefficient of order $\ell \in \{0, 2, 4, 6, 8\}$ respectively, as used in [Tournier et al. (2004, 2007); Sakaie and Lowe (2007)]. That is, terms of order 0, 2 and 4 are not altered, the terms of order 6 are multiplied by 0.8 and the terms of order 8 are multiplied by 0.1. The SH c_j coefficients representing the signal are estimated using a standard least-squares solution

as in Eq. 5.29. In the next sections, we describe other possible regularization of the deconvolution procedure to reduce negative values and spurious peaks.

Regularization of Spherical Deconvolution Methods

fSD proposes a heuristic low-pass filtering of the high order harmonics. Alternatively, [Sakaie and Lowe (2007)] propose to use a gradient-based Tikhonov regularization procedure. As mentioned in the previous two chapters, this regularization is similar in spirit to our Laplace-Beltrami regularization. Intuitively, the gradient of a spherical function is large at the peaks of the function. Hence, a possible choice of constrained function E is the squared norm of the gradient of the fiber ODF for all directions on the sphere, i.e.

$$E(F(\theta, \phi)) = \int_{\Omega} |\nabla F(\theta, \phi)|^2 d\Omega = \sum_{j=1}^R \ell(j)(\ell(j) + 1) f_j, \quad (9.17)$$

which again simplifies nicely due to the orthonormality of the SH basis (Eq. 5.19). Note that our Laplace-Beltrami regularization (Eq. 6.10) is the squared version of this gradient-based approach. Then, as before, the quantity we wish to minimize can be expressed in matrix form as

$$M(\mathbf{F}) = (\mathbf{S} - \mathbf{B}\mathbf{R}\mathbf{F})^T(\mathbf{S} - \mathbf{B}\mathbf{R}\mathbf{F}) + \lambda \mathbf{F}^T \mathbf{L} \mathbf{F}, \quad (9.18)$$

where \mathbf{B} is the SH basis matrix, \mathbf{S} is the $N \times 1$ discrete HARDI signal on the sphere, \mathbf{F} is the $R \times 1$ SH representation of the fiber ODF, \mathbf{R} is the $R \times R$ rotational harmonic decomposition of the response function, λ the regularization weight and \mathbf{L} is the matrix from the gradient-based approach with $\ell(j)(\ell(j) + 1)$ on the diagonal or the Laplace-Beltrami based regularization with $\ell(j)^2(\ell(j) + 1)^2$ on the diagonal. Note that it is also possible to use a zeroth order Tikhonov regularization with $\mathbf{L} = \mathbf{I}$, the identity matrix, as described in the previous chapter and in [Hess et al. (2006); Descoteaux et al. (2007a); Jian and Vemuri (2007b)].

The coefficient vector minimizing the expression in Eq. 9.18 can then be determined by setting each of the $\partial M / \partial f_j = 0$, from which we obtain the generalized expression for the desired SH series coefficient vector of the fiber ODF \mathbf{F} as

$$\mathbf{F} = ((\mathbf{B}\mathbf{R})^T(\mathbf{B}\mathbf{R}) + \lambda \mathbf{L})^{-1} (\mathbf{B}\mathbf{R})^T \mathbf{S}. \quad (9.19)$$

For the rest of this chapter, let LB-SD and GB-SD represent the Laplace-Beltrami (LB) and Gradient-Based (GB) regularized spherical deconvolution respectively.

Figure 9.4 compares fSD, GB-SD and LB-SD reconstructions on a very noisy synthetic simulation with SNR 10, sampling density $N = 60$ and $b = 3000$ s/mm². We also vary the estimation order between $\ell = 4, 6$ and 8. We see that fSD and the regularized GB-SD and LB-SD behave similarly and have the desired effect of reducing noise in the high order harmonic coefficients. Although not shown here, on higher




Figure 9.4: Effect of varying estimation order ℓ on the unfiltered SD, the fSD, the GB-SD and the LB-SD. Signal is generated with two orthogonal fibers each with FA = 0.7, $b = 3000$ s/mm², $N = 60$ data points and high noise level of SNR 10. We fix $\lambda = 5 \times 10^{-3}$ for GB-SD and $\lambda = 5 \times 10^{-5}$ for LB-SD. The opaque surface corresponds to the mean FOD over 100 noise trials, whereas the transparent surface corresponds to the mean + 2 standard deviations. Blue and red lines correspond to ground truth fiber directions and detected maxima respectively.

SNR simulations (SNR > 10) all three regularization methods give more or less the same fiber ODF shape. However, in this very noisy simulation, we see that both LB-SD and GB-SD reduce spurious peaks better than fSD. Theoretically, the LB regularization method imposes a stronger weight on the high order harmonics than the GB regularization and is thus more aggressive. A more aggressive regularization can be better suited when the signal is very noisy. In our experiments, we observe that the regularization parameter λ can be adapted so that both the LB-SD and GB-SD methods essentially give the same level of regularization and a very similar fiber ODF shape. Here, an optimal λ was used and computed with the L-curve method [Hansen (2001)], as in Chapter 6 (the generalized cross validation method [Wahba (1990)] is used in [Sakaie and Lowe (2007)]). It is not the purpose of this chapter to perform an extensive quantification and assessment of the advantages and disadvantages of both the LB and GB or Tikhonov SD regularization. It will be part of future work to do so. We refer the interested reader to [Jian and Vemuri (2007b)] for more on this issue.

Therefore, both the LB-SD and GB-SD are well-defined regularization procedures that can be used to replace the heuristic low-pass filtering used in fSD. For the rest of the chapter, we more often choose to compare against the classical fSD [Tournier et al. (2004)] because it has been extensively studied in the literature and also used for comparisons in [Sakaie and Lowe (2007); Tournier et al. (2007)].

Constrained and Super-Resolution Spherical Deconvolution

In our experience, negative values do not appear in large numbers when performing a low order reconstruction such as reconstructions at order $\ell = 4$ and $\ell = 6$. Moreover, we find that a large number of spurious negative peaks are reduced with our Laplace-Beltrami regularization of the signal (estimation of the c_j as described in Section 6.2.3) as well as when incorporating a regularization of the SD with a Tikhonov, gradient-based or Laplace-Beltrami method, as described in the previous section. This is also observed by [Jian and Vemuri (2007b)] with a classical Tikhonov regularization. Nonetheless, if one wants high order reconstructions with $\ell \geq 8$, such as reconstructions in [Tournier et al. (2007)] up to order $\ell = 14$, negative values and noise are very problematic. It is also physically meaningless to have negative values on the sphere and it does make sense to want to suppress them just as negative eigenvalues need to be suppress in diffusion tensor imaging [Tschumperlé and Deriche (2001); Ched'hotel et al. (2004); Arsigny et al. (2006); Lenglet et al. (2006b)].

Here, we describe the iterative approach of [Tournier et al. (2007)], which imposes a non-negativity constraint on the reconstructed fiber ODF using spherical harmonics. The idea is to look for an improved fiber ODF, using an optimization method that incorporates a term in which the fiber ODF is assumed to be zero in directions where the current estimate of the fiber ODF is negative. The problem is formulated as follows

$$\mathbf{F}_{i+1} = \arg \min \{ \|\mathbf{BRF}_i - \mathbf{S}\|^2 + \lambda^2 \|\mathbf{LF}_i\|^2 \}, \quad (9.20)$$

where as before \mathbf{B} is the SH basis matrix, \mathbf{S} is the HARDI signal on the sphere, \mathbf{F}_i is the SH representation of the fiber ODF at iteration i , \mathbf{R} is the rotational harmonic decomposition of the response function and λ the regularization weight. It is the \mathbf{L} matrix that is different from the other formulations we have seen so far. \mathbf{L} is formulated as

$$\mathbf{L}_{mn} = \begin{cases} (\mathbf{BR})_{mn} & \text{if } (\mathbf{BRF})_m < \tau \\ 0 & \text{otherwise,} \end{cases} \quad (9.21)$$

where τ is a user-specified threshold controlling the amplitude of the current fiber ODF under which the fiber ODF is assumed to be zero. Hence, \mathbf{L} provides the amplitude of the current estimate of the fiber ODF for the set of directions along which the amplitude fiber ODF is assumed to be zero. Hence, a stronger weight will be imposed on the coefficients of \mathbf{F}_i that produce the negative values on the sphere.

Implementation: The initial estimate of the fiber ODF \mathbf{F}_0 is computed from standard fSD of order ℓ and then the iterative procedure is repeated until the matrix \mathbf{L} does not change, which typically takes between 5 to 10 iterations for a wide range of imaging parameters [Tournier et al. (2007)]. We have implemented the CSD version with our modified SH basis framework and thus, cannot we certain that our implementation is the same as in [Tournier et al. (2007)].




Figure 9.5: Comparing the q-ball dODF, the fSD, the CSD and the super-CSD. Signal generated with two fibers with 60° separation each with $FA = 0.7$, with $b = 1000 \text{ s/mm}^2$, $N = 60$ data points and SNR 30. Results agree with simulations of [(Tournier et al., 2007, Fig.4)]. Note that we cannot compute fSD and CSD for $\ell \geq 10$ because in these cases, there are more unknown parameters to estimate than there are measurements. The opaque surface corresponds to the mean FOD over 100 noise trials, whereas the transparent surface corresponds to the mean + 2 standard deviations. Blue and red lines correspond to ground truth fiber directions and detected maxima respectively.

Tournier et al point out that it is possible to reconstruct high order fiber ODF even when the number of SH coefficients in the representation of \mathbf{F} is larger than the actual number of data points N of the signal \mathbf{S} . For example, if we have a spherical sampling scheme with $N = 60$ data points and seek an order $\ell \geq 10$ reconstruction, where there are $R \geq 66$ SH coefficients to estimate, the normal formulation of Eq. 9.20 gives an under-determined system of equations. To overcome this problem, Tournier et al rewrite Eq. 9.20 as

$$\mathbf{F}_{i+1} = \arg \min \left\{ \left\| \begin{pmatrix} \mathbf{BR} \\ \lambda \mathbf{L} \end{pmatrix} \mathbf{F}_i - \begin{pmatrix} \mathbf{S} \\ 0 \end{pmatrix} \right\|^2 \right\} \quad (9.22)$$

and thus, one notes that the system is over-determined as long as the number of rows in \mathbf{L} , i.e. the number of negative fiber ODF values on the sphere, is always greater than $|R - N|$, where R is number of SH coefficients in \mathbf{F}_i and N is number of data points in \mathbf{S} . This solution is called super-resolution CSD (super-CSD) [Tournier et al. (2007)].

Figure 9.5 reproduces the behavior of fSD, CSD and super-CSD satisfactorily on a synthetic signal example generated with $N = 60$ data points, $b = 1000 \text{ s/mm}^2$ and SNR 30. In contrast to standard fSD, very high order estimations are possible while reducing effects of noise and spurious peaks in the reconstruction. Moreover, note that for $\ell \geq 10$, $R > 60$ and there are more unknowns than measurements in the CSD formulation. However, super-CSD is able to recover a stable fiber ODF that becomes sharper as order ℓ increases.

The CSD and super-CSD methods reduce negative values on the sphere but do not force their elimination completely. The argument of [Tournier et al. (2007)] is that imposing a very strict non-negative constraint with high λ results in the reduc-

tion of the angular resolution of the fiber ODF, which for tractography application is not desirable. Moreover, Tournier et al note that the fiber ODF expressed with a truncated SH series is expected to contain some negative lobes due to the spherical equivalent of Gibb’s ringing observed in classical Fourier analysis.¹ Hence, Tournier et al prefer to impose a softer regularization weight λ . Nonetheless, a very recent method of [Jian and Vemuri (2007b)], eliminates the negative values completely with a non-negative least squares optimization problem stated as

$$\min\{||\mathbf{BRF} - \mathbf{S}||^2\} \quad \text{subject to} \quad \mathbf{BF} \geq 0, \quad (9.23)$$

which is solved as a quadratic programming problem. After the minimization, the fiber ODF estimation, \mathbf{BF} , cannot have negative values on the sphere. It is part of future work to implement this non-linear method and compare it with CSD and super-CSD. It is of interest to investigate the difference made on the fiber ODF reconstructions under soft and strict non-negativity constraints and really decide if negative values on the sphere are desirable or not.

9.3 EVALUATION OF THE SPHERICAL DECONVOLUTION METHODS

In the synthetic simulations of [Tournier et al. (2004); Anderson (2005); Dell’Acqua et al. (2007)], values of $[\lambda_2, \lambda_2, \lambda_1]$ are chosen from physiological data. Here, we prefer to estimate the single fiber diffusion ODF kernel and the single fiber response function directly from our real datasets, as also done in [Tournier et al. (2004, 2007)] for the single fiber response function. In practice, the average prolate tensor profile is estimated from 300 voxels with the highest FA values, as these tensors can each be assumed to contain a single fiber population. This average prolate tensor profile is used for values of λ_2 and λ_1 in the fODF and FORECAST. For fSD, CSD and super-CSD, the average prolate tensor is used to generate DW signal and a SH series of order ℓ is fitted to this signal to obtain the SH coefficients of the response function R (the r_j coefficients in Eq. 9.15).

9.3.1 Synthetic Data Experiment

First, we show qualitative reconstruction results of the dODF, the sharp fODF, the fSD, the FORECAST and the CSD from synthetic signal simulations varying the b -value, SNR, separation angle between fibers and fiber volume fraction in Eq. A.1 of Appendix A.1. We use a prolate tensor profile of $[3.55, 3.55, 13.9] \times 10^{-4}$ mm²/s, as estimated from the MPI dataset.

¹In the Cartesian case, the Dirac delta function expressed as a truncated Fourier series becomes a sinc function which has some negative lobes. Similarly, the spherical Dirac delta function expressed as a truncated SH series also contains small negative lobes. Hence, it seems that completely eliminating negative lobes is not desired and smoothes the reconstructions too much.

Second, we concentrate on the dODF, the sharp fODF and fSD to quantify the angular resolution limitation, the fiber detection success and angular error made on the detected maxima. To evaluate angular resolution limitations, we generate noise-free synthetic data for two fibers where we vary the crossing angle between fibers to determine the critical angle at which only a single maximum is detected instead of two.

Then, to evaluate fiber detection success, we use noisy synthetic data generated with SNR 35 and with 1, 2, or 3 fibers chosen randomly with equal volume fraction and random angle between fibers set above 45° . We generate 1000 such HARDI profile separately and count the number of times we correctly detect the number of ODF maxima. For the simulations, we also vary estimation order $\ell = 4, 6$ and 8 , use b -values of 1000, 3000 and 5000 s/mm^2 and use sampling densities of $N = 81$ and 321 on the hemisphere, corresponding to a 3rd and 7th order tessellation of the icosahedron respectively. We also record the angular error made in degrees on each direction.

Finally, to evaluate angular error made on the detected maxima, we fix the fODF, fSD and CSD estimation order at $\ell = 6$ and use a simulation with $b = 3000 \text{ s}/\text{mm}^2$, $N = 60$ (same gradient directions as the MPI acquisition of Appendix A.3.3), separation angle of 60° , SNR = 30 and volume fraction $p_1 = p_2 = 0.5$. We use a fine mesh, a 16th order tessellation of the icosahedron with 1281 directions on the hemisphere to perform fiber ODF maxima detection.

9.3.2 Real Data Experiment

We illustrate the q-ball dODF, our sharp fODF, the fSD, the FORECAST and the CSD reconstructions on the real datasets of Appendix A. We use the same ROIs as used in the previous chapter and show reconstructions on the biological phantom, the BIC dataset, the MPI dataset and the CMRR dataset. The acquisition parameters are all described in Appendix A. For all the figures, we estimate the c_j SH coefficients representing the signal S using $\ell = 4$ and Laplace-Beltrami regularization with $\lambda = 0.006$. The average prolate tensor profile was estimated for each dataset from 300 voxels in the white matter with the highest FA value, as suggested in [Tournier et al. (2004)]. For the MPI dataset, we estimated the profile $[\lambda_2, \lambda_2, \lambda_1] = [3.55, 3.55, 13.9] \times 10^{-4} \text{ mm}^2/\text{s}$, for the BIC dataset: $[2.00, 2.00, 7.64] \times 10^{-4} \text{ mm}^2/\text{s}$ and for the CMRR dataset: $[3.50, 3.50, 13.5] \times 10^{-4} \text{ mm}^2/\text{s}$. The BIC profile is markedly different with a much lower mean diffusivity. Nonetheless, it is interesting to note that all three profiles have a ratio of smallest to largest eigenvalue of $\lambda_2/\lambda_1 = 0.26$.

9.4 RESULTS

We now qualitatively show the effect of varying the HARDI signal parameters on the q-ball dODF, the sharp fODF, the fSD, the FORECAST and the CSD re-




Figure 9.6: Effect of varying estimation order ℓ and b -value on the q-ball dODF, the fSD, the FORECAST, the CSD and our fODF. W The signal is generated with $N = 60$ data points, $b = 3000 \text{ s/mm}^2$ and SNR 30. The opaque surface corresponds to the mean FOD over 100 noise trials, whereas the transparent surface corresponds to the mean + 2 standard deviations. Blue and red lines correspond to ground truth fiber directions and detected maxima respectively.

constructions. In these synthetic simulations, we have used simulations with $N = 60$, $b = 3000 \text{ s/mm}^2$, separation angle of 60° , equal volume fraction $p_1 = p_2 = 0.5$ and then have varied independently the b -value, SNR, separation angle and volume fraction.

9.4.1 Effect of Varying Signal Generation Parameters

Figures 9.6-9.9 all show a striking angular resolution gain of all SD methods reconstructing a fiber ODF over the q-ball diffusion ODF. There is a clear improvement in angular resolution of QBI with our new deconvolution operation. Also, we see that the fODF profile is overall more sharp than the fSD profile while as sharp as the




Figure 9.7: Effect of varying SNR on the q-ball dODF, our fODF, the CSD, the FORECAST and the fSD. Estimation order is fixed to $\ell = 6$ and the signal is generated with $N = 60$ data points, $b = 3000$ s/mm² and SNR 30. The opaque surface corresponds to the mean FOD over 100 noise trials, whereas the transparent surface corresponds to the mean + 2 standard deviations. Blue and red lines correspond to ground truth fiber directions and detected maxima respectively.

FORECAST and CSD profiles. However, it is possible to see that the fODF peaks are not always as perfectly aligned with the true direction as are the fSD and CSD peaks.

In particular, Figure 9.6 shows the difference in the reconstruction when changing estimation order ℓ and b -value. As expected, the best angular resolution is obtained for higher ℓ and higher b -values. Low order reconstructions with $\ell = 4$ have trouble discriminating the two fiber compartments. Note also that at the low b -value of 1000 s/mm² and $\ell = 4$, the regularized solutions of fSD produces a smoother profile that cannot discriminate the two fiber compartments. Overall, as ℓ and b -value increase, angular resolution is improved for all SD methods. We fix the order at $\ell = 6$

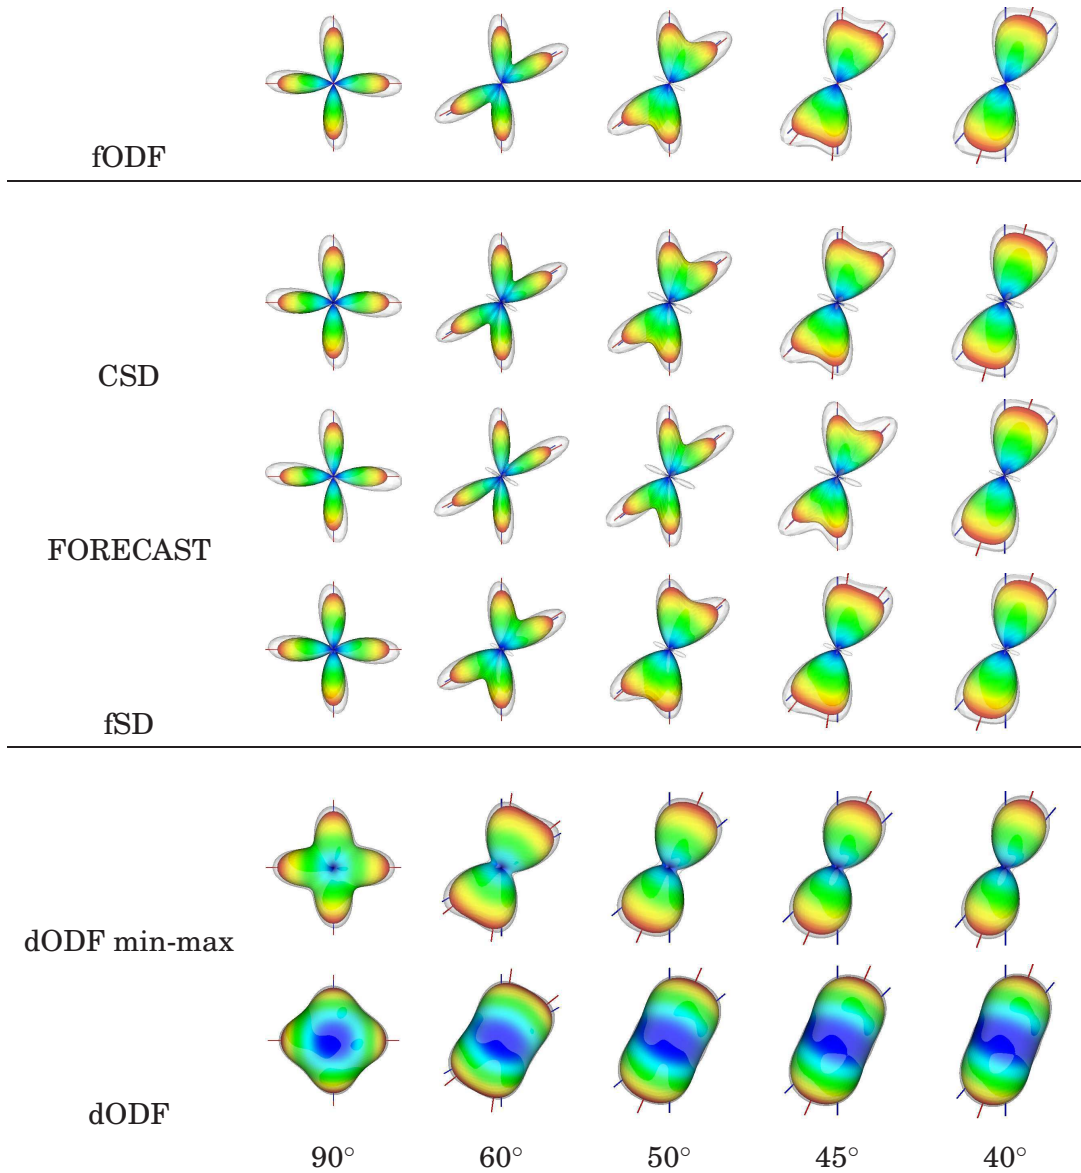


Figure 9.8: Effect of varying separation angle on the dODF, our fODF, the CSD, the FORECAST and fSD. Estimation order is fixed to $\ell = 6$ and the signal is generated with $N = 60$ data points, $b = 3000$ s/mm² and SNR 30. The opaque surface corresponds to the mean FOD over 100 noise trials, whereas the transparent surface corresponds to the mean + 2 standard deviations. Blue and red lines correspond to ground truth fiber directions and detected maxima respectively.

and $b = 3000$ s/mm² for the other simulations.

Figure 9.7 shows that all SD methods reconstructing a fiber ODF are robust to noise. Even at low SNR, the angular resolution is conserved and the spurious peaks effect is not dramatic for SNR > 10. However, for higher noise levels such as when SNR is 10, there are small spurious lobes in the center of most SD methods and a larger standard deviation of the surfaces. We fix SNR to 30 for other simulations.

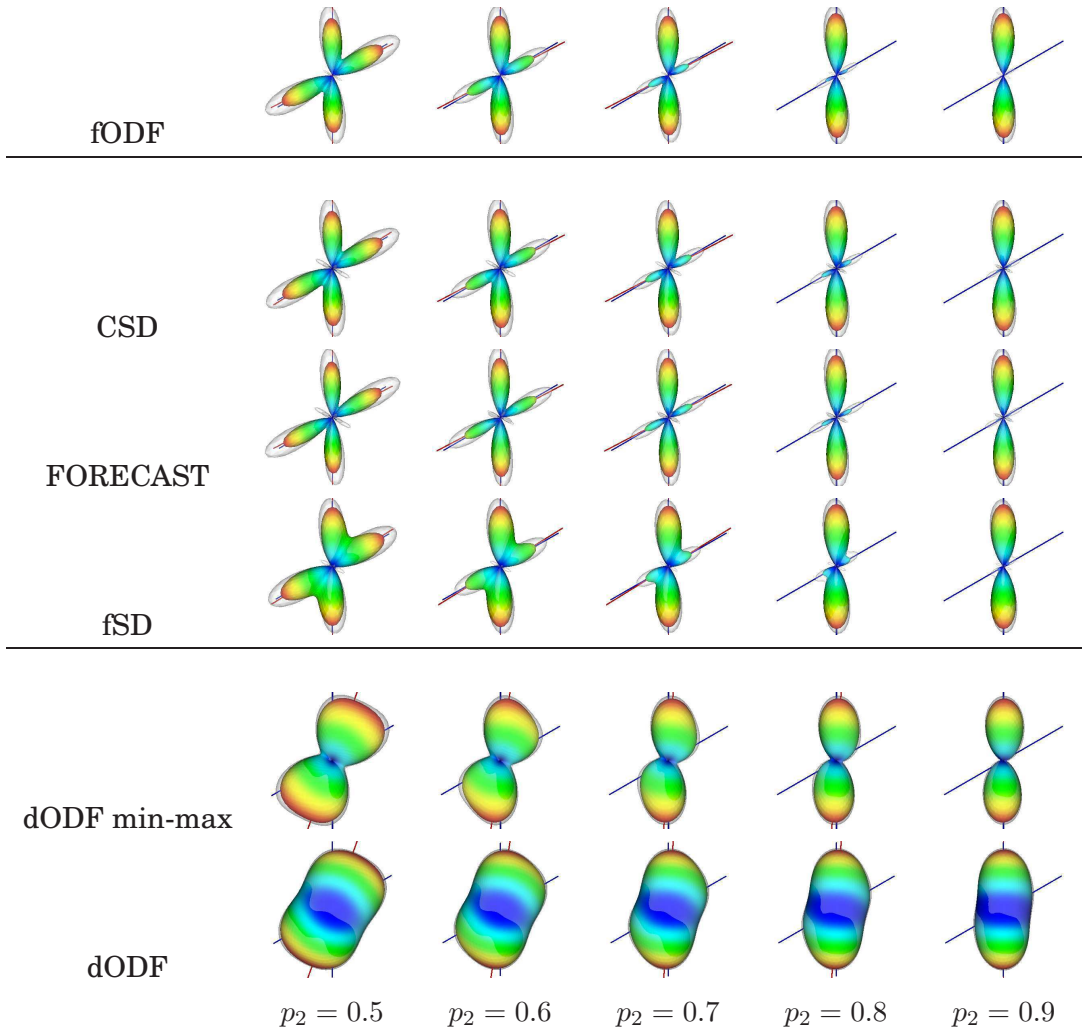


Figure 9.9: Effect of varying volume fraction on the q-ball dODF, our fODF, the CSD, the FORECAST and fSD. Estimation order is fixed to $\ell = 6$ and the signal is generated with $N = 60$ data points, $b = 3000$ s/mm² and SNR 30. The opaque surface corresponds to the mean FOD over 100 noise trials, whereas the transparent surface corresponds to the mean + 2 standard deviations. Blue and red lines correspond to ground truth fiber directions and detected maxima respectively.

Figure 9.8 qualitatively shows the gain in angular resolution for the fODF, the FORECAST and the CSD as compared to the fSD. The fODF, fSD, FORECAST and CSD are able to better discriminate the two fiber compartments at a separation angle of 45° whereas the dODF is limited to the separation angle of 60°. We fix the separation angle to 60° for other simulations.

Finally, Figure 9.9 shows that all SD methods reconstructing a fiber ODF are able to discriminate fiber compartments clearly for equal fractions up to volume fraction of 30%. For smaller fractions than 30% only a single compartment is detected for all fiber ODFs.

Overall, Figures 9.6, 9.7, 9.8 and 9.9 are in agreement with published results that can be found in [Tournier et al. (2004); Anderson (2005); Sakaie and Lowe (2007); Tournier et al. (2007); Dell’Acqua et al. (2007)].

9.4.2 Fiber Detection and Angular Resolution

Table 9.1 confirms qualitative observations made from Figures 9.6-9.9. Table 9.1(a) first shows that the angular resolution of QBI is improved with our new fiber ODF by more than 10° over all simulations. The table also shows that fODF has a slightly better angular resolution than fSD, with an average 5° difference in favor of the fODF. The improvement for order $\ell = 8$ is the most apparent with an increase of approximately 25° between fODF and dODF. At order $\ell = 6$, the increase is approximately 15° between fODF and dODF. As expected, the less apparent improvement is for lowest order $\ell = 4$, where the ODF is too smooth for the deconvolution sharpening to make an important difference. Moreover, it is also expected that the angular resolution increases considerably for higher b -values. However, note that the difference is very small when comparing $b = 5000$ s/mm² and $b = 3000$ s/mm² columns. Note also that increasing the sampling scheme does not make an important gain in angular resolution in this experiment. There is a difference of only a few degrees when going from sampling $N = 81$ to $N = 321$. This suggests that increasing the number of sampling directions is not important if one uses a low spherical harmonic order estimation. Finally, it is important to point out that when both the q-ball dODF and our new fODF successfully detect the underlying fiber populations, we record no noticeable difference in angular error made on the detected maxima, i.e. the deconvolution sharpening does not introduce errors on the detected maxima.

Then, Table 9.1(b) shows that the fiber ODF estimate increases the success rate of fiber detection in the synthetic data simulation described in Section 9.3.1. We also see that the fODF and fSD have very similar success rates in this simulation even though the fSD profile is smoother qualitatively than the fODF profile in Figures 9.6-9.9. The fODF has better success rate than fSD in some cases because it has a slightly better angular resolution. The more important increase in success rate occurs for high estimation order. In particular, for b -value 1000 s/mm², order $\ell = 8$ and sampling $N = 321$ there is an increase of 40% and for $b = 3000$ s/mm², $\ell = 6, 8$ and $N = 81$ there is an increase of more than 30% between fODF and dODF. Note that sharpening the b -value 1000 s/mm² data has the effect of improving fiber detection to above the level of b -value 3000 s/mm² data without sharpening for most order estimations and similarly, when comparing the dODF column at $b = 5000$ s/mm² and fODF column at $b = 3000$ s/mm². Moreover, increasing the spherical sampling density N increases the success rate of fiber detection. This is more apparent at high order $\ell = 8$ and for b -value > 1000 s/mm² for the dODF. Finally, note also that, for a given b -value, the fODF column at a low sampling $N = 81$ has better fiber detection success than for the dODF column at high sampling $N = 321$.

Angular resolution limitations (a)										
		$b = 5000 \text{ s/mm}^2$			$b = 3000 \text{ s/mm}^2$			$b = 1000 \text{ s/mm}^2$		
ℓ		fODF	fSD	dODF	fODF	fSD	dODF	fODF	fSD	dODF
$N = 81$	8	30°	33°	55°	31°	38°	58°	52°	54°	74°
	6	39°	42°	55°	42°	46°	59°	52°	58°	74°
	4	51°	52°	59°	52°	54°	63°	57°	63°	75°
$N = 321$	8	29°	32°	44°	30°	36°	50°	45°	50°	69°
	6	37°	40°	47°	38°	43°	52°	45°	55°	69°
	4	50°	50°	56°	52°	55°	60°	56°	62°	72°
Fiber detection success (b)										
		$b = 5000 \text{ s/mm}^2$			$b = 3000 \text{ s/mm}^2$			$b = 1000 \text{ s/mm}^2$		
ℓ		fODF	fSD	dODF	fODF	fSD	dODF	fODF	fSD	dODF
$N = 81$	8	100%	100%	78%	94%	94%	61%	86%	85%	56%
	6	99%	99%	76%	91%	90%	60%	69%	67%	54%
	4	70%	70%	62%	63%	63%	55%	62%	60%	52%
$N = 321$	8	100%	100%	96%	100%	100%	87%	95%	95%	55%
	6	100%	100%	88%	100%	100%	84%	66%	70%	53%
	4	83%	83%	62%	78%	76%	62%	58%	56%	52%

Table 9.1: The fiber ODF estimation improves fiber detection and angular resolution of QBI. We change estimation order ℓ , b -value and spherical sampling density N . Simulations were done on 1000 HARDI profiles. In (a) HARDI profiles are generated with SNR 35 and with a random number of crossing fibers between 1, 2 and 3 and with random angle between fibers above 45° . In (b), noise-free HARDI simulations are generated with 2 fibers crossing. We report the angle between fibers under which only a single ODF maxima is detected.

Table 9.1 shows results for the dODF, fODF and fSD only. Although not the same, fODF and fSD have a similar behavior. Note that for clarity and to make the table more readable, we have not included results for the FORECAST and CSD. In practice, the quantitative results for these two also agree with the qualitative observations made in the previous synthetic simulations. In fact, the angular resolution and fiber detection success of the FORECAST, CSD and our sharp fODF are nearly the same. Hence, we have not allocated a distinct column for both the FORECAST and the CSD methods.




Figure 9.10: Effect of varying SNR on the detected maxima of our fODF, the fSD and the CSD. We plot the mean and standard deviation of the estimated angle error for the fiber simulation with separation angle of 60° , $b = 3000 \text{ s/mm}^2$ and $N = 81$. The reconstructions used $\ell = 6$.

9.4.3 Robustness to Noise

Figure 9.10 shows the effect of noise on the angular error of detected maxima from the sharp fODF, the fSD and the CSD. The three curves have a similar profile but we see a small improvement of approximately 1 to 2° in angular error for the CSD method over the sharp fODF and fSD. Overall, the mean and standard deviation of the angular error is more or less $2.8^\circ \pm 1.5^\circ$ for the sharp fODF, $2.5^\circ \pm 0.9^\circ$ for the fSD and $2.1^\circ \pm 0.8^\circ$ for the CSD on average over all SNR noise levels. Note that the ODF maxima were detected using a fine mesh with 1281 directions on the hemisphere, giving roughly 4° between each mesh point. Hence, the angular error for all methods is less than the angle separation between two points of the spherical grid.

Conclusion of Synthetic Data Experiments

There are several important messages from this section. Overall, SD reconstruction methods have a striking angular resolution gain over the q-ball diffusion ODF of more than 15° . Moreover, as expected, the best angular resolution of SD methods is obtained for higher estimation order and b -value. However, surprisingly, increasing

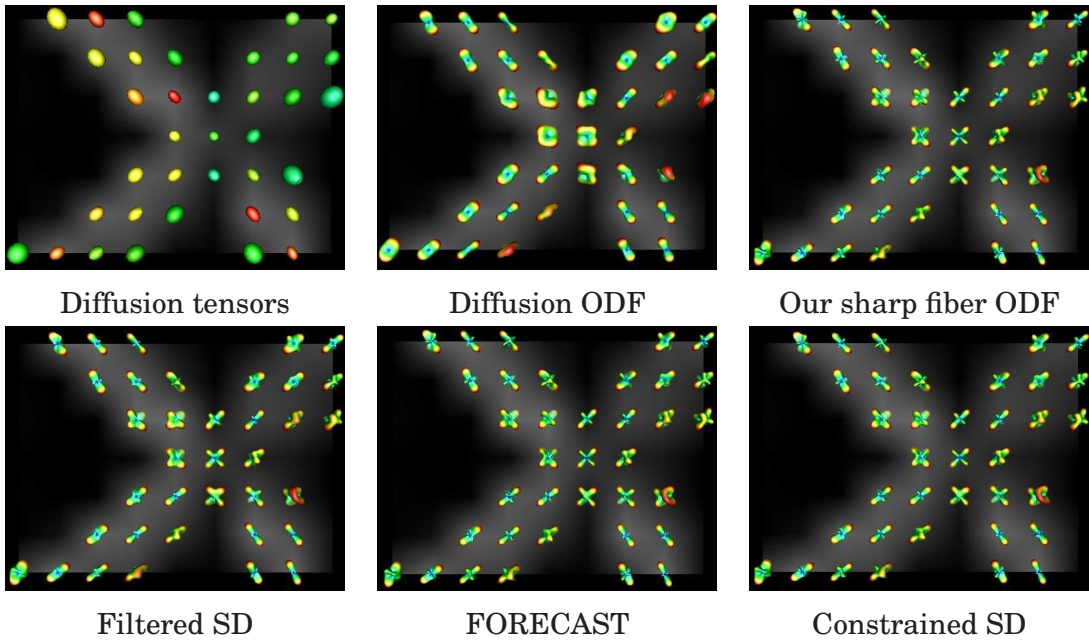


Figure 9.11: Fiber ODF reconstructions from SD methods presented in this chapter. We show the diffusion tensors, the q-ball dODF, the our fODF, the fSD, the FORECAST and the CSD overlaid on the GFA map.

sampling density from $N = 81$ to $N = 321$ only improves the angular resolution by 2 to 5 degrees.

Focusing on the SD deconvolution sharpening, the SD operation applied on the q-ball diffusion ODF improves QBI and our synthetic data simulations show that crossing fibers are more easily detected while noise effect is kept under control. The deconvolution transformation has the desired effect of enhancing the underlying fiber population, which makes it easier to detect crossing fiber configurations with smaller separation angle. Compared to fSD, FORECAST and CSD, our fODF behaves very similarly while working on the q-ball dODF and not directly on the signal. It is thus another option. Moreover, our new fiber ODF has a slightly better angular resolution and fiber detection than the original fSD but behaves essentially in the same way as FORECAST and in a similar way to CSD. Finally, regarding angular error made on the detected maxima, SD methods make less than 4° , which is less than the separation angle between two spherical grid points. Note however that CSD profiles in [Tournier et al. (2007)] seem to be even sharper with better angular resolution than the ones produce by our implementation. Numerical schemes used in the different implementations might have an impact on the end results and thus, it would be worth investigating what gives the optimal implementation of these spherical deconvolution.

9.4.4 Real Data Spherical Deconvolution Reconstructions

Biological Phantom Dataset

Figure 9.11 shows that the fiber ODFs from all SD methods enhance the multiple fiber compartments in the crossing region of the biological phantom. Fiber ODFs from all SD methods have multiple peaks that agree with the known underlying fiber populations. Again, we see that the diffusion tensors are more flat and sphere-like in the crossing region. As expected, the fiber ODF profiles are also much sharper than the diffusion ODF profiles, at the cost of small spurious peaks appearing in the fiber ODFs, especially in voxels on the borders of structure and background. Most spherical deconvolution fiber ODF reconstructions qualitatively look similar in this dataset.

Max Planck Institute (MPI) Dataset

Figure 9.12a qualitatively shows the effect of the deconvolution sharpening transformation on a single voxel of the MPI dataset. This voxel was selected at the interface between the fibers to the lateral motor stripe and the SLF (Talairach -34 -4 29). We see that the q-ball diffusion ODF finds only one maxima but there seems to be another single fiber compartment with a smaller volume fraction. The sharp fiber ODF is able to discriminate the second fiber compartment and the recorded separation angle is 62° . As the estimation order increases, the second fiber compartment is more evident at the cost of spurious peaks appearing for $\ell = 8$. We again see that fiber detection and angular resolution are improved with the deconvolution sharpening transformation.

Figure 9.12b shows the multi-directional information coming from the diffusion ODF and the sharp fiber ODF on a region of interest in a coronal slice (Talairach -4) of the human brain dataset. In this ROI, the CC forms the roof of the lateral ventricles and fans out in a massive collateral radiation, the corticospinal tract (CST) lies lateral to the ventricle and is directed vertically and the SLF crosses the base of the precentral gyrus in anterior-posterior direction. The lateral projections of the CC cross the CST and the SLF. Fibers of the SLF partly intersect with the fibers of the CST and the CC. Some voxels of the q-ball diffusion ODF and fiber ODF in area (a,a') contain these three fiber bundles crossing. It is thus surprising that recent work [Behrens et al. (2007)] report no voxels with three crossings. In fact, area (a) in this ROI, contains a large strip of voxels with low FA < 0.15 running at the medial border of the SLF, where crossings with three fiber populations are detected in the fiber ODF. Overall, our fiber ODF recovers more voxels with 2-fiber crossings than the diffusion ODF.

Figure 9.13 shows the other spherical deconvolution methods in the same ROI as Figure 9.12b. Again, we note that all SD methods improve angular resolution and fiber detection. As in the synthetic simulations, we note that our fODF profiles are as




Figure 9.12: Spherical deconvolution sharpening improves fiber detection of QBI by increasing angular resolution on coronal slice ROI of the MPI dataset. In (1), the second maxima is missed in the q-ball dODF and min-max normalized dODF. With SD, the second fiber direction is identified even at low orders of $\ell = 4$ and $\ell = 6$. The real data voxel in (a) is at the interface between motor stripe and superior longitudinal fasciculus. It was manually selected by an expert. In (b), there are more crossings detected using the fiber ODF (a,b) than dODF (a',b'). The ROI shows crossings between the cortical spinal tract / corona radiata (cst/cr) (going up in the plane), superior longitudinal fibers (slf) (coming out of the plane) and the lateral projections of the corpus callosum (cc) (in the plane). (a,a') are tilted to see the fiber directions more clearly.




Figure 9.13: We show the q-ball dODF, our fODF, the fSD, the FORECAST and CSD crossings in the same ROIs as in Figure 9.12b.

sharp as the CSD profiles and almost identical to the FORECAST profiles. We also see that our fODF profiles are sharper than the fSD profiles.

Brain Imaging Center (BIC) Dataset

As for the biological phantom and the MPI dataset, if we compare the fiber ODFs reconstructed from all SD methods on the BIC dataset, we see that the overall shapes of the fiber ODFs are nearly the same. The crossing fibers are between the CST and CC (in the plane) and the CST and SLF (coming out of the plane). Again, we see that the fiber ODF reconstructions from the different SD methods improve angular resolution while remaining robust to noise levels.

Center for Magnetic Resonance Research (CMRR) Dataset

We reconstruct the fiber ODF from all SD methods for the same axial, coronal and sagittal ROIs of the CMRR dataset in Figures 9.15 and 9.16, as in the previous chapter. The spherical deconvolution and sharpening operations are even more important on this dataset because the voxels are quite large ($3 \times 3 \times 3$) and hence, there is higher averaging of the diffusion signal due to each fiber compartment. The spherical deconvolution effect is thus quite striking in these ROIs. The QBI reconstructions are improved considerably by our sharp fiber ODF and by the fSD, FORECAST and CSD methods.

Percentage of Crossing Voxels in the Human Brain Datasets

As a final experiment on the real datasets, we count the number of voxels with $FA > 0.1$ that have two or three maxima in their respective diffusion and fiber ODF. As in the previous chapter, we use an order $\ell = 6$ reconstruction for the diffusion ODF and use an order $\ell = 4$ reconstruction for the fiber ODF from all SD methods, as in this chapter. If there are more than 3 maxima, the voxel is ignored in the statistics because it is considered as part of background voxel. In practice, for low order




Figure 9.14: Crossings in the BIC dataset between the CST and CC (in the plane) and the CST and SLF (coming out of the plane). We show the diffusion tensors, the q-ball dODF, our fODF, the fSD, the FORECAST and the CSD on the GFA map.

estimation, voxels with more than 3 maxima are rare.

Results are shown in Table 9.2. The lowest percentage is systematically recorded for the q-ball diffusion ODF, with less than 10% of voxels for datasets with $b = 1000$ s/mm^2 and nearly one fourth of the voxels in the BIC dataset with $b = 3000$ s/mm^2 .




Figure 9.15: Axial slice showing intersection between the genu of the corpus callosum, the capsule fibers and the superior frontal gyrus fibers. We show the diffusion tensors, the q-ball dODF, our fODF, the fSD, the FORECAST and the CSD overlaid on the GFA map. This is the same region of interest as Figure 7.11.

human brain datasets				diffusion and fiber ODF reconstructions				
name	N	b (s/mm ²)	voxel	dODF	fODF	fSD	FORECAST	CSD
MPI	60	1000	(1.7mm) ³	7.3%	41.4%	29.0%	39.9%	40.4%
CMRR	81	1000	(3mm) ³	9.3%	47.6%	29.9%	43.12%	53.2%
BIC	99	3000	(2mm) ³	24.0%	56.7%	55.1%	56.1%	59.3%

Table 9.2: Fraction of crossing voxels in the white matter of the CMRR, BIC and MPI datasets. We count the number of voxels with FA > 0.1 with two or three maxima in their diffusion and fiber ODF.

The percentages are considerably higher when looking at the fiber ODF reconstructions from most SD algorithms. Again, the percentages are lower for the lower b -value datasets than for the BIC dataset. For the $b = 1000$ s/mm² datasets, the fiber ODF from fSD records nearly one third of the voxels as crossings and the other fiber ODFs




Figure 9.16: (a) Coronal slice in the centrum semiovale, where the intersections between the corpus callosum commissural fibers and the corona radiata and superior longitudinal fasciculus are seen. (b) Sagittal slice showing the corona radiata diverging fibers and crossings with the superior longitudinal fasciculus.

(fODF, FORECAST and CSD) record more or less two fifth of the voxels as crossings. For the $b = 1000 \text{ s/mm}^2$ datasets, note that percentages are systematically higher for the CMRR dataset. We believe this is because the MPI dataset is smoother due to 3 signal averages and fewer diffusion encoding directions. For the higher b -value BIC dataset, more than half of the voxels are classified as crossings and approach nearly 60% for all fiber ODF reconstructions. These percentage results surely depend on the scanner and acquisition parameters but are interesting to report because we know that it is in these voxels that the classical DTI will provide unreliable results. To our knowledge, only [Behrens et al. (2007)] have performed such an analysis. Behrens et al report that one third of voxels with $FA > 0.1$ were classified as crossing voxels in their high-order model applied to a dataset acquired with a 1.5T scanner with $b = 1000 \text{ s/mm}^2$, $N = 60$ directions and 2 mm isotropic voxel size.

9.5 DISCUSSION

We have developed a new SD sharpening operation that transforms the q-ball diffusion ODF to a sharper fiber ODF from our regularized and analytical estimation of the ODF. We have shown that this fiber ODF considerably improves fiber detection and angular resolution compared to the q-ball diffusion ODF by approximately 15° . Results show that our ODF deconvolution transformation is a valid choice to obtain a stable fiber ODF. We have also shown that our new fiber ODF behaves very closely to the fSD, the FORECAST, the GB-SD and the CSD methods. We have extensively compared these SD methods and reproduced simulated results published in the literature. We have also improved and extended some of these methods by adding our Laplace-Beltrami regularization term in the signal estimation.

Our new fiber ODF is an important improvement to QBI. The classical q-ball diffusion ODF reconstruction is smooth with large contributions outside the principal directions of diffusion. Even after min-max normalization, the diffusion ODF has poor angular resolution, especially at lower b -values. Hence, the q-ball diffusion ODF is not necessarily the best function to be using if one is interested in the underlying fiber population. To our knowledge, our SD is the first attempt to transform the diffusion ODF into a fiber ODF. It is this fiber ODF that will be useful for later tractography developments in the rest of the thesis.

The SD sharpening is well-defined and the diffusion ODF kernel is estimated directly from the real data (values of λ_1 and λ_2). Therefore, we have very few parameters to set in this approach. One only needs to choose the truncated SH series order ℓ of the estimation and the regularization parameter λ . This makes our new fiber ODF estimation fast and robust when working with spherical harmonics. The powerful tool in the derivation was again the Funk-Hecke theorem, as in the derivation of the analytical ODF solution in the previous chapter. The theorem allows to solve the integral on the sphere between the single fiber diffusion ODF kernel and any spherical harmonic. We have also shown that the Funk-Hecke theorem can be

used to solve the FORECAST spherical deconvolution problem. In the end, our new fiber ODF is a linear transformation of the coefficients of the raw HARDI signal, as for the other linear SD methods (fSD and FORECAST).

Therefore, our diffusion ODF SD improves QBI considerably. Anybody working with QBI and/or a diffusion ODF can apply the sharpening transformation to obtain the fiber ODF. Results in this chapter show that crossing fibers are more easily detected and that the angular resolution limitation is improved with all fiber ODF reconstructions (fODF, FORECAST, fSD and CSD). All SD methods enhance the underlying fiber population, which makes it easier to detect crossing fiber configurations with smaller separation angles. Compared to fSD, FORECAST and CSD, our fiber ODF behaves very similarly while working on the diffusion ODF and not directly on the signal. Overall, our fiber ODF has a slightly better angular resolution and fiber detection than the original fSD but behaves essentially the same as CSD.

We have shown that, on a noise-free example, our new fiber ODF and FORECAST are the same up to small perturbations in the high order harmonics if the SH coefficients of the signal are estimated in the same way. Intuitively, this makes sense as both methods are based on the same underlying diffusion model, the axially symmetric tensor, and both use the exact same spherical harmonic derivation to obtain the respective SD solution. The principal difference is that FORECAST acts directly on the HARDI signal and uses the single fiber signal response as the deconvolution kernel whereas our new fiber ODF acts on the q-ball diffusion ODF and uses the single fiber diffusion ODF response as the deconvolution kernel. Hence, in the FORECAST solution, a single linear transformation of the signal is needed to obtain the fiber ODF whereas in our solution, two linear transformations are needed. One to transform the signal into the q-ball diffusion ODF and another to transform this q-ball diffusion ODF into a fiber ODF. Nonetheless, we have shown that FORECAST and our fiber ODF are qualitatively the same on synthetic simulations, on the rat biological phantom and on the real datasets. We have also shown that the two methods are quantitatively the same in a noise-free simulation even though small numerical differences are observed for higher order reconstructions. Further analysis needs to be done to understand the behavior of FORECAST and our fiber ODF. In our simulations, our fiber ODF seems to be less sensible to negative values on the sphere and spurious peaks in the reconstructions. This could be due to the two linear transformations included in our fiber ODF solution. The q-ball step might be adding an extra low-pass filtering in the reconstruction.

In our fiber ODF reconstruction, we have two mechanisms that limit the spurious peaks and negative-values-effect problems observed in the SD methods. First, we choose to focus on low order estimations of the ODF, i.e. we use a 4th or 6th order estimation of the fiber ODF. We find that the trade-off between noise in the higher order frequencies and gain in angular resolution is not necessary for our later tractography application. Moreover, the best results in our real data experiments are obtained

for estimation order $\ell = 4$. It also has the advantage of a very fast fiber ODF estimation (less than 20 seconds on the real datasets) and good compression properties because we only need to store 15 coefficients at every voxel. Note that [Tournier et al. (2004, 2007)] mostly consider high estimation orders of $\ell \geq 8$. Thus they need to filter high order frequencies and to do so with an iterative non-negativity constraint on the spherical deconvolution [Tournier et al. (2007)]. However, the SD operation then becomes non-linear and reduces computational efficiency of the solution. A second mechanism that limits the effect of spurious peaks in our solution is the Laplace-Beltrami regularization included in the estimation of the spherical harmonic coefficients describing the HARDI signal. This regularization is similar in spirit as the gradient constraint regularization recently proposed in [Sakaie and Lowe (2007)]. However, note that Sakaie and Lowe (2007) estimates the optimal regularization parameter at every voxel whereas we have fixed our regularization parameter to keep a fast solution. We could have also used L-curves to determine the optimal regularization λ , as in Chapter 6. Overall, we have shown that our new fiber ODF is robust to spurious peaks and negative values on the sphere by using a low approximation order and a regularization parameter. We prefer to keep our reconstruction linear and fast. Nonetheless, it is part of current work to see how we can theoretically deal with negative values and spurious peaks appearing at higher harmonic orders while working directly on the harmonic coefficients. This would avoid having to project to many data points on the sphere and also avoid the need for a non-linear algorithm [Tournier et al. (2007); Dell’Acqua et al. (2007); Jian and Vemuri (2007b)].

Finally, this chapter also showed that it is possible to resolve and improve fiber crossing detection on our real datasets, even on datasets with low spherical sampling schemes and a low b -value (60 and 81 DW images and $b = 1000$ s/mm²). We have shown that the fiber ODF reconstruction dramatically increases the number of voxels with crossing configurations computed from q-ball diffusion ODF. This is even more the case for a higher b -value acquisition. It is thus important to improve quality of high b -value acquisitions. We believe it is also important for the HARDI community to report the proportion of crossing voxels in their datasets. These voxels are locations where a classical DTI is limited and even standard single shell HARDI acquisition such as QBI reconstruction is also limited for fiber populations crossing at angles less than 90°. Proportions of crossing voxels in each dataset used should always be given, just as the SNR of the unweighted image S_0 is often given as well.

In theory, q-space imaging methods need the narrow pulse approximation and in particular, QBI approximation of the diffusion ODF is best for high b -values. Therefore, the fact that the sharp fiber ODF estimation is fast and simple to implement and that it can resolve more fiber crossings in datasets with low b -values and low spherical sampling schemes is promising for potential clinical applications of QBI. We will now see how our sharp fiber ODF is a good function to be used for tractography.

9.5.1 Contributions of this chapter:

- New deconvolution sharpening operation transforming the q-ball diffusion ODF into a sharper fiber ODF.
- Derivation, implementation and comparison of the state-of-the-art spherical deconvolution methods based on the spherical harmonic basis, namely fSD, FORECAST, GB-SD, LB-SD, CSD and super-CSD [Tournier et al. (2004); Anderson (2005); Sakaie and Lowe (2007); Tournier et al. (2007)].
- Extension and improvement of the FORECAST reconstruction using a Laplace-Beltrami regularized estimation of the SH coefficients describing the signal.
 1. Solution to the FORECAST reconstruction proved to be equivalent using a derivation using the Funk-Hecke theorem.
 2. Parallel between the FORECAST solution and our new fiber ODF solution. We have shown the numerical equivalence of the FORECAST and our new fiber ODF reconstructions in noise-free synthetic simulations.
- Extensive validation study of our new fiber ODF on numerical simulations, on the biological phantom and on the real human brain datasets.

Parts of the contributions from this chapter appear in [Descoteaux et al. (2007c); Descoteaux and Deriche (2007c); Descoteaux et al. (2007b)].

ACKNOWLEDGMENTS

Many thanks to Alfred Anwander for productive discussions and for his participation on the development of the diffusion ODF spherical deconvolution operations. Thanks to Jennifer Campbell for ideas and discussions on this problem as well. Part of this work was supported by PAI Procope and the INRIA-FQRNT fund.

9.6 APPENDIX A: DIFFUSION ODF KERNEL FOR SHARPENING

As in [Tournier et al. (2004); Anderson (2005)], we assume the simplest model of diffusion in axon fiber bundles, i.e. diffusion in each voxel follows an axially symmetric diffusion tensor \mathbf{D} . As before, the corresponding diffusion signal that would be measured in direction \mathbf{u} is $S(\mathbf{u}) = e^{-b\mathbf{u}^T\mathbf{D}\mathbf{u}}$. Without loss of generality, we assume that the fiber is aligned with the z -axis and it has a profile with eigenvalues $[\lambda_2, \lambda_2, \lambda_1]$ ($\lambda_1 \gg \lambda_2$),

$$\mathbf{D} = \begin{pmatrix} \lambda_2 & & \\ & \lambda_2 & \\ & & \lambda_1 \end{pmatrix},$$

This profile is estimated directly from the real data, taking 300 voxels with highest FA values as we are confident that there is a single fiber population at those locations.

Now, recall that the general analytical expression for the diffusion ODF Ψ , for any diffusion tensor profile \mathbf{D} , is given by Eq. 7.24

$$\Psi(\mathbf{u}) = \frac{1}{Z} \sqrt{\frac{1}{\mathbf{u}^T\mathbf{D}^{-1}\mathbf{u}}},$$

where Z is the normalization constant defined in Eq. 7.25. Here, in the case of an axially symmetric tensor, we have

$$\mathbf{u}^T\mathbf{D}^{-1}\mathbf{u} = \frac{1}{\lambda_2} \sin^2 \theta + \frac{1}{\lambda_1} \cos^2 \theta, \quad (9.24)$$

where $\mathbf{u} = [x, y, z]^T$ with $x = \cos \phi \sin \theta$, $y = \sin \phi \sin \theta$, $z = \cos \theta$.

$$\Psi(\mathbf{u}) = \frac{1}{Z} \left(\frac{1}{\lambda_2} \sin^2 \theta + \frac{1}{\lambda_1} \cos^2 \theta \right)^{-\frac{1}{2}} = \frac{1}{Z} \left[\frac{1}{\lambda_2} (1 - \cos^2 \theta (1 - \lambda_2/\lambda_1)) \right]^{-\frac{1}{2}}, \quad (9.25)$$

With the expression of Z in Eq. 7.25, we expand it to find a new normalization factor in the case of an axially symmetric tensor. The new normalization Z' is defined as

$$Z' = 8\pi b \sqrt{\lambda_2^2 \lambda_1}. \quad (9.26)$$

Letting $t := \cos \theta$ represent the dot product between the direction of the fiber and the point of evaluation, we define the single fiber diffusion ODF kernel R' as

$$R'(t) := \frac{1}{8\pi b \sqrt{\lambda_1 \lambda_2}} \frac{1}{\sqrt{1 - (1 - \lambda_2/\lambda_1)t^2}}, \quad (9.27)$$

Therefore, the Funk-Hecke integral between the diffusion ODF kernel and spherical harmonic of Eq. 9.10 is

$$\begin{aligned} r'_j &= 2\pi \int_{-1}^1 P_{\ell(j)}(t) R'(t) dt = \frac{2\pi}{8\pi b \sqrt{\lambda_1 \lambda_2}} \int_{-1}^1 P_{\ell(j)}(t) (1 - \alpha t^2)^{-1/2} dt \\ &= \frac{1}{4b \sqrt{\lambda_1 \lambda_2}} A'_{\ell(j)}(\alpha), \end{aligned} \quad (9.28)$$

where $\alpha = (1 - \lambda_2/\lambda_1)$ has the property that $0 < \alpha < 1$ and $A'_{\ell(j)}(\alpha)$ is

$$A'_{\ell(j)}(\alpha) = \int_{-1}^1 \frac{1}{\sqrt{1 - \alpha t^2}} P_{\ell(j)}(t) dt. \quad (9.29)$$

To solve this integral, the Legendre polynomial $P_{\ell(j)}(t)$ can be expressed as a power series and $A'_{\ell(j)}(\alpha)$ can be evaluated by integrating each term of the power series. We list $A'_{\ell(j)}(\alpha)$ up to order 8

$$\begin{aligned} A'_0(\alpha) &= \frac{1}{\alpha^{1/2}} \cdot 2 \arcsin(\sqrt{\alpha}) \\ A'_2(\alpha) &= \frac{-1}{2\alpha^{3/2}} \cdot [(-3 + 2\alpha) \arcsin(\sqrt{\alpha}) + 3\sqrt{1 - \alpha}\sqrt{\alpha}] \\ A'_4(\alpha) &= \frac{1}{32\alpha^{5/2}} \cdot [(105 - 120\alpha + 24\alpha^2) \arcsin(\sqrt{\alpha}) + (-105 + 50\alpha)\sqrt{1 - \alpha}\sqrt{\alpha}] \\ A'_6(\alpha) &= \frac{-1}{128\alpha^{7/2}} \cdot [(-1155 + 1890\alpha - 840\alpha^2 + 80\alpha^3) \arcsin(\sqrt{\alpha}) + (1155 + \\ &\quad -1120\alpha + 196\alpha^2)\sqrt{1 - \alpha}\sqrt{\alpha}] \\ A'_8(\alpha) &= \frac{1}{8192\alpha^{9/2}} \cdot [(225225 - 480480\alpha + 332640\alpha^2 - 80640\alpha^3 + 4480\alpha^4) \arcsin(\sqrt{\alpha}) + \\ &\quad (-225225 + 330330\alpha - 132440\alpha^2 + 12176\alpha^3)\sqrt{1 - \alpha}\sqrt{\alpha}] \end{aligned} \quad (9.30)$$

9.7 APPENDIX B: COEFFICIENTS OF A IN THE FORECAST SOLUTION

Here, we list $A_{\ell(j)}(a)$ up to order-8. Recall that $A_{\ell(j)}(a)$ is given by

$$A_{\ell(j)}(a) = \int_{-1}^1 e^{-at^2} P_{\ell(j)}(t) dt, \quad (9.31)$$

Note that these are slightly different than the A coefficients in [(Anderson, 2005, Appendix A)] because we have used a different derivation and the normalization in front of the integral is simplified. Nonetheless, the ultimate fiber ODF reconstructions of

the f_j coefficients are the same and the $A_\ell(a)$ coefficients are given by

$$\begin{aligned}
A_0(a) &= \frac{1}{a^{1/2}} \cdot \sqrt{\pi} \operatorname{erf}(\sqrt{a}) \\
A_2(a) &= \frac{-1}{4a^{3/2}} \cdot [6\sqrt{a}e^{-a} + (-3 + 2a)\sqrt{\pi} \operatorname{erf}(\sqrt{a})] \\
A_4(a) &= \frac{1}{32a^{5/2}} \cdot [-10\sqrt{a}(21 + 2a)e^{-a} + 3(35 - 20a + 4a^2)\sqrt{\pi} \operatorname{erf}(\sqrt{a})] \\
A_6(a) &= \frac{-1}{128a^{7/2}} \cdot [42\sqrt{a}(165 + 20a + 4a^2)e^{-a} + 5(-693 + 378a - 84a^2 + \\
&\quad 8a^3)\sqrt{\pi} \operatorname{erf}(\sqrt{a})] \\
A_8(a) &= \frac{1}{2048a^{9/2}} \cdot [-6\sqrt{a}(225225 + 30030a + 7700a^2 + 248a^3)e^{-a} + 35(19305 - \\
&\quad 10296a + 2376a^2 - 288a^3 + 16a^4)\sqrt{\pi} \operatorname{erf}(\sqrt{a})]
\end{aligned} \tag{9.32}$$

9.8 APPENDIX C: RELATION BETWEEN OUR FIBER ODF AND THE FORECAST

Recall the SH coefficients of our fiber ODF and the FORECAST from Eqs. 9.13 and 9.12

$$\text{fODF}_j = \frac{8\pi b \sqrt{\lambda_2 \lambda_1} P_{\ell(j)}(0)}{\underbrace{S_0 A'_{\ell(j)} (1 - \lambda_2/\lambda_1)}_{t'_j}} c_j \quad \text{and} \quad \text{FORECAST}_j = \frac{e^{b\lambda_2}}{\underbrace{2\pi S_0 A_{\ell(j)} (b(\lambda_1 - \lambda_2))}_{t_j}} c_j, \tag{9.33}$$

with $A'_{\ell(j)}$ and $A_{\ell(j)}$ defined in Eqs. 9.30 and 9.32 respectively and $P_{\ell(j)}(0)$ given by

$$P_\ell(0) = \begin{cases} 0 & \ell \text{ odd} \\ (-1)^{\ell/2} \frac{1 \cdot 3 \cdot 5 \cdots (\ell - 1)}{2 \cdot 4 \cdot 6 \cdots \ell} & \ell \text{ even.} \end{cases}$$

Now one quickly realizes that the respective coefficients of t'_j and t_j are not equal. However, we can show that they are both zero, non-zero, positive and negative for each coefficient j . Hence, they have the same global effect on the reconstructed spherical functions. In fact, we note that both reconstructed fiber ODF are the same after projecting them on the sphere and normalizing them between 0 and 1. They are the same up to small perturbations propagating in the higher order harmonics.

We illustrate this numerically on the following synthetic simulation.

- We generate a noise-free signal with two orthogonal fibers, each with FA = 0.7, $N = 60$ DW images and $b = 3000$ s/mm².
- We estimate the SH representation of the c_j with our Laplace-Beltrami regularization of Eq. 6.13 as seen many times in this thesis (it could be any technique, as long as it is the same for both the fODF and FORECAST).

Mean difference between fODF and FORECAST

$\ell = 2$	$\ell = 4$	$\ell = 6$	$\ell = 8$
2.86×10^{-8}	6.22×10^{-3}	9.98×10^{-3}	0.0502

Table 9.3: Mean difference between the fiber ODFs reconstructed from our sharp fiber ODF and from FORECAST.

- We compute fODF_j and FORECAST_j given in Eq 9.33.
- We map the SH representation of the fODF and FORECAST to the sphere with the modified SH basis of Eq. 5.24 (the B matrix).
- We normalize both the fODF and FORECAST spherical function between 0 and 1. We compute the point-wise mean difference between the two functions and we visualize them.

In Figure 9.17, we compare the two spherical functions represented as 1D signals for each of the 81 points on the hemisphere. We also visualize them on the sphere for a 16th order tessellation of the sphere. Both our sharp fiber ODF and the FORECAST are nearly identical. The greatest difference appears for order $\ell = 8$ around the small amplitudes of the signal. These are due to the small spurious peaks appearing for high order SD. Quantitatively, Table 9.3 confirms these observations. Both the fODF and the FORECAST are in close agreement. There is less than 1% difference at reconstruction orders $\ell \leq 6$ and 5% difference at the higher order of $\ell = 8$. Finally, in Figure 9.17, we show the bar graph and signature profile of the SH representations of fODF_j and FORECAST_j . We plot $\text{fODF}_j / \max(\text{fODF}_j)$ and $\text{FORECAST}_j / \max(\text{FORECAST}_j)$ for all j up to order $\ell = 8$. We see that the SH representations are quite different but are both zero, non-zero, positive and negative for the same coefficients. It is also interesting to note that for order $\ell \leq 6, j \leq 28$, the normalized FORECAST is always higher in absolute value than fODF whereas for $\ell = 8, j \in [29, 45]$, the fODF becomes always higher in absolute value. This coefficient analysis needs further investigation numerically and needs a theoretical explanation.

Overall, we can conclude that our SD operation generating the sharp fiber ODF is behaviorally equivalent to the FORECAST SD method on a noise-free example.


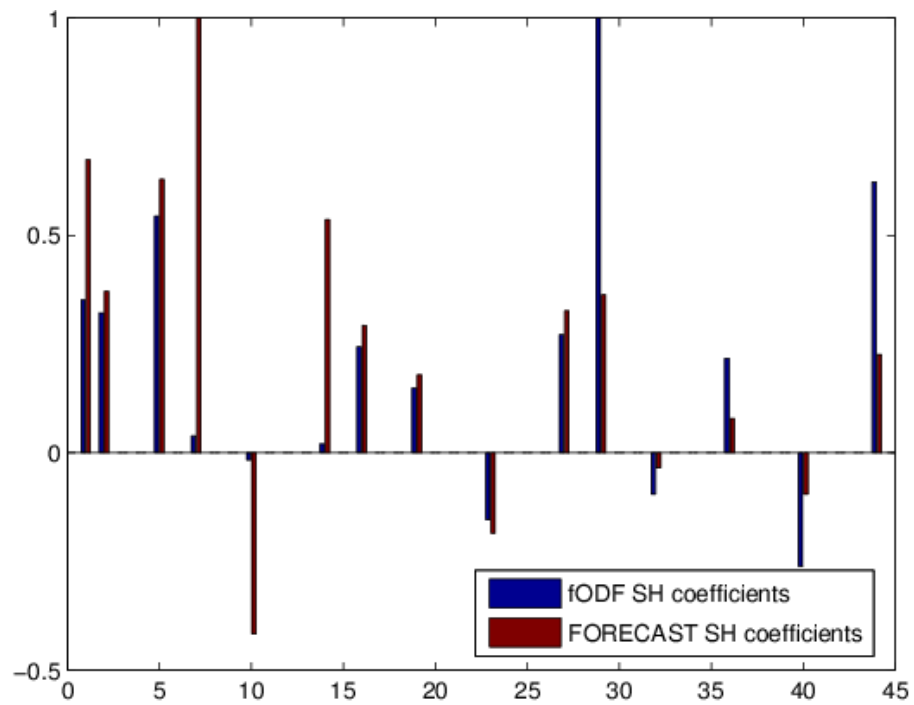



Figure 9.17: Our fiber ODF and the fiber ODF reconstructed from FORECAST are in close agreement.



Bar graph of the normalized SH coefficients, $fODF_j$ and $FORECAST_j$



Signature profile of the normalized SH coefficients, $fODF_j$ and $FORECAST_j$

Figure 9.18: Our fiber ODF versus the fiber ODF reconstructed from FORECAST represented in SH coefficients. We have normalized the SH coefficients according the respective maxima of $fODF_j$ and $FORECAST_j$. Recall that for $j \in [2, 5]$ $\ell = 2$, for $j \in [6, 15]$ $\ell = 4$, for $j \in [16, 28]$ $\ell = 6$ and for $j \in [29, 45]$ $\ell = 8$.

DETERMINISTIC AND PROBABILISTIC TRACTOGRAPHY IN HIGH ANGULAR RESOLUTION DIFFUSION MRI

“Follow the leader”

Contents

10.1 Introduction	207
10.2 Background on Tractography	208
10.3 Tractography	210
10.3.1 Deterministic Multidirectional ODF Tracking	210
10.3.2 Probabilistic fiber ODF Tracking	212
10.4 Tractography Results	213
10.4.1 Synthetic Simulation Results	214
10.4.2 Biological Phantom Results	215
10.4.3 Human Brain HARDI Data Results	216
10.5 Discussion	223
10.5.1 Contributions of this chapter:	226

OVERVIEW

Now that we have robust fiber ODF estimates, how can we integrate this local information into new tractography algorithms? This chapter covers a new deterministic tractography algorithm and a new probabilistic algorithm based on the fiber ODF. In particular, how can the deterministic DTI tracking be generalized? Can deterministic tracking take into account all the maxima of the fiber ODF? Can DTI probabilistic tracking be generalized? How does deterministic and probabilistic tracking compare? First, a new deterministic tracking algorithm and a new probabilistic tracking algorithm are developed to extend the classical DTI tractography algorithms. To do so, our new fiber ODF is used. Then, accurate results of complex fiber bundles with crossing, fanning and branching configurations are obtained from our new tracking algorithms. Finally, an extensive comparison of the new deterministic and the new probabilistic tracking algorithms is performed a human brain dataset.

Keywords: fiber tractography, deterministic tractography, probabilistic tractography, DTI, HARDI, QBI, diffusion ODF, fiber ODF

Contributions of this chapter:

- Extensive literature review of the existing state-of-the-art deterministic and probabilistic tractography methods in DTI and HARDI.
- New deterministic HARDI tractography based on the full multidirectional information of the fiber ODF.
- New probabilistic HARDI tractography based on the full fiber ODF distribution.
- Qualitative comparison and validation of our new deterministic and probabilistic HARDI tractography algorithm using the DT, the diffusion ODF and the fiber ODF. The results are shown on synthetic data, on the biological phantom and on a real human brain HARDI dataset.

Organization of this chapter:

The chapter is organized as follows. We motivate the tractography problem on HARDI data in Section 10.1. Then, we review the existing methods in DTI and HARDI deterministic and probabilistic tractography in Section 10.2. Next, Section 10.3 describes our new deterministic and probabilistic q-ball tractography algorithms. Results on synthetic HARDI data, on a biological phantom and on the MPI dataset are presented in Section 10.4. Finally, we conclude with a discussion of the results and our contributions in Section 10.5.

10.1 INTRODUCTION

Now that we have developed state-of-the-art local reconstruction methods able to deal with multiple fiber compartments, we need to integrate this local information to obtain more global information about the white matter geometry. Diffusion MRI fiber *tractography* or fiber *tracking* is the only non-invasive tool to obtain information on the neural architecture in vivo of the human brain white matter. Tractography is needed to understand functional coupling between cortical regions of the brain and is important for characterization of neuro-degenerative diseases, for surgical planning and for many other medical applications [Mori and van Zijl (2002)]. Tractography and tracking will be used interchangeably and as synonyms in this chapter. Currently, white matter fiber tractography is most commonly implemented using the principal diffusion direction of the DTI data. The DT model [Basser et al. (1994a)] characterizes the orientation dependence of the diffusion probability density function (PDF) of the water molecule. However, as seen many times in this thesis, an important limitation of the DT model is the Gaussian diffusion assumption, which implies that there can only be a single fiber population per voxel. At the resolution of DTI acquisitions, this is an important problem since it is known that many voxels have low anisotropy index due to non-Gaussian diffusion coming from multiple fibers crossing, branching, fanning or in a bottleneck. In fact, the resolution of DTI acquisitions is usually between 3 mm^3 and 15 mm^3 , while the diameter of bundles of axons considered in fiber tractography are on the order of 1mm and individual physical fibers on the order of 1-30 μm [Mori and van Zijl (2002)]. In the previous chapter, it was shown that between one third and two thirds of white matter voxels with $\text{FA} \geq 0.1$ in our brain datasets contain evidence for a multiple fiber configurations. These are locations where we know that the DT model is unreliable. Thus, tractography algorithms based on the DT can follow false tracts due to DT profiles that are prolate or can prematurely stop in regions of isotropic tensors.

To overcome these limitations of the DT, we have seen that HARDI techniques have been proposed to estimate the ODF. All the existing HARDI techniques are developed to deal with non-Gaussian diffusion process and reconstruct spherical functions with potentially multiple maxima aligned with the underlying fiber distribution. Hence, one naturally wants to generalize existing DT-based tractography algorithms with HARDI-based techniques to better deal with some fiber crossings.

In tractography, two families of algorithms exist: deterministic and probabilistic algorithms. Research groups have recently started to generalize both deterministic and probabilistic DT-based tractography algorithms to use some of the HARDI reconstruction methods mentioned in Chapter 4. Popular high order functions used in the literature are the ODF [Tuch (2002); Hagmann et al. (2004); Campbell et al. (2005); Perrin et al. (2005); Campbell et al. (2006); Haroon and Parker (2007)], the PAS function [Parker and Alexander (2005); Seunarine et al. (2006, 2007)], variants of the multi-tensor fitting models [Kreher et al. (2005); Guo et al. (2006);

Ramirez-Manzanares and Rivera (2006); Behrens et al. (2007)] and parametric spherical deconvolution [Kaden et al. (2007)]. The latter techniques show improvement in tracking results where the DT model fits the data poorly and show preliminary fiber bundles with some crossing and branching configurations handled.

Normally, an advantage of the probabilistic tractography techniques is that they are based on the full spherical function considered (DT, ODF, PAS, etc...) and not only on the principal direction or maximum(a) extracted. However, this creates a problem when the spherical function profiles are smooth and have significant isotropic parts. In that case, tractography produces diffusive tracking results that leak into unexpected regions of the white matter. This is a well-known problem in probabilistic DT tractography and has not been thoroughly studied in the literature. Typically, one simply takes a power of the diffusion tensor [Lazar et al. (2003); Koch et al. (2002); Tournier et al. (2003); Descoteaux et al. (2007c); Anwander et al. (2007b)] to increase the ratio of largest to smallest eigenvalue and thus have enhanced and more elongated tensors. To avoid diffusive tracking and leaking in ODF-based tracking methods, one needs to use the fiber ODF to obtain more complete and accurate tracts. In this chapter, the high order spherical function used is the fiber ODF developed in the previous chapter. We will compare the behavior of DTI tracking, diffusion ODF tracking and fiber ODF tracking.

The goal of this chapter is thus to integrate the full multidirectional information of the fiber ODF in both a new deterministic and a new probabilistic tractography algorithm with emphasis on a comparison on real human brain fiber bundles where classical deterministic and probabilistic DT techniques fail.

10.2 BACKGROUND ON TRACTOGRAPHY _____

The most intuitive tracking algorithms are the classical deterministic streamline (STR) tracking algorithms [Mori et al. (1999); Conturo et al. (1999); Basser et al. (2000); Poupon (1999)] and slightly more complex tensor deflection (TEND) algorithms [Lazar et al. (2003); Westin et al. (2002)] used in many applications [Horsfield and Jones (2002)]. Many other DT-based streamlines and flow-based approaches also exist. A good review and discussion of DT-based algorithms can be found in [Mori and van Zijl (2002)]. Here, we focus on HARDI-based tractography algorithms.

Recently, [Tuch (2002); Hagmann et al. (2004)] have proposed a generalized streamline tracking algorithm based on the principal direction of the diffusion ODF computed from DSI. In [Kreher et al. (2005)] and in [Bergmann et al. (2007)], a multi-tensor local model of the data is used to extend the fiber assignment by continuous tracking (FACT) [Mori et al. (1999)] algorithm and in [Chao et al. (2007b)], multiple maxima of the q-ball ODF are used to also extend FACT into a multiple ODF maxima FACT (M-FACT). Moreover, to deal with more complex fiber configurations, [Parker and Alexander (2003)] extended streamline tracking with a mixture

of Gaussian densities and similarly, [Guo et al. (2006)] recently extended the TEND model with a bi-Gaussian model. Finally, based on the classical diffusion ODF reconstructed from QBI [Tuch (2004)] and the very recent regularized version of the diffusion ODF [Campbell et al. (2006)] proposed an extension of the streamline approach with curvature constraint that allows to deal with fibers crossing. In this chapter, we propose another extension to streamline tractography based on the multiple maxima information of the fiber ODF. From this fiber ODF, we extract all available maxima and allow for splitting in multiple directions at each step. Not only can the tracking propagate through crossing fibers but it can also deal with some fibers fanning and branching.

Existing deterministic HARDI-based techniques mostly show that tracking is improved where the DT model fits the data poorly. However, these deterministic tractography algorithms inherit the classical limitations of deterministic algorithms such as choice of initialization [Jones and Pierpaoli (2005)], sensitivity to the estimated principal direction, lack of a straightforward way to compute statistics on tracts, and lack of connectivity information between regions of the brain [Tuch (2002)]. To overcome limitations of deterministic tractography, DT-based probabilistic [Koch et al. (2002); Parker and Alexander (2003); Behrens et al. (2003); Lazar and Alexander (2005); Friman et al. (2006); Ramirez-Manzanares and Rivera (2006)] and geodesic [Lenglet (2006); Jbabdi et al. (2004, 2007a)] algorithms have been used. This also motivates the development of new HARDI-based probabilistic algorithms. Probabilistic algorithms are computationally more expensive than deterministic ones but can better deal with partial volume averaging effects and noise uncertainty in underlying fiber direction. Most importantly, the output of the algorithms are usually designed to give a connectivity index measuring how probable two voxels are connected to one another.

HARDI-based probabilistic tractography have recently been published in the literature [Perrin et al. (2005); Parker and Alexander (2005); Seunarine et al. (2006); Behrens et al. (2007); Jbabdi et al. (2007b); Savadjiev et al. (2007); Chao et al. (2007a); Seunarine et al. (2007); Haroon and Parker (2007); Kaden et al. (2007)] to generalize several existing DT-based methods. First, in [Kaden et al. (2007)] parametric spherical deconvolution is used and in [Behrens et al. (2007)] a mixture of Gaussian model is used to extend the probabilistic Bayesian DT-based tracking [Behrens et al. (2003)]. Related to these techniques, [Jbabdi et al. (2007b)] uses a Bayesian framework to do global tractography instead of tracking through local orientations. In [Perrin et al. (2005)], Monte Carlo particles move inside the continuous field of q-ball diffusion ODF and are subject to a trajectory regularization scheme. In [Parker and Alexander (2005); Haroon and Parker (2007)], an extension to their DT-based approach [Parker and Alexander (2003)] is also proposed using a Monte Carlo estimation of the white matter geometry. Their implementation is based on PAS-MRI with a new noise modeling component and recently, a

Bingham distribution is used to model the peak anisotropy in the fiber distributions [Seunarine et al. (2007)]. In [Chao et al. (2007a)], large number of M-FACT QBI streamlines are reconstructed and all pathways are reversed-traced from their end points to generate of map of connection probability. Overall, these methods show successful tracking of several fiber bundles difficult to recover with DT-based techniques. Although [Seunarine et al. (2007)] does take fanning configurations into consideration, most of these methods have mainly focused on fiber bundles with crossing configurations and have not attempted to account for bundles demonstrating high curvature or points where fiber populations branch or fan. Only very recent work by [Savadjiev et al. (2007)] show preliminary results that can take into account branching fiber configurations. In this chapter, our new probabilistic algorithm attempts to account for branching and fanning fiber populations as well as fibers crossing. The novelty is to use our new q-ball fiber ODF.

As seen in Chapter 9, anybody working with QBI and diffusion ODFs has the problem of dealing with smooth ODFs that have a large diffusion part not aligned with the principal fiber directions. One needs to use a fiber ODF to obtain more complete and more accurate fiber tracts. The sharper peaks of the fiber ODF will make the probabilistic tracking have a smaller fiber orientation range to choose from at each step and will thus improve the spatial resolution of the connectivity measures.

10.3 TRACTOGRAPHY

In this section, we use our analytical solution to QBI of Chapter 7 to compute the diffusion ODF and our spherical deconvolution operation of Chapter 9 to obtain a sharper fiber ODF. We now introduce a new deterministic and a new probabilistic tracking algorithm that will be used with the DT, the diffusion ODF and the fiber ODF.

10.3.1 Deterministic Multidirectional ODF Tracking

We extend the classical streamline techniques [Mori et al. (1999); Mori and van Zijl (2002); Conturo et al. (1999); Basser et al. (2000)] based on the DT principal direction to take into account multiple fiber ODF maxima at each step. We denote $p(s)$ as the curve parameterized by its arc-length. This curve can be computed as a 3D path adapting its tangent orientation locally according to vector field \mathbf{v} . Hence, for a given starting point p_0 , we solve

$$p(t) = p_0 + \int_0^t \mathbf{v}(p(s)) ds. \quad (10.1)$$

The integration is typically performed numerically with Euler or Runge-Kutta schemes of order 2 or 4. In the Euler case, we have the discrete evolution equation

$$p_{n+1} = p_n + \mathbf{v}(p_n) \Delta s, \quad (10.2)$$

where Δs is a small enough step size to obtain subvoxel precision. A continuous linear, cubic, spline or geodesic [Lenglet et al. (2006a)] interpolation of the vector field can be done at each step for the subvoxel points. A good review is found in [Mori and van Zijl (2002)] and more recently in [Campbell et al. (2005); Hagmann et al. (2006b)].

For our deterministic algorithm, we need one or more seed points p_0 ; we need an anisotropy measure map A (FA, GFA or any other anisotropy measure); we need an anisotropy measure threshold t_{aniso} to make sure we stay in the white matter; we need a curvature threshold t_θ to make sure curves do not come back on themselves; we need a function $\text{ExtractMax}(\Psi, p)$ that returns the list l of directions l_j along each fiber ODF maxima Ψ' at point p ; we need a function $\text{size}(l)$ that returns the size of list l .

The deterministic fiber ODF tracking algorithm can be described as follows:

(0) Estimate field of fiber ODF, Ψ' , at every point p

(1) **Set** seed p_0 and **set** $\mathbf{v}(p_0) = \underset{\mathbf{u}}{\text{argmax}} \Psi'(\mathbf{u})_{p_0}$

(2) Update curve according to Eq. 10.2.

If $A(p_n) < t_{aniso}$

then STOP;

If $\frac{\mathbf{v}(p_n) \cdot \mathbf{v}(p_{n-1})}{\|\mathbf{v}(p_n)\| \|\mathbf{v}(p_{n-1})\|} > t_\theta$

then STOP;

Let $l = \text{ExtractMax}(\Psi', p_n)$. **If** $\text{size}(l) > 1$

then SPLIT curve;

for $i = 1$ to $|l|$

do (1) with $p_0 = p_n$ and $\mathbf{v}(p_0) = l_i$;

The stopping criteria of the algorithm is based on the anisotropy measure threshold and the curving angle threshold provided. In the DT tracking, we normally provide the FA measure. However, for diffusion and fiber ODF tracking, we prefer to provide the GFA or GA measure because they are more stable in the fiber crossing regions, as seen in Chapter 6. The ExtractMax function is the same as described earlier in Chapter 7, Section 7.4.7. Recall that to extract the fiber ODF maxima, we use a 16th order tessellation of the icosahedron, which gives 1281 sample directions on the sphere. Then, a finite difference method is applied on the mesh points. If a mesh point is above all its neighbors and if this point has a normalized ODF value above 0.5, we keep the mesh point direction as a maxima.

Implementation We use $t_{aniso} = 0.1$ as a mask to prevent tracks to leak outside white matter, we set curving angle threshold $t_\theta = 75^\circ$ and $\Delta s = 0.1$ unless specified. We also use a Euler integration and trilinear interpolation to obtain the diffusion ODF, the fiber ODF and the DT at subvoxel precision. The diffusion and fiber ODF are interpolated directly on their spherical harmonic representation, which is also

very efficient. For the rest of the chapter, DT-STR refers to this algorithm using the DT principal eigenvector, dODF-STR and fODF-STR refers to this algorithm using a single diffusion/fiber ODF maxima that is the closest to the incoming direction, and SPLIT-STR refers to this algorithm using all available fiber ODF maxima with splitting at each step.

10.3.2 Probabilistic fiber ODF Tracking

In this section, we propose an extension of the random walk method proposed for DTI by [Koch et al. (2002)]. We want to use the full multidirectional of the fiber ODF. That is, we want to use the whole fiber ODF and not only the maxima detected. We imagine a particle in a seed voxel moving in a random manner with a constant speed within the brain white matter, where the transition probability to a neighboring point depends on the local fiber ODF, Ψ' . This fiber ODF is discretized into N directions evenly distributed and at every time step, one of these N direction is picked at random according to the fiber ODF distribution. This yields higher transitional probabilities along the main fiber directions. Hence, the particle will move with a higher probability along a fiber direction than perpendicular to it. We start a large number of particles from the same seed point, let the particles move randomly according to the local fiber ODF and count the number of times a voxel is reached by the path of a particle. The random walk is stopped when the particle leaves the white matter mask provided by the user.

For each elementary transition of the particle, the probability for a movement from the seed point x to the target point y in one of the N directions \mathbf{u}_{xy} is computed as the product of the local fiber ODFs in direction \mathbf{u}_{xy} , i.e.

$$P(x \rightarrow y) = \Psi'(\mathbf{u}_{xy})_x \cdot \Psi'(\mathbf{u}_{xy})_y \quad (10.3)$$

where $P(x \rightarrow y)$ is the probability for a transition from point x to point y , $\Psi_{\text{sharp}}(\mathbf{u}_{xy})_x$ is the fiber ODF at point x in direction xy and $\Psi_{\text{sharp}}(\mathbf{u}_{xy})_y$ is the fiber ODF at point y in direction yx (by symmetry, direction xy and yx are the same). Note that in the original DTI formulation of [Koch et al. (2002)] and implementation of [Anwander et al. (2007b)], the transitional probability from point x to point y is given by

$$P(x \rightarrow y) = (\mathbf{D}(\mathbf{u}_{xy})_x \cdot \mathbf{D}(\mathbf{u}_{xy})_y)^\alpha, \quad (10.4)$$

where $\mathbf{D}(\mathbf{u})_x$ is the DT at point x in direction \mathbf{u} and the heuristic tensor sharpening factor¹ used is $\alpha = 7$.

The transition directions in the local model are limited to N discrete directions corresponding to the angular sampling resolution of the acquired data, and the step

¹This DTI sharpening operation can be done more formally with a spherical deconvolution on the spherical harmonic representation of the DT [Descoteaux et al. (2007c)], in a similar way has done in the previous chapter.

size of the particle step was fixed to 0.5 times the voxel size. We used trilinear interpolation of the fiber ODF for the subvoxel position and projected to N discrete directions (same as our real data gradient encodings). Moreover, voxels within the CSF and voxels containing mainly gray matter were excluded from the tracking using a mask computed from a minimum FA value of 0.1 and a maximum ADC value of $0.0015 \text{ mm}^2/\text{s}$. These values were the best to produce a good agreement with the white matter mask from the T1 anatomy. The mask was morphologically checked for holes in regions of low anisotropy due to crossing fibers. Finally, a total of 100000 particles were tested for each seed voxel. To remove artifacts of the random walk, only voxels which were reached by at least 100 particles were used for further processing. For visualization purposes, the dynamic range of the connectivity values was changed by logarithmic transformation and the entire connectivity measure is normalized between 0 and 1. For the rest of the chapter, the 3D distribution of connected voxels to the seed voxel is called a *tractogram*.

The main novelties in the probabilistic algorithm compared to previous published versions [Koch et al. (2002); Anwander et al. (2007b)] are the use of the fiber ODF, the higher angular sampling N combined with a continuous interpolation of the data and the subvoxel interpolation. The use of the fiber ODF is shown to have an important impact on the quality of the tractogram. For the rest of this thesis, the method just described is referred to as DT-PROBA, dODF-PROBA and fODF-PROBA depending on how the transitional probabilities (Eq. 10.3 and Eq. 10.4) are used with the DT, the diffusion ODF or the fiber ODF respectively.

10.4 TRACTOGRAPHY RESULTS ---

In this section, we compare and study tracking results from the deterministic tracking algorithms, DT-STR, dODF-STR, fODF-STR, and SPLIT-STR of Section 10.3.1, and from the probabilistic tracking algorithm, DT-PROBA, dODF-PROBA and fODF-PROBA of Section 10.3.2. The tracking is performed on a synthetic branching data example and on human fiber bundles with known crossings in regions where DT-based algorithms are limited. We also compare deterministic and probabilistic tracking methods on complex fiber bundles from the MPI human HARDI data. We reconstruct several commissural fiber tracts (see Chapter 3) from different seed points. Overall, we observe three results: 1) fODF-PROBA and SPLIT-STR are able to track through fiber crossings and recover crossing, fanning and branching fiber configurations, 2) SPLIT-STR and fODF-PROBA are better than dODF-based and DT-based tracking in regions of fiber crossings, and 3) fODF-PROBA has the classical advantages of probabilistic algorithms over deterministic SPLIT-STR and also is able to recover more projections from the CC to lateral areas of the cerebral cortex.





Figure 10.1: SPLIT-STR is able to recover the known fiber fanning configurations. In the brain, example of fanning configurations occur in the corticospinal tract and corona radiata where fibers fan out to different parts of the motor cortex. The deterministic tracking was started at the bottom of the branch in all cases. HARDI data was generated with b -value of 3000 s/mm^2 , $\text{FA} = 0.7$, $N = 81$ and $\text{SNR} 35$.



fODF-PROBA with (left, middle, right) initialization voxel

Figure 10.2: fODF-PROBA tracking done with 3 different initializations overlaid on DT principal eigenvector. fODF-PROBA is not particularly sensitive to initialization.

10.4.1 Synthetic Simulation Results

Figure 10.1 shows the limitations and differences of DT-STR results compared to the fODF-STR and SPLIT-STR results. Tracking was started at the bottom of the branch in all cases. Note that where DTs are prolate with principal direction not agreeing with the true fiber orientations, the fiber ODFs have multiple maxima that match with the underlying fiber population. Hence, the path followed by DT-STR is wrong and follows a false direction that takes it to the middle of the branch. Had there been




Figure 10.3: Deterministic tracking on the biological phantom. DT-STR cannot track through the crossing area with a standard FA threshold of 0.1 whereas dODF-STR, fODF-STR and SPLIT-STR can. Moreover, HARDI tracking results are able to recover parts of the high curvature fibers.

another structure behind the branching fibers, the tract could have easily leaked in the other structure and diverged. On the other hand, fODF-STR has the advantage of following the right direction. If there are two possible orientations, it goes in direction closest to its incoming direction. On the other hand, SPLIT-STR splits and follows both ODF directions when possible which recovers the full branching structure. In the brain, example of similar fanning configurations occur in the corticospinal tract and corona radiata where fibers fan out to different parts of the motor cortex. It is interesting to note the difference between dODF and fODF streamlines in the crossing area. The fODF-STR splitting occurs several voxels lower in the branching where the separation angle between the two fiber compartments is lower than 60° . This is expected from comparison tables from the previous chapter.

In Figure 10.2, we see the tractograms of the fODF-PROBA for three different initializations. The fODF-PROBA is not very sensitive to initialization and is able to recover the branching structure starting from voxels at the left, middle and right of the structure. Note that there is a fanning of tracts from one parallel fiber to the neighboring fibers within the synthetic bundle. This is expected from probabilistic tracking and this is why one typically thresholds the tractogram to obtain the most probable tracts.

10.4.2 Biological Phantom Results

Figure 10.3 shows deterministic tracking on the biological phantom. DT-STR cannot track through the crossing area with a standard FA threshold of 0.1 whereas

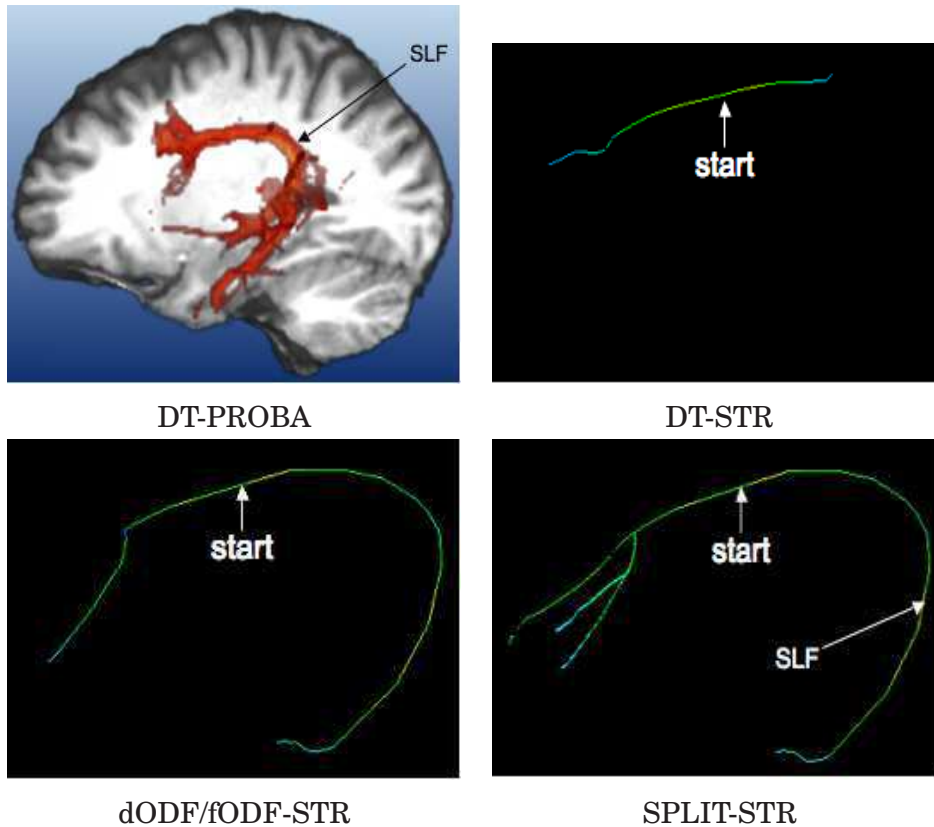


Figure 10.4: Deterministic and probabilistic tracking of the arcuate fasciculus part of the superior longitudinal fasciculus (SLF). The tracking is started from only a single seed indicated with an arrow for all methods. See Chapter 3 for comparison with the known cerebral anatomy.

dODF-STR, fODF-STR and SPLIT-STR can. It can go through the crossing if the threshold is lowered to 0.05 but in this case, the tracking steps out of the fiber bundle at the top of the structure. Deterministic tracking from the diffusion ODF and the fiber ODF produce the same qualitative result, even if allowed to split. Not only can dODF/fODF/SPLIT-STR go through the crossing area, it can also recover part of the curving section of the spinal cord at the top. Finally, if there is no curvature constraint, i.e. $t_\theta = 90^\circ$, SPLIT-STR is able to recover parts of both spinal cord bundles. In this case, most of the phantom structure is recovered.

10.4.3 Human Brain HARDI Data Results

Probabilistic Tracking Versus Deterministic Tracking

We now show tracking results on fiber bundles that are successfully reconstructed from classical DTI probabilistic tracking and that are well documented in Mori's white matter brain atlas [Mori et al. (2005)]. We implement the DT-PROBA algorithm of [Koch et al. (2002); Anwander et al. (2007b)] using Eq. 10.4. We want to see




Figure 10.5: Deterministic and probabilistic tracking of the thalamic radiations. The tracking is started from four voxel seeds in the ATR in a small area indicated with an arrow. The anterior thalamic radiations (ATR), the posterior TR (PTR), the superior TR (STR) and the corticopontine tract are recovered by SPLIT-STR. See Chapter 3 for comparison with the known cerebral anatomy.

how the DT-PROBA results compare with deterministic DTI and HARDI tracking from a low number of seed points.

The top left subfigure of Figure 10.4 and 10.5 are the tracking obtained from DT-PROBA tracking of Eq. 10.4 ([Anwander et al. (2007b)]). They agree with Mori's white matter brain atlas and known fiber bundle cerebral anatomy seen in Chapter 3. The SLF and ATR fiber bundles are well identified by DT-PROBA. Although not shown here, tractograms produced by fODF-PROBA also qualitatively agree with them as well. In Figure 10.4, the tracking is started from only a single seed in the SLF. Note that DT-STR is unable to recover the high curvature part of the fiber whereas dODF/fODF-STR can, as in the biological phantom results. Moreover, SPLIT-STR can recover the splitting part at the end of the SLF. Then, in Figure 10.5, the tracking is started from four voxel seeds in the ATR. As expected, DT-STR stops in the area of high complexity whereas dODF/fODF-STR is able to step through the crossings and recover the posterior TR (PTR) and corticopontine tract (CPT). Again,




Figure 10.6: SPLIT-STR recovers known fanning/crossing configurations to the two motor gyri from both seed points. We have colored fibers starting from the CC in red and fibers starting from the CST in green/yellow. The two motor areas of the cortex are also reached by the fODF-STR tracking whereas DT-STR and dODF-STR only captures one of the motor area.

SPLIT-STR is able to recover more by also splitting into the superior TR (STR). Overall, SPLIT-STR results qualitatively agree with the gold standard and known cerebral anatomy of both the SLF and TR obtained from DT-PROBA.

Deterministic fiber ODF Tracking

Figure 10.6 shows a branching fiber configuration in the same ROI as seen in Figure 9.12. One set of tracts (red fibers) are started from a voxel in the CC and another set of tracts are started from a voxel in the CST (green/yellow fibers). As expected, SPLIT-STR recovers the branching configuration of both fiber tracts and recovers fibers projecting in motor areas in two gyri. On the other hand, fODF-STR is able to step through the crossings whereas dODF-STR and DT-STR are limited. Both fibers starting from the CST and from the CC take the average direction of the two ODF orientations and projects only the medial motor cortex. tracking. The angle between fiber compartments varies between 60° and 80° in that area and the diffusion ODF cannot discriminate clearly the two fiber populations.

Probabilistic fiber ODF Tracking

Anterior Commissural (AC) Fibers For the reconstruction of the fibers passing through the AC fibers, a seed voxel was placed in the mid-sagittal cross section of the AC and a second tracking was started from two seeds, one on the left and right side of the mid-sagittal cross section. In Figure 10.7, the tracking results of the AC fibers shows the advantages of the fODF-PROBA tracking over dODF-PROBA and DT-STR




Figure 10.7: Deterministic and probabilistic tracking of the anterior commissure (AC) fibers. Probabilistic tracking is shown on coronal and axial slices through the AC. We also show a 3D rendering of the tractogram iso-surface. fODF-PROBA reaches the temporal pole close to the uncinata fasciculus (UNC) whereas dODF-PROBA can only reconstruct a small part. Deterministic tractography needs multiple seeding initialization in the DT and dODF case due to the small diameter of the AC on this subject.

tracking. dODF-PROBA and DT-STR are blocked close to the seed point by low FA areas. Particles of dODF-PROBA cannot propagate to the temporal poles because the paths are diffusive and leak outside the anterior commissural bundle, which is only a few voxel wide around the seed point. As a consequence, dODF-PROBA mostly recovers shorter parts of the fiber bundles. However, with a multiple seeding approach, DT-STR and dODF-STR are able to recover both paths to the temporal poles, as published in [Catani et al. (2002)] and as expected from cerebral white matter anatomy of Chapter 3. In contrary, deterministic fODF-STR and SPLIT-STR tracking can reconstruct the fibers connecting the temporal pole via the AC from a single seed point in the mid-sagittal cross section. Probabilistic tractography done with the fiber ODF suggests additional projections to more posterior parts of the temporal lobe and through the anterior subinsular white matter to the inferior fronto-occipital (IFO) fasciculus and inferior longitudinal fasciculus (ILF). Moreover, a second anterior pathway was found on the right hemisphere for this subject (Figure 10.7, top row




Figure 10.8: Deterministic and probabilistic tracking of the tapetum. Probabilistic tractogram of a single seed point in the tapetum (TAP) reconstructs the fibers crossing the CC and connecting both temporal lobes. Deterministic tracking was initialized from 4 seed points in the tapetum. SPLIT-STR reconstructed this complex structure.

middle and right).

Callosal fibers to the temporal lobe For reconstruction of fibers in the tapetum and temporal commissural fibers, we selected 4 seed voxels between the left lateral ventricle and the optic radiation close to the splenium of the CC (Talairach -22, -42, 24). Figure 10.8 shows the fibers in the splenium of the CC, which sweep inferiorly along the lateral margin of the posterior horn of the lateral ventricle which form the tapetum. These fibers are in close contact to the commissural fibers connecting the precuneus of both hemispheres. On the left hemisphere the splitting is close to the 4 voxel seed region and all methods can bind both parts. Only SPLIT-STR and fODF-PROBA can reconstruct the splitting in both hemispheres to the bundles connecting the temporal lobe and the parietal lobe. The fibers to the temporal lobe split to the transverse temporal (Heschel's) gyrus and to the ILF. Probabilistic tractography also suggest a close relation to the posterior TR.

Callosal fibers to the inferior and middle frontal gyrus For reconstruction of the commissural fibers connecting the contralateral inferior and middle frontal gyrus





Figure 10.9: Deterministic and probabilistic tracking of the projections of the corpus callosum to Broca's area. The probabilistic tractogram is shown on a coronal and sagittal slice and as 3D rendering. The tractogram shows asymmetry with stronger connections to the left inferior and middle frontal gyrus than to the homologue area. We also show a selection of the probabilistic fibers colored differently depending on their end point projections to the lateral or medial areas. From the deterministic methods only SPLIT-STR can reconstruct this complex structure.

a seed voxel was defined in the mid-sagittal section of the rostral body of the CC (Talairach 0, 18, 18). Figure 10.9 shows the tracking results. DT-STR and fODF-STR tracking can only find the commissural fibers connecting the medial parts of the frontal lobe (Figure 10.9, second row, left and middle). Fanning of the fiber bundle to the inferior and middle frontal gyrus was found with the SPLIT-STR method on the left hemisphere and to a lower extent on the right. The tractogram computed with the fODF-PROBA method reveals a strong interhemispheric connection of the lateral parts of the frontal lobe. Additional fibers are found branching to the anterior thalamic radiation (Figure 10.9, top left). Figure 10.9, top right shows sample fiber tracts included in the probability map (left, middle).

As a final comparison, we used the fiber ODF obtained from filtered spherical deconvolution (fSD) [Tournier et al. (2004)] and from constrained spherical deconvolution (CSD) [Tournier et al. (2007)] as input to our probabilistic tracking algorithm. Results are shown in Figure 10.10. As expected from the numerical simulations and real data reconstructions of the previous chapter, the tractogram produced by fODF-PROBA, fSD-PROBA and CSD-PROBA are qualitatively similar. Same regions are



Tractograms of the fODF-PROBA, fSD-PROBA and CSD-PROBA methods




Figure 10.10: Comparison of the tracking results of the projections of the corpus callosum from our sharp fiber ODF (fODF-PROBA), from the filtered spherical deconvolution (fSD-PROBA) and from the constrained spherical deconvolution(CSD-PROBA). The probabilistic tractograms are shown on a coronal, sagittal and axial slice. We also show the isosurface corresponding to the tractograms thresholded at 0.3.

connected in all tractograms with slightly different connection score. This result confirms that different local reconstruction fiber ODFs do not have a dramatic difference on the global behavior of the probabilistic tracking connectivity measures. Above all, the importance is to use a fiber ODF and not a diffusion ODF.

10.5 DISCUSSION

We have proposed an integral concept for tractography of crossing and splitting fiber bundles in the brain based on the fiber ODF. We have extensively compared both new tracking algorithms and have performed expert visual validation on complex fiber bundles from the MPI human brain dataset.

The standard q-ball diffusion ODF reconstruction is smooth with large contributions outside the principal directions of diffusion. Even after min-max normalization, the diffusion ODF has poor angular resolution, especially at lower b -values. Hence, this diffusion ODF is not the best function if one is interested in the underlying fiber population. Without the fiber ODF, the tracking is diffusive, stops prematurely and favors shorter fiber tracts. From our experience, it is crucial to use a fiber ODF. Which fiber ODF? It seems that the choice of spherical deconvolution algorithm used to obtain a fiber ODF does not greatly influence the tractogram results. After all, this is expected from our quantitative comparisons of Chapter 9.

The fiber ODF seems the right candidate function for tractography algorithms. We have used it in both our deterministic and probabilistic tracking extensions of previous DT-based methods. We clearly showed that the results are improved over standard DT-based methods. Moreover, when crossing areas have separation angle less than 90° , the diffusion ODF averages the two fiber compartments and has the same limitations as DT tracking [LeBihan et al. (2006)]. In the deterministic case, the fiber ODF has a better angular resolution and the maxima of crossing and splitting fiber bundles are detected more easily. Hence, we showed that SPLIT-STR can follow more multiple maxima and recover most of fanning structures in the regions studied. However, more investigation and better characterization of crossing, kissing, fanning and branching fiber configurations remains to be done in the human brain [Parker and Alexander (2005); Savadjiev et al. (2007)]. Although sensitive to initialization, SPLIT-STR tracking was able to recover similar bundles as the fODF-PROBA method in most cases.

SPLIT-STR is thus an efficient and easy way to obtain a good idea of fiber tracts starting from only a few seeds. The underlying assumption of SPLIT-STR is that all multiple peaked fiber ODF have an underlying branching structure, which makes it reasonable to follow all available maxima at each step. This is the reason for using a curvature threshold of the tracts of 75° instead of 90° . This threshold avoids following tracts through “pure” crossing configurations, where we know that we are then stepping into other fiber bundles. This was clearly seen in the biological phantom tracking in Figure 10.3. When the curvature threshold was set to 90° , both fiber cords were

recovered even if the tracking was initialized in one of the two single fibers. This raises questions for multidirectional deterministic tractography: Should the tracking algorithm split as much as possible to recover as much fiber structure as possible before clustering and post-processing the tracts to separate them into bundles? Or, should the tracking have a built-in scheme to differentiate the different sub-voxel crossing possibilities and decide whether or not a tract should be split in the tracking process? For instance, split in the case of a branching bundle but not split in the case of a crossing fiber because it is assumed that it steps into a different fiber bundle. Information about the local geometry, curvature and torsion of tracts can help for this problem [Savadjiev et al. (2006)] and in [Savadjiev et al. (2007)], preliminary results are obtained discriminating and labeling crossing and fanning sub-voxel fiber configurations.

To deal with the uncertainty in the diffusion ODF or fiber ODF maxima, a probabilistic approach is more robust and gives more complete results. In a probabilistic case, the sharpening operation or use of a fiber ODF has its greatest importance just as it is in probabilistic DTI tracking [Lazar et al. (2003); Koch et al. (2002); Tournier et al. (2003); Descoteaux et al. (2007c); Anwander et al. (2007b)]. For DTI probabilistic tracking, heuristic sharpening taking a power of the diffusion tensor is done to obtain a tensor representing the fiber distribution, which is crucial to obtain decent tracking results.² Similarly, in ODF-based probabilistic tracking, the tractograms are diffusive and leak into unwanted bundles if done without sharpening. However, this leaking effect is not as dominant in probabilistic ODF tracking as it is probabilistic DTI tracking. The min-max normalization of the sampled diffusion ODF values to the range between 0 and 1 removes some of the isotropic part and the shape of the diffusion ODF is sharpened. However, this min-max normalization subtracts a different value from each ODF and might introduce artifacts to the probabilistic tracking result. We believe that the spherical deconvolution methods, either based on the q-ball ODF or based on the raw HARDI signal itself, are the natural operations to obtain a fiber ODF. The spherical deconvolution methods really recover an estimation of the fiber distribution and really improve the angular discrimination of crossing fibers.

In our method, we employ an estimate of the fODF and this function is sampled directly to account for the fact that the fibers in a bundle are not all strictly parallel. Our algorithm follows all possible fiber directions from this fODF. It is, however, quite difficult to disentangle the uncertainty and the actual spreading of fiber orientations. We are therefore conservative by tracking all directions that are possible given the data and the assumed model assumptions. Other methods use calibration [Parker and Alexander (2003); Seunarine et al. (2007)], statistical techniques like Bayesian modeling [Friman and Westin (2005)], Markov Chain Monte

²Note that this DTI sharpening can also be done with spherical deconvolution applied on the spherical harmonic representation of the diffusion tensor [Descoteaux et al. (2007c)].

Carlo [Behrens et al. (2003); Kaden et al. (2007)], bootstrap [Haroon and Parker (2007)] to infer a peak uncertainty of the fiber distributions. It is now important to compare these different HARDI-based probabilistic approaches and see how the different integration of the local reconstruction information impacts the resulting tractograms.

Microscopy of white matter shows that the nerve fibers are not completely parallel within a bundle. Probabilistic tractography includes this uncertainty. The reconstructed fibers spread out to neighboring voxels within the cross section of a fiber bundle. This makes the algorithm robust to the initialization. Different starting points within an area to the same bundle lead to similar tracking results. The spreading of the tractography to the whole bundle enables the segmentation of the selected fiber bundle in its total extent. This is a clear advantage over deterministic tracking. In deterministic tracking, the sensitivity to initialization has to be solved with multiple initialization using a larger starting region and a combination of target regions.

The probabilistic tractogram represents a *connectional fingerprint* of a piece of brain tissue contained in a voxel. The information provided by the tractograms can be employed in several ways. By summing the tractograms of several neighboring seed voxel, one can obtain connectional fingerprints for greater regions (joint tractograms). It is then straightforward to derive quantitative connectivity measures between two different brain regions A and B , e.g. by just integrating the joint tractogram of region A over region B and dividing this by the norm of the entire joint tractogram. Such a measure reflects an estimate of the proportion of the fibers starting in A that reach B (see also [Kaden et al. (2007)]). Another use of the tractogram is based on the assumption that the functionality of a neuron or a population of neurons is strongly determined by its pattern of connectivity to the rest of the brain. If seed points are placed in the cortex or at the boundary between cortex and white matter, the associated tractograms would yield estimates of the global connectivity patterns of the respective pieces of cortex. By using the above hypothesis one can infer from the connectional similarity between two voxels to its functional similarity and thereby parcellate the cortex (or other grey matter structures) into internally relatively homogeneous, but mutually distinct functional-anatomical areas [Anwander et al. (2007b)]. Finally, tractograms yield a very intuitive representation of the connectivity pattern of the seed point or region when overlaid on anatomical slices. This visualization allows to look at the exact location of the fiber bundle in an anatomical context. Moreover, a 3D rendering of the iso-surface of the probabilistic tractogram gives an overview of the topology and the branching structure of the fiber bundles. In counterpart, probabilistic tractography suffers from connectivity values decreasing with the distance from the initial seed point due to the fanning of fiber bundles. It is part of current work to study this open problem.

For all these reasons, although slower than SPLIT-STR, we find that FODF-PROBA is more convenient to use for tractography. Overall, the tracking is less

sensitive to initialization and the tractogram output more informative than the 3D curves output by SPLIT-STR. fODF-PROBA is thus a very good tool to study specific fiber bundles and compare the projections for many subjects. In particular, for fiber tracts which are known to intersect with other bundles. DT-based tractography can only reconstruct a part of the fibers or might even find wrong connections in the intersection position. It is now important to continue this type of experiment and focus on complex fiber bundles involved in specific functional tasks where DTI tracking is in most cases very limited. Chapter 11 will present one such tractography application, where the transcallosal fibers are reconstructed and quantified on multiple subjects.

10.5.1 Contributions of this chapter:

- Extensive literature review of the existing state-of-the-art deterministic and probabilistic tractography methods in DTI and HARDI.
- New deterministic HARDI tractography based on the full multidirectional information of the fiber ODF.
- New probabilistic HARDI tractography based on the full fiber ODF distribution.
- Qualitative comparison and validation of our new deterministic and probabilistic HARDI tractography algorithm using the DT, the diffusion ODF and the fiber ODF. Most results are compared against classical FACT tracking [Mori et al. (1999)] and are shown on synthetic data, on the biological phantom and on a real human brain HARDI dataset.

Parts of the contributions from this chapter appear in [Deriche and Descoteaux (2007); Descoteaux et al. (2007c)].

ACKNOWLEDGMENTS _____

Thanks to Alfred Anwander for the DTI probabilistic tracking implementation and the corresponding figures he generated for that part. The development of the probabilistic HARDI tracking algorithm was closely done with his collaboration.

Part IV

Application

QUANTITATIVE ASSESSMENT OF TRANSCALLOSAL FIBERS USING PROBABILISTIC HARDI TRACTOGRAPHY

“If you find a path with no obstacles, it probably doesn’t lead anywhere”
–Frank A. Clark

Contents

11.1 Introduction	230
11.2 Quantifying Lateral Projections of the Corpus Callosum . . .	233
11.2.1 Methods	233
11.2.2 Results	234
11.3 Discussion	235

OVERVIEW

Now that we have state-of-the-art local reconstruction methods and tractography algorithms, can we track through regions of complex fiber crossings for a neuroscientific application? This chapter covers the tractography of transcallosal fibers. Current DTI based methods neglect these transcallosal fibers, which might lead to wrong interpretations of the brain functions. In particular, can transcallosal fibers be reconstructed with our new fiber ODF-based probabilistic tractography? Are the tracking results consistent and reproducible across subjects? First, we apply our new fiber ODF probabilistic tractography algorithm to reconstruct the transcallosal fibers. Then, we see how we improve the tractography results of transcallosal fibers intersecting with the corona radiata and the superior longitudinal fasciculus. Finally, the tracking results are reproduced and compared on a group of 8 subjects.

Keywords: Transcallosal fibers, corpus callosum (CC), corona radiata (CR), superior longitudinal fasciculus (SLF), fiber ODF, probabilistic tractography, DTI, HARDI

Contributions of this chapter:

- Transcallosal fibers reconstruction from our new fiber ODF probabilistic tractography.
- Consistent and reproducible tractography results on 8 subjects of the MPI HARDI database.
- Quantification and labeling of the areas of the corpus callosum that are able to project to the ventral and lateral parts of the cortex.

Organization of this chapter:

The chapter is organized as follows. We briefly introduce the importance of transcallosal fibers reconstruction in Section 11.1. Then, Section 11.2 describes our method for quantifying the transcallosal projections with results shown on the MPI HARDI database. Finally, we conclude with a discussion of the results and our contributions in Section 11.3.

11.1 INTRODUCTION

The corpus callosum (CC) is the largest fiber bundle in the human brain with 200-250 million fibers interconnecting both hemispheres. The CC is involved in the interhemispheric interaction of cortical regions and the reconstruction of fibers connecting the cerebral hemispheres is of major interest for cognitive research and clinical practice. We know that the projections of the CC play an important role in functional integration of perceptual, cognitive and learned information.




Figure 11.1: DTI tracking misses transcallosal fibers to the lateral cortex. Only connections to the medial and dorsal cortex are recovered. Images borrowed from [(Hofer and Frahm, 2006, Fig.1)].

In this short application chapter, we focus on the complex projections of callosal fibers to the cortex. These bundles have been studied extensively in the neuroanatomical and DTI literature [Hofer and Frahm (2006)]. The problem is interesting for us because the transcallosal fibers of the CC are normally missed by DTI tracking algorithms, as seen in Figure 11.1 borrowed from [(Hofer and Frahm, 2006, Fig.1)]. While




Figure 11.2: The corona radiata (CR) and superior longitudinal fasciculus (SLF) are major fiber bundles that block transcallosal fibers in standard DTI tracking. Images are borrowed and adapted from [Leemans (2006)] and from the 'ExploreDTI' open-source software [Leemans et al. (2005)]. Red fibers correspond to the CC, blue fibers to the CR and orange fibers to the uncinate fasciculus (UNC).




Figure 11.3: Diffusion ODF and fiber ODF on the MPI dataset in a complex crossing region between the transcallosal fibers, the CR and the SLF.

DT-based tractography finds fibers passing through the CC connected with the medial and dorsal cortex, lateral and ventral fibers are not found, since these fibers cross the corona radiata (CR) and the superior longitudinal fasciculus (SLF). The CR and SLF are major fiber bundles that block transcallosal fibers in standard DTI tracking. This is seen in Figure 11.2 with the CR in dark blue. The CR act as a wall for the CC fibers in red. Images are borrowed from [Leemans (2006)] and they are masked to only see the CC, CR and uncinate fasciculus.

Hence, local reconstruction and tractography of transcallosal fibers are a good test for local HARDI reconstruction and HARDI-based tractography. Is it possible to recover the multiple fiber information in these regions? Only a recent study [Price et al. (2007)] proposed to use a HARDI-based method to reconstruct fibers of the genu and splenium of the CC but they have not focus on the transcallosal projections to the lateral cortex. We have already seen the interest of QBI and spherical deconvolution in that area of the brain. Local reconstruction of Chapters 7 and 9 find multiple maxima diffusion ODFs and fiber ODFs in that region, as seen and reproduced in Figure 11.3. Moreover, our new fiber ODF tracking results of Chapter 10 are able to propagate through this complex fiber region, as seen and reproduced in Figure 11.4.

In this chapter, we investigate how the reconstruction of transcallosal fiber con-




Figure 11.4: Tracking from a single seed point in the splenium of the CC. The lateral fibers intersect with the CR and SLF.

nections can be done with our fiber ODF tracking algorithm. We want to see if the tracking results are consistent and reproducible across many subjects.

11.2 QUANTIFYING LATERAL PROJECTIONS OF THE CORPUS CALLOSUM

We want to investigate how the reconstruction of transcallosal fibers can be improved with our new fiber ODF tractography in a group of 8 subjects. The 8 datasets were all acquired with the same MPI HARDI protocol described in Appendix A.3.3. We want to show in which parts of the CC we can reconstruct fibers connecting the ventral and lateral parts of the cortex. These fibers cross the CR and parts of the SLF [Talairach and Tournoux (1988)] and cannot be detected with the simple DT model and are limited and not well recovered with the diffusion ODF.

11.2.1 Methods

1. Regions of interest (ROI) for the white matter tractography in each subject were defined manually in the sagittal cross section of the CC.
2. For each seed voxel in the ROI, the fODF-PROBA tractography of Chapter 10 was performed separately.
3. To evaluate the connectivity to lateral and ventral cortical areas, the percentage of random fiber tracts reaching lateral parts of the brain (Talairach > 30) was recorded.
4. The resulting quantitative count is color coded on the mid-sagittal plane of the




Figure 11.5: Some transcallosal fibers project to ventral and lateral parts of the cortex. These are fODF-PROBA tracking results on one subject of the MPI HARDI database. The redder the color, the more connections cross and project to the ventral and lateral cortex. The color corresponds to the percentage of fibers which cross a para-sagittal plain (Talairach > 30) as indicated by the lines on the coronal slice on the left subfigure.

CC to have a labeling indicating which parts of the CC connect to the ventral and lateral parts of the cortex.

11.2.2 Results

Figure 11.5 shows the color labeling of the ROI indicating regions with strong lateral connectivity in the MPI dataset. The color labeling is shown for all 8 subjects of the MPI HARDI database in Figure 11.6. The color corresponds to the percentage of fibers which cross a para-sagittal plain (Talairach > 30) as indicated by the lines on the coronal slice in Figure 11.5. For all subjects, a maximum number of lateral fibers was found in the genu or the rostral body of the CC connecting the inferior and middle frontal cortex and the premotor cortex (see Figure 3.6 of Chapter 3.2 for the anatomy of the CC). The callosal fibers interdigitates with the fibers of the corona radiata. The tractogram for the voxel with the most lateral connections shows fibers to the left and right middle frontal gyrus and the left inferior frontal gyrus in addition to the connections of the medial frontal cortex. The proportion of lateral fibers in this

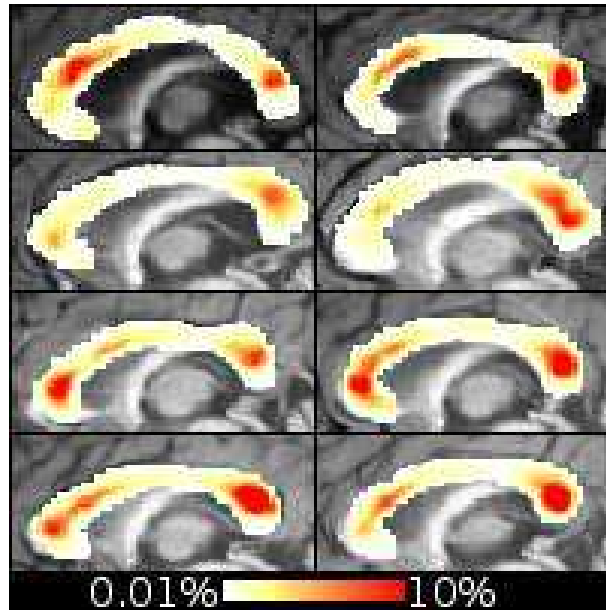


Figure 11.6: Quantification of the ratio of lateral fibers (Talairach > 30) in the different parts of the corpus callosum (cc) typically completely missed in DT-based tractography. All 8 subjects show strong lateral connectivity in the genu or the rostral body of the cc and a second peak in the splenium. fODF-PROBA finds more than 10% of lateral fibers in large parts of the cc.

voxel is $20\% \pm 6\%$ for the 8 subject study.

A second peak (bright red color) was found in the splenium of th CC connecting the temporal areas. The corresponding fibers lie deep to the optic radiation in the tapetum and connect the temporal lobe and in particular to the auditory cortex. The proportion of temporal fibers in this voxel is $30\% \pm 10\%$ for the 8 subject study.

Note that both the stronger posterior and anterior responses in the labeling of Figure 11.6 are consistent and reproducible for all subjects. These regions are areas where we are certain that DTI tracking is limited and most probably inaccurate.

11.3 DISCUSSION

We have illustrated an important application for the modeling of inter-hemispheric cognitive networks with a quantification of the projections of the corpus callosum. We have quantified regions of the corpus callosum where we are able to recover fibers projecting to the lateral cortex. Therefore, we have showed that we can find fibers from the ventral and lateral cortex in a major part of the corpus callosum. Current DT-based methods, such as those used in [Hofer and Frahm (2006)], [Leemans (2006)] and references therein, neglect these fibers completely (at least in the frontal lobe), which might lead to wrong interpretations of the brain functions.

The chapter also highlights remaining limitations of tractography algorithms, even if one uses HARDI and probabilistic tractography. From the labeling of Figure 11.6, we clearly see that we are unable to find lateral projections from medial parts of the corpus callosum. From post-mortem dissection and atlases, we know that there are interhemispheric connections passing through these voxels that project to parts of the ventral and lateral cortex. Why are we not able to recover these connections in HARDI-based tractography? When we look more closely in these regions, we see very round local diffusion ODF and DT profiles. This is because there are potentially multiple major fiber bundles crossing (more than two), which results in low anisotropy measurements and local ODF profiles that are almost isotropic or fiber ODF profiles with more than three maxima. Improvement in spatial resolution of current diffusion MRI acquisitions might help overcome this problem. However, we believe that it is also worth investigating tractography algorithms that are less local. All current tracking methods are based on a very local integration of DT or ODF information at the single voxel level. We need to look at more global tracking methods that not only integrate the local fiber ODF information but that also integrate information that is further away in neighborhoods where other fiber ODF support the local geometry of the fiber tracts we are interested in.

Nonetheless, we have showed the added value of current HARDI-based probabilistic tractography in complex fiber configuration regions and a specific application of tractography. The multiple subject study of this chapter opens to many other neuroscience perspectives and applications. It is now important to continue this type of experiment and focus on complex fiber bundles involved in specific functional tasks where DTI tracking is in most cases very limited.

Contributions of this chapter:

- Transcallosal connections reconstruction from fiber ODF tractography.
- Consistent and reproducible tractography results on 8 subjects of the MPI HARDI database.
- Quantification and labeling of the areas of the corpus callosum that are able to project to the ventral and lateral parts of the cortex.

Parts of the contributions from this chapter appear in [Descoteaux et al. (2007c); Anwander et al. (2007a)].

ACKNOWLEDGMENTS

Many thanks to Alfred Anwander for the multiple subject study. The development of the probabilistic fiber ODF tracking algorithm and generated results were closely done with his collaboration.

Part V

Conclusion

CONCLUSION

“The end crowns all, And that old common arbitrator, Time, Will one day end it.”

–William Shakespeare

In this thesis, we have first proposed new methods for local HARDI reconstruction using spherical harmonics: from the ADC modeling, the analytical q-ball diffusion ODF reconstruction, to sharper fiber ODF reconstructions using spherical deconvolution. Then, we have proposed two new methodological contributions that integrated this local HARDI reconstruction information for more global inference of the white matter architecture: a new segmentation algorithm and new deterministic and probabilistic tractography algorithms. Finally, we have focused on a specific neuroscientific application to illustrate the added value of our new HARDI-based fiber ODF probabilistic tractography algorithm. We were able to successfully reconstruct transcallosal fibers and quantify regions of the corpus callosum where the DTI-based algorithms have failed. These theoretical and methodological contributions were described and developed in the Methods part of the thesis and the tractography application was shown in a separate Application part.

These contributions required some background knowledge on cerebral white matter anatomy, diffusion MRI principles, HARDI reconstruction algorithms and the mathematical properties of the spherical harmonics. All these topics were reviewed and covered in the Background part of the thesis.

All along this thesis, we have tried to make the good mathematical and algorithmic choices to model and solve the problems of interest. First, we used the spherical harmonics basis with linear and regularized methods to perform HARDI estimations. This was done with the Funk-Hecke theorem and the Laplace-Beltrami operator. Then, we developed a new and efficient region-based surface evolution method to segment fiber bundles efficiently from the HARDI data. Finally, we proposed a fast deterministic HARDI tractography algorithm and a slower but more robust and more informative fiber ODF probabilistic tractography algorithm. Overall, we have tested and compared our methods on synthetic HARDI simulations, on a biological phantom and real datasets described in the Appendix A. We have made a special effort

to implement and discuss many existing state-of-the-art methods in the literature to highlight the strengths and limitations of our proposed methods.

Throughout the thesis, we have enumerated our major contributions and more specific contributions at the beginning and at the end of each chapters. We believe some of these contributions have had and will have more impact than others. In summary, the important and original contributions of the thesis are:

1. A robust and linear spherical harmonic estimation of the HARDI signal. This contribution is the foundation of the thesis and has been used at the basis of all our algorithms.
 - (a) A physically meaningful modified spherical harmonic basis taking into account the physical constraints of the HARDI data.
 - (b) A closed-form Laplace-Beltrami regularization term appropriate to smooth spherical functions.
 - (c) Optimal regularization parameter estimation with the L-curve method.
 - (d) Estimation of the ADC and HARDI signal and successful validation against ground truth.
 - (e) Theoretical link between the modified SH basis and the high order diffusion tensor.
 - (f) HARDI anisotropy measures computation and classification of voxels containing isotropic, single fiber and multiple fiber distributions.
2. A regularized, fast and robust analytical QBI solution. This is probably the most important contribution of the thesis. The QBI solution using spherical harmonics is now the standard in the literature.
 - (a) The analytical QBI solution was up to 15 times faster than numerical QBI solution.
 - (b) The proof of a new corollary of the Funk-Hecke theorem.
 - (c) The analytical QBI solution validation and comparison against classical numerical QBI solution and other analytical solution that use Tikhonov regularization.
3. A new statistical region-based segmentation framework for HARDI data segmentation based on the SH representation of the q-ball diffusion ODF. This contribution will have some applications for shape analysis of major fiber bundles in subject and patient populations.
 - (a) Efficient distance measure definition that captures similarities and dissimilarities between diffusion ODFs.

- (b) Robust and reproducible results of the segmentation algorithm on multiple subjects.
4. A new method to estimate the fiber ODF. This contribution is of interest for anyone working with the diffusion ODF. Moreover, an important message is that the fiber ODF is crucial to do tractography.
 - (a) A new analytical spherical deconvolution method that transforms the q-ball diffusion ODF into a sharper fiber ODF.
 - (b) Improvement of the angular resolution of QBI by 15° .
 - (c) The fiber ODF validation and extensive comparison against spherical deconvolution methods using spherical harmonics.
 - (d) Existing spherical deconvolution methods extension and improvement using our Laplace-Beltrami regularized HARDI signal estimation of 1.
 5. A new deterministic tractography algorithm and a new probabilistic tractography algorithm based on the fiber ODF of 4. This contribution is likely to have the most impact for applications in neuroscientific problems.
 - (a) Deterministic and probabilistic HARDI tractography comparisons on complex fiber bundles where the DTI tractography algorithms fail.
 6. Reconstruction of the transcallosal fibers intersecting with the corona radiata and superior longitudinal fasciculus. This contribution shows the success of our fiber ODF-based probabilistic tracking and can convince people to use our method.
 - (a) Quantification of regions on 8 subjects where probabilistic fiber ODF tracking can recover transcallosal fibers projecting to the lateral cortex, regions normally neglected by most DTI-based methods.

We believe that these contributions meet the initial goal of this thesis that was to propose new HARDI processing techniques able to recover fiber crossing information and overcome limitations of the DTI.

A central part of this thesis was the diffusion ODF and the fiber ODF. These spherical functions are able to solve the crossing problem and overcome limitations of the diffusion tensor. However, it is important to point out that the diffusion ODF and fiber ODF estimates also have limitations. Many different sub-voxel configurations can lead to the same diffusion ODF and fiber ODF profile (problem also mentioned in [Tuch et al. (2003); Campbell (2004); Perrin (2006)]). This is illustrated and sketched in Figure 12.1. We can see that within an imaging voxel, fibers can curve, fan/merge, cross, kiss/bottleneck or branch/merge and still produce a similar looking ODF. The exact diffusion ODF or fiber ODF that these fiber configurations produce on real data





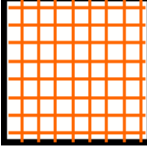


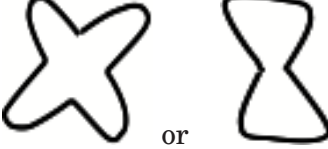

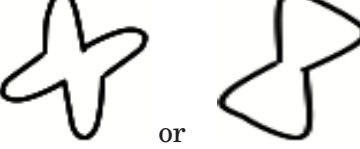
Sub-voxel fiber configurations	Expected ODF profile
 <p data-bbox="454 344 613 373">single curve</p>	
 <p data-bbox="415 554 659 583">fanning / merging</p>	
 <p data-bbox="480 764 591 793">crossing</p>	
 <p data-bbox="406 974 672 1003">bottleneck / kissing</p>	
 <p data-bbox="399 1184 672 1213">branching / merging</p>	

Figure 12.1: Sketch of the possible sub-voxel configurations and their expected diffusion ODF profile. We see curving single fibers, fanning or merging fibers, multiple crossing fibers, fibers in a bottleneck and multiple branching or merging fibers. Images courtesy of J. Campbell of McGill University in Montreal, Canada.

remains unknown and ambiguous. An even more important problem is that it is still unknown how to generate synthetic datasets, ex-vivo phantoms or biological phantoms that mimic these sub-voxel configurations. These datasets are crucial for the development and understanding of these sub-voxel configuration problems. It is a very local problem that can have global repercussions on tractography results. To our knowledge, only a few research groups have started to address these problems. [Savadjiev et al. (2007)] and [Seunarine et al. (2007)] propose methods that can account for fanning configurations in the tracking but results are still preliminary and the more complicated sub-voxel configurations are not addressed.

The problem with most of the sub-voxel configuration cases, except for the "pure"




Figure 12.2: A schematic illustration of sub-voxel fiber configurations that can cause ambiguous fiber ODF shapes and thus confound fiber tracking algorithms. Both a single curving fiber tract (top row) and a fanning fiber tract (bottom row) can yield a similar ODF with a single broad peak (middle column in each row). In the case of the curving fiber bundle, regardless of the incoming direction, tracking should follow the medial direction in the peak (top row, right). In the case of the fanning fiber bundle, however, when tracking enters the voxel in the fan direction (green arrow, bottom row, right panel), it should follow all directions within the extent of the fan (green cone within the voxel). If tracking enters the voxel from the opposite side of the fan (blue arrow), it should only follow the medial direction of the peak. This illustrates the importance of recovering the polarity of the fan in addition to its extent. Image taken from article [Savadjiev et al. (2008)] *under review*.

multiple fiber crossing, is that the sub-voxel geometry can produce a single peak ODF with a broad maxima. In this case, it is impossible to discriminate the multiple fiber orientations (see Figure 12.2). Hence, although the problem is very local, one must look in a neighborhood around the voxel of interest to infer the sub-voxel configuration. If we do so, we can obtain information about the curving angle of the tract locally (row 1 of Figure 12.2) as well as information about the possible cone of directions that the fanning tracts cover (row 2 in Figure 12.2). We believe that this information about the local geometry, curvature and torsion of the local curve estimates [Savadjiev et al. (2006)] will be needed to solve this important problem to

obtain more precise anatomical connectivity information.

Perspectives

We believe that the contributions from the thesis can now be applied to answer more neuroscientific questions. In fact, our algorithms are starting to be used by other groups in clinical settings, in part because they are available on demand through the BrainVisa software¹. In collaboration with A. Anwander and T. Knösche of the Max Planck Institute in Leipzig, Germany, we are using the diffusion ODF and fiber ODF reconstructions as well as probabilistic tractography to characterize complex fiber bundles involved in the language areas of the brain. Furthermore, we are interested in performing multiple subject study of the shape of major fiber bundles. We would like to be able to say what is a “normal” or “abnormal” corpus callosum or cingulum (or other bundles) and study how a population of subjects and patients can have different fiber bundles. Another collaboration is with S. Lehericy at the Pitié Salpêtrière Hospital in Paris. We are interested in the fiber tracts involved in the motor tasks. Both our deterministic and probabilistic tractography algorithms from QBI are being used to reconstruct tracts projecting from the basal nuclei to the motor cortex [Thiebaut et al. (2007)]. Finally, another new application is q-ball imaging of the spinal cord with J. Cohen-Adad and H. Benali of the INSERM in Paris, France and S. Rossignol and R. Hoge of the Université de Montréal in Canada ([Adad et al. (2008)]). Not only are we able to recover main longitudinal directions of the spine but we can also recover smaller dorso-ventral paths in the gray matter of a cat’s spinal cord.

Aside from these applications of HARDI, we can expect, in the near future, to see more research development in modeling, in new data acquisition techniques, in better visualization and in new algorithms. First, better models of diffusion in biological tissues such as white matter are needed. To do so, mathematicians and computer scientists will have to interact and surround themselves with biologists. The relation between the measured diffusion signal of a fibrous biological tissue and the actual fiber distribution needs to be refined and better known. Better models will allow for better reconstruction of fiber distributions.

Second, the arrival of 7 Tesla scanners will generate new datasets that will raise new image processing problems. New and better acquisitions techniques are needed for high magnetic field imaging, which have access to more powerful and faster gradients. We will be able to obtain better spatial resolution and higher quality DWI. Our current HARDI experiments in the gray matter of a cat gives us hope that we can soon attempt to do the same in the human brain gray matter with high resolution 7T diffusion MRI. We can also hope to get down to the cyto-architecture level to better study the structures of the basal nuclei.

¹Available on demand on <https://gforge.inria.fr> for the BrainVisa software <http://brainvisa.info>

Then, faster visualization algorithms will be needed to better interact with complex datasets coming from HARDI, q-ball and spherical deconvolution data. Work by [Kindlmann et al. (2007a,b)] aims at facilitating the visualization of HARDI measurements and spherical functions on the sphere. Note also that it is part of our intention to add q-ball diffusion and fiber ODF visualization and HARDI tracking as part of MedINRIA² [Toussaint et al. (2007a,b)] and 3D Slicer³ developed at Harvard University.

Next, improvement of local HARDI reconstruction methods need to be pushed further. As acquisition techniques become faster, the limitation in the number of gradient encoding directions will become less important. It is thus of interest to find the best trade-off between the number of directions used in the sampling scheme, the number of shell acquisitions and the b -value used to obtain best angular resolution while reducing effects of noise. This kind of study has been done in DTI-like acquisitions Jones (2004); Jones and Basser (2004) but needs to be revisited for HARDI-based reconstruction techniques. We believe that multiple shell acquisitions ([Khachaturian et al. (2007)]) will become more popular and provide better fiber orientation distributions. Real-time HARDI acquisitions also has a bright future, and in particular, real-time DTI and QBI [Poupon et al. (2007)] will give more insight into the problem of trade-off between the right sampling scheme and achievable angular resolution. Moreover, reconstruction algorithms need better noise models incorporated in their estimation. In diffusion tensor analysis, the latest state-of-the-art estimation include a Rician noise model [Fillard et al. (2007a)]. Surprisingly, while it is well-known that noise is Rician in DW images, none of the HARDI methods take this noise profile into account in the HARDI reconstruction.


Furthermore, even if we now have robust estimates of fiber orientation distributions from the HARDI data, there are remaining questions to be answered regarding post-processing of these spherical functions. First, peaks of these fiber distributions used for tracking have uncertainty in them due to noise in the data and actual spreading of the underlying fibers in the biological tissue. Hence, it is important to model this uncertainty in the peaks the fiber ODF estimates for calibration of the tractography algorithms. Work by [Parker and Alexander (2003); Friman and Westin (2005); Seunarine et al. (2007); Fonteijn et al. (2007); Haroon and Parker (2007)] have started investigating this issue, which might have important implications on tracking results. Another important question in post-processing HARDI data is to find the right distance metric to use to compare and compute statistics on spherical functions. In this thesis, we have used Euclidean distance between the spherical harmonic representation of the spherical functions but other distances such as the Kullback-Leibler divergence might be better suited. We believe that better distance measures will greatly improve segmentation and clustering algorithms of HARDI

²<http://www-sop.inria.fr/asclepios/software/MedINRIA/>


³<http://www.slicer.org/>

data. Another open problem is how to extract maxima from spherical functions obtained from HARDI. There are simple numerical finite difference methods and more complicated gradient-descent optimization methods to do so. Can maxima extraction be done analytically? We believe that exploring properties of the high order tensor (HOT) and its corresponding decomposition in eigen tensors, eigen vectors and eigen values might give insight into this problem. This research might lead to other sets of basis functions to better describe spherical functions. Currently, spherical harmonics are the popular mathematical tool used but other radial basis functions such as spherical Gaussians, HOTs or spherical wavelets might have other attractive properties.


Finally, we can also expect many more developments in segmentation, clustering and fiber tractography. It is our hope to study and compare these three different classes, i.e. segmentation, clustering and tractography algorithms. We have not discussed the clustering problem in this thesis (see [O'Donnell (2006); O'Donnell and Westin (2007); Leemans (2006); Kouby et al. (2005)]) but it does provide valuable information about white matter architecture. It is important to understand what information can be extracted precisely by HARDI clustering, HARDI segmentation and HARDI tractography and how they can improve one another. What is in common and complementary about these techniques? For example, Figure 12.3 shows an example of clustering, tracking and segmentation done on the same regions of complex fiber crossings between the corpus callosum (CC), corona radiata (CR) and superior longitudinal fasciculus (SLF). To do so, we will ultimately need better tools for validation. Ex-vivo and biological phantoms such as [Lin et al. (2003); Campbell et al. (2005); Perrin (2006)] are useful but their configurations are not complex enough. We believe that further development of realistic and complex phantoms will greatly help the validation problem. A special effort should also be put to make these datasets publicly available so that the community can have benchmarks to test and validate their algorithms. Only then, our work and algorithms will possibly convince clinicians and be integrated in hospitals and used for many neuroscientific applications.




GFA in coronal slice



Diffusion ODF clustering



Diffusion ODF tracking



Diffusion ODF segmentation

Figure 12.3: Diffusion ODF clustering, diffusion ODF tracking and diffusion ODF segmentation in similar regions of complex fiber architecture. Clustering was done as in [Wassermann et al. (2007b,a)] from the field of diffusion ODFs. 7 clusters are found indicated by different colors and overlaid on the diffusion ODFs. Deterministic tracking was done on the diffusion ODF from a few seed points placed in the CR and multiple seed points placed in the CC. The tracts are overlaid on a GFA with reduced opacity. Segmentation was done on the image of diffusion ODF in the corona radiata, as shown in Chapter 8. The information recovered from each algorithm is both in agreement and complementary. It is now important to make these methods interact with one another.

CONCLUSION (FRANÇAIS)

“En toute chose il faut considérer la fin”

–Jean de la Fontaine

Dans cette thèse, nous avons proposé de nouvelles méthodes pour le traitement de l'IRM à haute résolution angulaire (HARDI). Nous avons apporté des contributions au niveau de l'estimation locale du signal HARDI, de l'estimation du coefficient de diffusion apparent (ADC), de l'estimation de la distribution d'orientation des fibres de diffusion (ODF) à partir de l'imagerie par q-ball (QBI) et de l'estimation de l'ODF de fibres avec les méthodes de déconvolution sphérique. Nous avons démontré avec succès que toutes ces estimations apportent une information importante servant à retrouver les croisements de fibres, valeur ajoutée importante par rapport à l'imagerie par tenseur de diffusion (DTI). Cette thèse propose également plusieurs contributions algorithmiques au niveau de l'intégration de cette information locale obtenue par les reconstructions HARDI. Nous avons développé un nouvel algorithme de segmentation par contour actif basé sur les statistiques des harmoniques sphériques des ODF de diffusion dans la matière blanche. Nos segmentations arrivent à retrouver le tronc principal de certains faisceaux de fibres importants. Il est maintenant possible d'envisager l'étude de forme de ces faisceaux de fibres segmentés sur plusieurs sujets et patients. Par ailleurs, nous avons aussi introduit de nouveaux algorithmes de suivi de fibres adaptés aux données HARDI. Nous avons présenté un nouvel algorithme de suivi de fibres déterministe qui utilise tous les maxima de l'ODF de fibres ainsi qu'un nouveau suivi de fibres probabiliste qui exploite toute la distribution de cette ODF de fibres. Nos résultats de suivi de fibres retrouvent des réseaux de fibres complexes ayant des configurations de fibres en éventail et des fibres qui se croisent et s'embranchent. Ces nouvelles techniques nous ont alors permis de reconstruire, de façon robuste, les fibres transcallosales traversant le corps calleux et qui intersectent les fibres du faisceau supérieur longitudinal et les fibres de la couronne rayonnante. Toutes ces contributions ont été décrites en détail dans la partie *Méthodes* et la partie *Applications*.

Afin d'obtenir ces différents résultats, il a d'abord fallu introduire la théorie et les notions de base sur l'anatomie cérébrale, sur la physique de l'IRM de diffusion et sur les propriétés mathématiques des harmoniques sphériques. Cette section de la thèse était décrite dans la partie *Préliminaires et Théorie*.

Tout au long de cette thèse, nous nous sommes efforcés de faire les bons choix mathématiques et algorithmiques ainsi que les bons choix de modélisation pour résoudre les problèmes HARDI d'intérêt. Pour toutes nos nouvelles approches, nous avons fait l'effort de les placer dans le contexte de cette thèse ainsi que de les situer par rapport aux autres méthodes existantes. Nous avons donc dû mettre en oeuvre beaucoup de techniques de la littérature afin de les comparer à nos méthodes. Finalement, tous nos résultats ont été testés sur des données HARDI simulées, sur un fantôme biologique et sur des données réelles variées, tous décrits dans l'annexe A.

Le début et la fin de chacun des chapitres commençaient par nos contributions majeures et nos contributions plus spécifiques. Nous voulons maintenant en rappeler les grandes lignes dans l'ordre des chapitres de la thèse. Évidemment, certaines contributions ont eu et auront plus d'impact que d'autres et nous l'indiquons lorsque c'est approprié. Les contributions de la thèse sont les suivantes:

1. L'estimation robuste et linéaire du signal HARDI avec la base des harmoniques sphériques. Cette estimation en harmoniques sphériques est la fondation de toute la thèse.
 - (a) Définition d'une base modifiée d'harmoniques sphériques modélisant les contraintes physiques des données HARDI.
 - (b) Introduction d'un terme de régularisation analytique sur la sphère avec l'opérateur de Laplace-Beltrami.
 - (c) Détermination du paramètre optimal de régularisation à l'aide de la méthode L-curve.
 - (d) Lien théorique entre le tenseur d'ordre supérieur et la base modifiée des harmoniques sphériques.
 - (e) Estimation de l'ADC et calcul des mesures d'anisotropie HARDI qui en découlent afin de faire la classification des voxels isotropes, d'une seule distribution de fibres et d'une distribution de plusieurs fibres.
2. La solution analytique pour le QBI. Cette méthode est maintenant le standard, dans la littérature, pour faire du QBI. C'est certainement une contribution très importante et celle qui a eu et qui aura le plus d'impact.
 - (a) La solution est régularisée et obtenue à partir de l'estimation robuste du signal HARDI.
 - (b) La solution est environ quinze fois plus rapide que le QBI numérique classique [Tuch (2004)].

- (c) La preuve d'un nouveau corollaire du théorème de Funk-Hecke.
 - (d) La comparaison et la validation du QBI analytique par rapport à l'approche classique de QBI numérique et aux autres méthodes QBI.
3. La nouvelle méthode de segmentation par contour actif des données HARDI à partir de l'ODF de diffusion obtenue par la solution analytique du QBI. Cette contribution aura certainement quelques applications pour l'analyse de forme de certains larges faisceaux sur des populations de sujets et de patients.
- (a) Nouvelle métrique euclidienne définie sur les harmoniques sphériques pour calculer la distance entre deux ODFs.
 - (b) Segmentation robuste et efficace sur des faisceaux de fibres importants de la matière blanche.
 - (c) Segmentation automatique et répliquable sur plusieurs sujets.
4. La reconstruction de l'ODF de fibres à partir du QBI. Cette contribution est d'intérêt pour tous ceux qui font du QBI. De plus, et c'est l'un des rôles les plus importants, l'ODF de fibres est cruciale afin d'effectuer un bon suivi de fibres.
- (a) Solution analytique de l'ODF de fibres basée sur la solution analytique du QBI.
 - (b) Amélioration de la résolution angulaire du QBI de 15°.
 - (c) Lien théorique entre la déconvolution sphérique à partir du QBI et la déconvolution sphérique faite directement à partir du signal HARDI.
 - (d) Implémentation de toutes les méthodes courantes de déconvolution sphérique et comparaison avec notre estimation de l'ODF de fibres.
5. Nouveau suivi de fibres déterministe et nouveau suivi de fibres probabiliste à partir de l'ODF de fibres. Cette contribution aura un grand impact et certainement plusieurs applications en neuroscience.
- (a) Comparaison des deux nouvelles méthodes de suivi de fibres et mise en évidence des inconvénients et avantages de chacune.
 - (b) Valeur ajoutée du suivi de fibres à partir de l'ODF de fibres par rapport à l'ODF de diffusion et au tenseur de diffusion.
6. Reconstruction des fibres transcallosales qui intersectent le faisceau supérieur longitudinal et les fibres de la couronne rayonnante. Cette contribution illustre la valeur ajoutée du suivi de fibres probabiliste à partir des données HARDI et devrait encourager les neuroscientifiques à tirer parti du potentiel de nos méthodes de suivi de fibres.

- (a) Expérience quantitative soulignant les régions du corps calleux où les fibres transcallosales se projettent dans le cortex latéral. En comparaison, le suivi de fibres DTI n’y arrive pas.

Toutes ces contributions répondent bien au but initial de cette thèse de proposer de nouvelles méthodes de traitement d’images HARDI afin de retrouver les croisements de fibres et de surmonter les limites du DTI.

Une partie centrale de la thèse a été l’ODF de diffusion et l’ODF de fibres. Nous avons montré ses avantages par rapport au tenseur de diffusion. En revanche, au jour d’aujourd’hui, les ODFs de diffusion ou de fibres ont encore certaines limites importantes qui peuvent avoir un effet négatif sur les résultats de suivi de fibres. Comme le démontrent les Figures 12.1 et 12.2, différentes configurations sub-voxelique peuvent engendrer la même ODF (également souligné dans les thèses de [Tuch (2002); Campbell (2004); Perrin (2006)]). L’effet exact sur le signal HARDI et sur les méthodes de suivi de fibres qui en découlent de ces fibres qui s’embrassent, s’embranchent, forment des éventails ou qui se joignent est encore inconnu. Ce qui s’avère encore plus problématique est le fait que nous ne savons pas comment générer des données synthétiques réalistes pour reproduire ces configurations. Nous n’avons pas non plus de fantômes ex vivo ou in vivo avec ce genre de configurations. Il est donc très important de raffiner nos modèles et de développer des méthodes qui regardent l’information de façon plus globale autour du voxel d’intérêt afin de pouvoir inclure des notions locales de courbure et de torsion des fibres ([Savadjiev et al. (2006)]). Il s’agit d’ailleurs d’un de nos récents projets [Savadjiev et al. (2007)] et nous arrivons à identifier et à faire un meilleur suivi de fibres dans les voxels contenant des éventails et des fibres avec une grande courbure (voir Figure 12.2 et [Savadjiev et al. (2007)]). Nous croyons que cette information géométrique plus globale permettra de discriminer un plus grand nombre de ces configurations sub-voxeliques problématiques.

Perspectives

Nous pensons fortement que beaucoup de nos méthodes sont maintenant assez mûres pour être utilisées dans les équipes de recherche en neuroscience, et également en milieu clinique. En fait, quelques uns de nos algorithmes sont déjà utilisés par d’autres groupes de recherche, en partie grâce à la distribution gratuite de nos méthodes dans BrainVisa¹. Nous collaborons activement, par exemple, avec T. Knösche et A. Anwander du Max Planck Institute à Leipzig en Allemagne. Nous appliquons les techniques de suivi de fibres dans les régions de traitement du langage. Nous tentons également de faire une étude de population des faisceaux de fibres segmentés par notre approche. Nous travaillons aussi avec S. Lehericy de l’hôpital Pitié Salpêtrière à Paris sur les faisceaux impliqués dans les fonctions motrices. Enfin, nous avons une nouvelle application intéressante avec les gens de l’INSERM et de l’Université de

¹Disponible sur demande au <https://gforge.inria.fr>

Montréal sur le QBI de la moelle épinière chez le chat. Nous arrivons actuellement à trouver des connections dorso-ventrales dans la matière grise qui croisent la direction principale longitudinale de la moelle [Adad et al. (2008)]. Toutes ces collaborations sont importantes car elles permettent de mettre en valeur et de tester nos méthodes de même que de spécifier leurs limites.

Mises à part ces perspectives d'applications, on peut s'attendre à voir de plus en plus de nouveaux modèles qui décrivent plus finement la diffusion dans les tissus biologiques. Nous verrons également de nouvelles méthodes d'acquisitions en IRM de diffusion, de meilleurs logiciels de visualisation pour mieux interagir avec les données HARDI et énormément de développements algorithmiques dans un futur proche.

D'abord, nous avons besoin de meilleurs modèles qui décrivent plus finement la diffusion dans les tissus biologiques. Cela nous permettra d'imposer de meilleures contraintes et de faire des reconstructions HARDI plus précises de la distribution réelle des fibres. Pour ce faire, il est évident que mathématiciens, informaticiens et biologistes devront se regrouper et interagir de plus près.

Ensuite, avec les débuts de l'IRM à très haut champ magnétique, il y aura de nouveaux jeux de données qui soulèveront de nouvelles problématiques. On trouve de plus en plus d'IRM 7 Tesla dans les laboratoires de recherche, qui nécessitent le développement de meilleures acquisitions avec des gradients de diffusion plus puissants et plus rapides. Ces nouveaux scanners à haut champ nous permettront d'atteindre une résolution spatiale plus fine. Avec les nouvelles images issues de nos résultats récents sur la moelle épinière du chat, nous avons bon espoir de pouvoir observer les connexions au sein de la matière grise. Nous pouvons aussi espérer imager la cyto-architecture des noyaux gris centraux.

Il est de plus évident que les logiciels courants de visualisation (BrainVISA², 3D Slicer³, MedINRIA⁴, etc...) devront s'adapter aux nouvelles données complexes HARDI. De nouvelles méthodes rapides et efficaces devront être développées pour mieux interagir avec ces données. Nous sommes actuellement à ajouter le module de visualisation des ODFs et le suivi de fibres à MedINRIA [Toussaint et al. (2007a,b)] et 3D Slicer.

Il est également clair que les méthodes locales d'estimation HARDI continueront de s'améliorer au cours des prochaines années. Avec les progrès technologiques des imageurs, le problème du nombre de directions ne sera plus aussi important et il sera donc essentiel de trouver le meilleur compromis entre nombre de directions, valeur de b et nombre de "shell" d'acquisition dans le q -space, afin d'obtenir le meilleur rapport signal sur bruit et la meilleure résolution angulaire. Nous entrevoyons de belles perspectives en ce qui concerne les méthodes d'acquisition HARDI en temps réel [Poupon et al. (2007)] ainsi que les méthodes à "multi-shell", qui nous permet-

²<http://brainvisa.info/>

³<http://www.slicer.org/>

⁴<http://www-sop.inria.fr/asclepios/software/MedINRIA/>

tront d'obtenir de meilleures estimations de la distribution des fibres. Il est également important d'introduire les bons modèles de bruit dans nos algorithmes de reconstruction. Il est bien connu que le bruit est ricien dans les images pondérées en diffusion. Par ailleurs, l'état de l'art en analyse de tenseur de diffusion injecte un modèle de bruit ricien dans l'estimation [Fillard et al. (2007a)]. Il faudra donc intégrer cette composante dans nos futurs algorithmes de reconstruction.

Les algorithmes de suivi de fibres et de segmentation seront eux aussi améliorés par les développements mentionnés. En revanche, il reste encore des problèmes non résolus concernant le traitement des fonctions sphériques provenant des données HARDI. Il faudra d'abord intégrer des modèles d'incertitude dans le suivi de fibres afin de le calibrer par rapport à l'incertitude dans les directions principales des distributions de fibres estimées. Les travaux de [Parker and Alexander (2003); Friman and Westin (2005); Seunarine et al. (2007); Haroon and Parker (2007); Fonteijn et al. (2007)] proposent déjà quelques idées sur ce sujet. Un autre problème est de trouver la bonne métrique pour estimer des distances et faire des statistiques entre les ODFs. D'autres métriques, telles que la divergence de Kullback-Leibler, méritent d'être explorées. Une métrique mieux adaptée aura pour effet d'améliorer la segmentation et les algorithmes de regroupement (*clustering*). D'autre part, dans cette thèse, nous avons privilégié les harmoniques sphériques comme fonctions de base pour représenter les fonctions sphériques. Cependant, il est aussi important d'étudier d'autres fonctions comme les ondelletes sphériques, les gaussiennes sphériques ou les tenseurs d'ordres supérieurs qui peuvent elles aussi avoir des propriétés intéressantes pour le traitement des données HARDI. Une dernière problématique importante pour la communauté est l'extraction des maxima de ces fonctions sphériques. À l'heure actuelle, des méthodes numériques de différences finies simples ou des méthodes de descente de gradient sont utilisées pour détecter les maxima. Nous pensons que la décomposition des tenseurs d'ordres supérieurs en leurs tenseurs propres, vecteurs propres et valeurs propres donnera peut-être un moyen d'extraire les maxima de façon plus analytique et automatique.

Enfin, nous sommes vraiment intéressés par l'unification ou du moins par le rapprochement des méthodes de regroupement, de segmentation et de suivi de fibres. Nous avons peu abordé la question du regroupement dans cette thèse (voir [O'Donnell (2006); O'Donnell and Westin (2007); Leemans (2006); Kouby et al. (2005)]) mais il fournit également une information précieuse sur l'architecture de la matière blanche. Nous pensons qu'il y a vraiment un intérêt à combiner l'information provenant de ces trois classes d'algorithmes. Ces algorithmes fournissent une information commune et complémentaire qui se doit d'être mieux exploitée (voir Figure 12.3). À cette fin, il est clair que nous aurons besoin de meilleurs outils pour valider nos méthodes. Les fantômes *ex vivo* and *in vivo* comme ceux de [Lin et al. (2003); Campbell et al. (2005); Perrin (2006)] sont vraiment utiles et doivent continuer à être développés avec une plus grande complexité géométrique. Un effort particulier doit être égale-

ment fait pour rendre ces données accessibles à la communauté, un peu comme nous, informaticiens, tentons de distribuer nos logiciels librement. C'est seulement à cette condition que nous pourrions valider et tester nos méthodes et ainsi, peut-être convaincre les cliniciens et les médecins d'utiliser nos approches dans leurs applications neuroscientifiques.

Part VI

Appendix

SYNTHETIC HARDI DATA GENERATION AND HARDI ACQUISITIONS

A.1 SYNTHETIC DATA GENERATION

To evaluate the performance of the ADC, q-ball diffusion ODF and fiber ODF estimation as well as the performance of the tractography and segmentation algorithms, we generate synthetic data using the multi-tensor model. This model was also used in [Alexander et al. (2002); Hess et al. (2006); Tuch (2004); Alexander and Barker (2005); Descoteaux et al. (2006b)]. This model assumes that single fiber responses can be described by a Gaussian (rank-2 tensor) and that the tissue in a voxel with more than one fiber is simply composed of multiple Gaussian fibers that do not exchange molecules. The precise steps to generate the DWI signal are as follows:

1. Set diffusion encoding gradient directions $\mathbf{g}_i = (\sin \theta_i \cos \phi_i, \sin \theta_i \sin \phi_i, \cos \theta_i)^T$ using certain order tessellation of the icosahedron with N samples on the sphere.
2. Set the number of fibers to a specific n between 0 and 3 (0 is for isotropic signal).
3. For each fiber $1 \leq k \leq n$, randomly choose an orientation (θ_k, ϕ_k) and a relative weight p_k for the k -th fiber, where $\sum_{k=1}^n p_k = 1$. In practice, we impose a minimum angle between fibers of 45° and relative weights between 0.3 and 0.7 for 2-fiber and between 0.2 and 0.4 for 3-fiber distributions. Otherwise, the fiber compartments are too weak to distinguish them.
4. For the isotropic voxels, we use a profile with eigenvalues $[700, 700, 700] \times 10^{-6}$ mm²/s. For the other distributions, we generate 3x3 tensors with eigenvalues $[\lambda_1, \lambda_2, \lambda_3] = [300, 300, 1700] \times 10^{-6}$ mm²/s and generate each fiber tensor \mathbf{D}_k using rotation matrix \mathbf{Rot} to orient the main axis of the

tensor in direction (θ_k, ϕ_k) , i.e. $\mathbf{D}_k = \mathbf{Rot}_{(\theta_k, \phi_k)}^T \text{diag}(\lambda_1, \lambda_2, \lambda_3) \mathbf{Rot}_{(\theta_k, \phi_k)}$. Note that we choose tensor profiles with anisotropy nearly as high as the most anisotropic regions found in the human brain. This facilitates the discrimination of 2- and 3-fiber distributions. As in [Tournier et al. (2004)], the underlying assumption is that lower anisotropy levels are due to noise and partial volume averaging from fibers crossing, kissing or diverging.

5. Pick a b -value and for each encoding direction i generate the diffusion signal

$$S(b, \mathbf{g}_i) = S_0 \sum_{k=1}^n p_k e^{-b \mathbf{g}_i^T \mathbf{D}_k \mathbf{g}_i}. \quad (\text{A.1})$$

6. Add complex Gaussian noise [Sijbers et al. (1998)] with a standard deviation (std) of $\sigma = S_0/\zeta$ to the raw signal of Eq. [A.1]. That is, to each $S(\mathbf{g}_i)$, we add a random complex number with independent real and imaginary parts coming from a zero mean Gaussian distributions with $\sigma = 1/\zeta$ ($S_0 = 1$) and take the modulus to obtain the noisy synthetic data. A typical value is $\zeta = 35$ so that the unweighted signal S_0 has a signal-to-noise ratio (SNR) of 35. [Alexander et al. (2002); Tuch (2002); Alexander and Barker (2005)] also use a similar noise level even though typical real HARDI acquisitions tend to be corrupted by more noise. In our simulations, we need a relatively low noise level in order to characterize 2- and 3-fiber distributions.

This makes computation using this model relatively straightforward. We are trying to obtain the ADC profile, however, rather than the signal itself and so, we wish to find $D(\mathbf{g})$ such that

$$D(\mathbf{g}) = -\frac{1}{b} \ln \left(\frac{S(b, \mathbf{g})}{S_0} \right) = -\frac{1}{b} \ln \left(\sum_{k=1}^n p_k e^{-b \mathbf{g}^T \mathbf{D}_k \mathbf{g}} \right). \quad (\text{A.2})$$

It is worth mentioning that the hindered cylinder model is also another popular synthetic data generation model used in other works [von dem Hagen and Henkelman (2002); Ozarslan et al. (2004b, 2005b)]. In this formulation, [Soderman and Jonsson (1995)] assume that fibers are perfect cylinders and that water molecules are confined to diffuse within the walls of these cylinders. In the presence of multiple fibers, the signal attenuation from the cylinders is additive. This gives a more physically-based mixture model but is less realistic than the multi-tensor model that corresponds more closely to empirical fiber data. In any case, the approaches have the same model-selection and fitting problems [Alexander (2006)].

Another advantage of the multi-tensor model is that the exact ODF can easily be computed analytically. It is given in [Tuch (2004)] for a single Gaussian fiber and by linearity, we can easily obtain the exact ODF for n fibers and normalization

constant Z ,

$$\Psi(\mathbf{u}_i) = \sum_{k=1}^n \frac{p_k}{Z} \sqrt{\frac{\pi b}{\mathbf{u}_i^T \mathbf{D}_k^{-1} \mathbf{u}_i}}. \quad (\text{A.3})$$

This expression is formally derived in Appendix 7.9.

A.2 BIOLOGICAL PHANTOM DATA ---

The biological phantom was produced by Jennifer Campbell at the McConnell Brain Imaging Center (BIC) [Campbell et al. (2005)], McGill University, Montréal, Canada on a 1.5T Sonata MR scanner using a knee coil. It was created from two excised Sprague-Dawley rat spinal cords embedded in 2% agar. The acquisition was done with a single-shot spin-echo planar sequence with twice-refocused balanced gradients, designed to reduce eddy current effects. The dataset is acquired with 90 pairs of points generated using electrostatic repulsion algorithm [Jones et al. (1999)]. The q-ball protocol was used with $b = 3000 \text{ s/mm}^2$, $q = 0.35 \mu\text{m}^{-1}$, $\text{TR} = 6.4 \text{ s}$, $\text{TE} = 110 \text{ ms}$, $\text{FOV} 360 \times 360 \text{ mm}^2$, 128×128 matrix, 2.8 mm isotropic voxels and four signal averages per direction. The SNR of the $b = 0$ image S_0 was estimated to be approximately 70 for the averaged phantom and around 10 for the cord at $b = 3000 \text{ s/mm}^2$.

A.3 HUMAN BRAIN DATA ---

Several datasets were generously given to us to test our new local reconstruction algorithms as well as our new tractography and new segmentation algorithms. We would like to thank to Bruce Pike, V. V. Rymar and Jennifer Campbell of the Brain Imaging Center at McGill University, Canada for the human brain and biological phantom datasets. Thanks also to Thomas Knösche, Timm Wetzel and Alfred Anwander of the Max Planck Institute for Human Cognitive and Brain Sciences, Leipzig, Germany for their human brain datasets. Thanks also to Cyril Poupon and Jean-François Mangin and colleagues from NeuroSpin / CEA for the HARDI database of 13 subjects. Finally, thanks to G. Sapiro of the University of Minnesota and Kamil Ugurbil and Stéphane Lehericy of the Center for Magnetic Resonance Research in Minneapolis, USA. Thanks again to all these collaborators, without whom this thesis would not have been possible.

A.3.1 Data from the BIC, McGill University Montréal, Canada

Just like the biological phantom, the human brain dataset was acquired by Jennifer Campbell from the BIC. A Siemens 3T MR scanner 8 channel head coil was used, using 99 pairs of points generated with the electrostatic repulsion algorithm [Jones et al. (1999)], has 63 slices of 2 mm each, covering the entire cerebrum. The FOV was $256 \times 256 \text{ mm}^2$, $\text{TR} = 11.1 \text{ s}$, $\text{TE} = 121 \text{ ms}$, $b = 3000 \text{ s/mm}^2$ ($q = 0.35 \mu\text{m}^{-1}$), $\text{BW} 1346 \text{ Hz/Pixel}$, 128×128 matrix and phase partial Fourier 7/8.

Ten $b = 0$ images were acquired and averaged to produce the S_0 image. The SNR in the white matter of this S_0 image was estimated to be approximately 41.

A.3.2 Data from the Center for Magnetic Resonance Research (CMRR)

DWI were acquired at the CMRR, University of Minnesota, Minneapolis, USA, on a 3T Siemens Magnetom Trio whole-body scanner. We generate 162 gradient directions with a 3rd order tessellation of the icosahedron, as seen in Figure 4.12. In practice, assuming symmetry of the diffusion process, we use only one direction out of two to obtain 81 gradient directions on the sphere. We acquired 3 repetitions per direction, each with $b = 1000$ s/mm², TR = 5100 s and TE = 109 ms. The three measurements are averaged by default by the scanner to produce 81 individual measurements to process. The voxel size was 3 mm³ cube and there were 24, 64 x 64 slices. The SNR of the T2-weighted data was estimated to be approximately 30.

A.3.3 Data from the Max Planck Institute (MPI), Leipzig, Germany

DWI data and high-resolution T1-weighted images were acquired in 8 healthy right-handed volunteers (25 ± 4 years, 4 females) on a whole-body 3T Magnetom Trio scanner (Siemens, Erlangen) equipped with an 8-channel head array coil [Anwander et al. (2007b)]. Written informed consent was obtained from all subjects in accordance with the ethical approval from the University of Leipzig. The spin-echo echo-planar-imaging sequence, TE = 100 ms, TR = 12 s, 128 x 128 image matrix, FOV = 220 x 220 mm², consists of 60 diffusion encoding gradients [Jones et al. (1999)] with a b-value of 1000 s/mm². Seven images without any diffusion weightings are placed at the beginning of the sequence and after each block of 10 diffusion weighted images as anatomical reference for offline motion correction. The measurement of 72 slices with 1.7mm thickness (no gap) covered the whole brain. Random variations in the data were reduced by averaging 3 acquisitions, resulting in an acquisition time of about 45 minutes. No cardiac gating was employed to limit the acquisition time. The issue of cardiac gating is discussed in [Jones et al. (2002)]. The SNR in the white matter of this S_0 image was estimated to be approximately 37. Additionally, fat saturation was employed and we used 6/8 partial Fourier imaging, a Hanning window filtering and parallel acquisition (generalized auto-calibrating partially parallel acquisitions, reduction factor = 2) in the axial plane.

The brain is peeled from the T1-anatomy, which was aligned with the Talairach stereotactical coordinate system [Talairach and Tournoux (1988)]. The 21 images without diffusion weightings distributed within the whole sequence were used to estimate motion correction parameters using rigid-body transformations [Jenkinson et al. (2002)], implemented in [FSL (2006)]. The motion correction for the 180 diffusion-weighted images was combined with a global registration to

the T1 anatomy computed with the same method. The gradient direction for each volume was corrected using the rotation parameters. The registered images were interpolated to the new reference frame with an isotropic voxel resolution of 1.72 mm and the 3 corresponding acquisitions and gradient directions were averaged.

A.3.4 Public HARDI Database from NeuroSpin / CEA, Paris, France

Finally, the public HARDI database of [Poupon et al. (2006)] is used. This HARDI database can be obtained by contacting Cyril Poupon directly by email, cyril.poupon@cea.fr. The 13 datasets were acquired on a 1.5T scanner with 200 encoding directions, $b = 3000 \text{ s/mm}^2$, 60 slices with 2 mm thickness, twenty five $b = 0 \text{ s/mm}^2$ images, 128 x 128 image matrix, TE = 93.2 ms, TR = 1.9 s.

PUBLICATIONS OF THE AUTHOR ARISING FROM THIS WORK

Journal papers

1. Maxime Descoteaux and Rachid Deriche. *Q-Ball Images Segmentation Using Region-Based Statistical Surface Evolution*. Journal of Mathematical Imaging in Vision, special issue on Mathematics in Image Analysis, in press 2008.
2. Peter Savadjiev, Jennifer S. W. Campbell, Maxime Descoteaux, Rachid Deriche, G. Bruce Pike and Kaleem Siddiqi *Labeling of ambiguous sub-voxel fibre bundle configurations in high angular resolution diffusion MRI*. NeuroImage, in press 2008.
3. Maxime Descoteaux, Elaine Angelino, Shaun Fitzgibbons, Rachid Deriche. *Regularized, Fast and Robust Analytical Q-Ball Imaging*, Magnetic Resonance in Medicine, Volume 58, Issue 3, Pages 497-510, 2007
4. Maxime Descoteaux, Elaine Angelino, Shaun Fitzgibbons, Rachid Deriche. *Apparent Diffusion Coefficients from High Angular Resolution Diffusion Images: Estimation and Applications*. Magnetic Resonance in Medicine Volume 56, Issue 2, Pages 395-410, August 2006.
5. Demian Wassermann, Maxime Descoteaux, Rachid Deriche. *Diffusion Maps Clustering for Magnetic Resonance Q-Ball Imaging Segmentation*. International Journal on Biomedical Imaging, special issue on Recent Advances in NeuroImaging Methodology, in press 2007.
6. (under review) Maxime Descoteaux, Rachid Deriche, Alfred Anwander. *Deterministic and Probabilistic Tractography: from Diffusion to Sharp Fiber Distributions*. Submitted to IEEE Transactions in Medical Imaging.

Conference papers

1. Maxime Descoteaux and Rachid Deriche. *Segmentation of Q-Ball Images Using Statistical Surface Evolution*. Medical Image Computing and Computer Assisted Intervention (MICCAI) 2007, Brisbane, Australia, April 2007.

2. Demian Wassermann, Maxime Descoteaux and Rachid Deriche. *Diffusion Maps Segmentation of Magnetic Resonance Q-Ball Imaging*. Mathematical Methods in Biomedical Image Analysis (MMBIA) Workshop 2007, held in conjunction with the 11th International Conference on Computer Vision (ICCV), Rio, Bresil, October 2007.
3. Rachid Deriche and Maxime Descoteaux. *Splitting Tracking through crossing fibers: Multidirectional Q-Ball Tracking*. 4th IEEE International Symposium on Biomedical Imaging: From Nano to Macro (ISBI'07) , Arlington, Virginia, USA, April 2007
4. Maxime Descoteaux, Peter Savadjiev, Jennifer Campbell, G. Bruce Pike, Kaleem Siddiqi, Rachid Deriche. *Validation and Comparison of Analytical Q-Ball Imaging Methods*. 4th IEEE International Symposium on Biomedical Imaging: From Nano to Macro (ISBI'07) , Arlington, Virginia, USA, April 2007.
5. Maxime Descoteaux, Rachid Deriche, Christophe Lenglet. *Diffusion Tensor Sharpening Improves White Matter Tractography*. SPIE Image Processing: Medical Imaging, San Diego, California, USA, February 2007.
6. Maxime Descoteaux, Elaine Angelino, Shaun Fitzgibbons, Rachid Deriche. *A Fast and Robust ODF Estimation Algorithm in Q-Ball Imaging* . 3rd IEEE International Symposium on Biomedical Imaging: From Nano to Macro (ISBI'06), April 2006.
7. Maxime Descoteaux, Elaine Angelino, Shaun Fitzgibbons, Rachid Deriche. *Apparent Diffusion Profile Estimation From High Angular Resolution Diffusion Images*. SPIE Image Processing: Medical Imaging, February 2006.

Conference abstracts

1. (under review) Julien-Cohen Adad and Maxime Descoteaux and Rachid Deriche and Serge Rossignol and Rick D. Hoge and Habib Benali. *Q-Ball Imaging of the Spinal Cord*. Submitted to ISMRM 2008.
2. Alfred Anwander, Maxime Descoteaux and Rachid Deriche. *Probabilistic Q-Ball Tractography Solves Crossings of Callosal Fibers*. Human Brain Mapping, Chicago, USA, June 2007.
3. Maxime Descoteaux and Rachid Deriche. *Sharpening Improves Clinically Feasible Q-Ball Imaging Reconstructions*. Joint Annual Meeting ISMRM-ESMRMB, Berlin, Germany, May 19-25th 2007.
4. Peter Savadjiev, Jennifer Campbell, Maxime Descoteaux, Rachid Deriche, G. Bruce Pike, Kaleem Siddiqi. *Disambiguation of Complex Subvoxel Fibre Configurations in High Angular Resolution Fibre Tractography*. Joint Annual Meeting ISMRM-ESMRMB, Berlin, Germany, May 19-25th 2007.

5. Maxime Descoteaux, Rachid Deriche , Peter Savadjiev, Jennifer Campbell, Bruce Pike, Kaleem Siddiqi. *Analytic ODF Estimation and Validation in Q-Ball Imaging*. 12th annual meeting of the Organization for Human Brain Mapping (HBM), Florence, Italie, June 11-15th 2006.

Research Reports

1. Demian Wassermann, Maxime Descoteaux and Rachid Deriche. *Recovering cerebral white matter structures with Spectral Clustering of Diffusion MRI Data*. INRIA Research Report 6351, November 2007.
2. Maxime Descoteaux, Rachid Deriche and Alfred Anwander. *Deterministic and Probabilistic Q-Ball Tractography: from Diffusion to Sharp Fiber Distributions*. INRIA Research Report 6273, August 2007.
3. Maxime Descoteaux and Rachid Deriche. *Q-Ball Images Segmentation Using Region-Based Statistical Surface Evolution*. INRIA Research Report 6257, July 2007.
4. Maxime Descoteaux, Elaine Angelino, Shaun Fitzgibbons, Rachid Deriche. *A Linear and Regularized ODF estimation algorithm to recover multiple fibers in Q-Ball Imaging*. INRIA Research Report 5768, November 2005.
5. Maxime Descoteaux, Elaine Angelino, Shaun Fitzgibbons, Rachid Deriche. *Apparent Diffusion Coefficients from High Angular Resolution Diffusion Images: Estimation and Applications* . INRIA Research Report 5681, September 2005.

Bibliography

- Adad, J.-C., M. Descoteaux, R. Deriche, S. Rossignol, R. D. Hoge, and H. Benali: 2008, 'Q-Ball Imaging of the Spinal Cord'. In: *Submitted to ISMRM*. [244, 253]
- Alexander, A. L., K. M. Hasan, M. Lazar, J. S. Tsuruda, and D. L. Parker: 2001, 'Analysis of partial volume effects in diffusion-tensor MRI'. *Magnetic Resonance in Medicine* **45**(4), 770–780. [55]
- Alexander, D.: 2005a, 'Multiple-fibre reconstruction algorithms for diffusion MRI'. *Annals of the New York Academy of Sciences* **1046**, 113–133. [36]
- Alexander, D.: 2006, *An Introduction to Diffusion MRI: the Diffusion Tensor and Beyond*. Springer. [260]
- Alexander, D. and G. Barker: 2005, 'Optimal imaging parameters for fiber-orientation estimation in diffusion MRI'. *NeuroImage* **27**, 357–367. [86, 259, 260]
- Alexander, D., G. Barker, and S. Arridge: 2002, 'Detection and Modeling of Non-Gaussian Apparent Diffusion Coefficient Profiles in Human Brain Data'. *Magnetic Resonance in Medicine* **48**(2), 331–340. [53, 54, 55, 68, 69, 75, 83, 84, 87, 259, 260]
- Alexander, D. C.: 2005b, 'Maximum Entropy Spherical Deconvolution for Diffusion MRI'. In: *Image Processing in Medical Imaging*. pp. 76–87. [55, 56, 59, 107, 167]
- Anderson, A.: 2005, 'Measurements of Fiber Orientation Distributions Using High Angular Resolution Diffusion Imaging'. *Magnetic Resonance in Medicine* **54**, 1194–1206. [8, 18, 56, 57, 69, 107, 109, 110, 114, 135, 166, 167, 168, 169, 170, 171, 172, 178, 184, 197, 198, 199]
- Andrews, G., R. Askey, and R. Roy: 1999, *Special Functions*. Cambridge University Press. [68]
- Anwander, A., M. Descoteaux, and R. Deriche: 2007a, 'Probabilistic Q-Ball Tractography Solves Crossings of Callosal Fibers'. In: *Human Brain Mapping*. Chicago, Illinois, USA. [236]
- Anwander, A., M. Tittgemeyer, D. Y. von Cramon, A. D. Friederici, and T. R. Knosche: 2007b, 'Connectivity-Based Parcellation of Broca's Area'. *Cerebral Cortex* **17**(4), 816–825. [208, 212, 213, 216, 217, 224, 225, 262]

- Arfken, G. B. and H. J. Weber: 1995, *Mathematical Methods for Physicists, 4th ed.* Academic Press. [63, 69, 152]
- Arsigny, V.: 2006, 'Processing Data in Lie Groups: An Algebraic Approach. Application to Non-Linear Registration and Diffusion Tensor MRI'. Ph.D. thesis, Ecole polytechnique. [45]
- Arsigny, V., P. Fillard, X. Pennec, and N. Ayache: 2006, 'Log-Euclidean Metrics for Fast and Simple Calculus on Diffusion Tensors'. *Magnetic Resonance in Medicine* **56**(2), 411–421. [99, 176]
- Assaf, Y. and P. Basser: 2005, 'Composite hindered and restricted model of diffusion (CHARMED) MR imaging of the human brain'. *NeuroImage* **27**(1), 48–58. [55]
- Assaf, Y., R. Z. Freidlin, G. K. Rohde, and P. J. Basser: 2004, 'New Modeling and Experimental Framework to Characterize Hindered and Restricted Water Diffusion in Brain White Matter'. *Magnetic Resonance in Medicine* **52**, 965–978. [55]
- Assaf, Y., A. Mayk, and Y. Cohen: 2000, 'Displacement Imaging of Spinal Cord Using q-Space Diffusion-Weighted MRI'. *Magnetic Resonance in Medicine* **44**, 713–722. [51]
- Backus, G.: 1964, 'Geophysical interpolation of measurements of average phase velocities of surface waves over great circular and semi-circular paths'. *Bulletin of the Seismological Society of America* **54**, 571–610. [110]
- Barmpoutis, A., B. Jian, and B. C. Vemuri: 2007, 'Symmetric Positive 4th Order Tensors & their Estimation from Diffusion Weighted MRI'. In: *Information Processing in Medical Imaging (IPMI)*. [100]
- Basser, P.: 2002, 'Relationships between Diffusion Tensor and q-Space MRI'. *Magnetic Resonance in Medicine* **47**, 392–397. [44]
- Basser, P., J. Mattiello, and D. LeBihan: 1992, 'Diagonal and off-diagonal components of the self-diffusion tensor: their relation to and estimation from the NMR spin-echo signal'. In: *11th Society of Magnetic Resonance in Medicine Meeting*. p. 1222. [44]
- Basser, P., J. Mattiello, and D. LeBihan: 1994a, 'Estimation of the effective self-diffusion tensor from the NMR spin echo'. *Journal of Magnetic Resonance* **B**(103), 247–254. [3, 12, 44, 45, 76, 107, 207]
- Basser, P., J. Mattiello, and D. LeBihan: 1994b, 'MR Diffusion Tensor Spectroscopy and imaging'. *Biophysical Journal* **66**(1), 259–267. [3, 12, 38, 44]
- Basser, P., J. Mattiello, R. Turner, and D. L. Bihan: 1993, 'Diffusion tensor echo-planar imaging of human brain'. In: *Proceedings of the SMRM*. p. 584. [3, 12, 15, 44]

- Basser, P., S. Pajevic, C. Pierpaoli, J. Duda, and A. Aldroubi: 2000, 'In Vivo Fiber Tractography Using DT-MRI Data'. *Magnetic Resonance in Medicine* **44**, 625–632. [46, 208, 210]
- Basser, P. and C. Pierpaoli: 1996, 'Microstructural and Physiological Features of Tissues Elucidated by Quantitative Diffusion Tensor MRI'. *Journal of Magnetic Resonance* **111**(3), 209–219. [75]
- Basser, P. J. and D. K. Jones: 2002, 'Diffusion-tensor MRI: theory, experimental design and data analysis - a technical review'. *NMR in Biomedicine* **15**, 456–467. [44]
- Basser, P. J. and S. Pajevic: 2007, 'Spectral decomposition of a 4th-order covariance tensor: Applications to diffusion tensor MRI'. *Signal Processing* **87**, 220–236. [100]
- Beaulieu, C.: 2002, 'The basis of anisotropic water diffusion in the nervous system - a technical review'. *NMR in Biomedicine* **15**, 435–455. [4, 24, 27, 38]
- Behrens, T. E. J., H. Johansen-Berg, S. Jbabdi, M. F. S. Rushworth, and M. W. Woolrich: 2007, 'Probabilistic diffusion tractography with multiple fibre orientations. What can we gain?'. *NeuroImage* **34**(1), 144–155. [4, 12, 55, 188, 194, 208, 209]
- Behrens, T. E. J., M. W. Woolrich, M. Jenkinson, H. Johansen-Berg, R. G. Nunes, S. Clare, P. M. Matthews, J. M. Brady, and S. M. Smith: 2003, 'Characterization and Propagation of Uncertainty in Diffusion-Weighted MR Imaging'. *Magnetic Resonance in Medicine* **50**, 1077–1088. [209, 225]
- Bergmann, Ø., G. Kindlmann, A. Lundervold, and C.-F. Westin: 2006, 'Diffusion k-tensor estimation from Q-ball imaging using discretized principal axes'. In: *Ninth International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI'06)*. Copenhagen, Denmark, pp. 268–275. [167]
- Bergmann, Ø., G. Kindlmann, S. Peled, and C.-F. Westin: 2007, 'TWO-TENSOR FIBER TRACTOGRAPHY'. In: *4th International Symposium on Biomedical Imaging*. Arlington, Virginia, USA, pp. 796–799. [208]
- Bloch, F.: 1946, 'Nuclear Induction'. *Physical Review* **70**, 460–474. [39]
- Blyth, R., P. Cook, and D. Alexander: 2003, 'Tractography with Multiple Fibre Directions'. In: *Proceedings of the International Society of Magnetic Resonance in Medicine*. Toronto, Canada, p. 240. [55, 84]
- Brechbuhler, C., G. Gerig, and O. Kubler: 1995, 'Parametrization of Closed Surfaces for 3D Shape Description'. *Computer Vision and Image Understanding* **61**, 154–170. [68]

- Brown, R.: 1828, 'A brief account of microscopical observations made in the months of June, July and August, 1827, on the particles contained in the pollen of plants; and on the general existence of active molecules in organic and inorganic bodies'. *Phil. Mag.* **4**, 161–173. [37]
- Bulow, T.: 2004, 'Spherical Diffusion for 3D Surface Smoothing'. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **26**(12), 1650–1654. [173]
- Callaghan, P. T.: 1991, *Principles of nuclear magnetic resonance microscopy*. Oxford: Oxford University Press. [42, 48, 51]
- Callaghan, P. T., C. D. Eccles, and Y. Xia: 1988, 'RAPID COMMUNICATION: NMR microscopy of dynamic displacements: k-space and q-space imaging'. *Journal of Physics E Scientific Instruments* **21**, 820–822. [48, 51]
- Campbell, J., K. Siddiqi, V. Rymer, A. Sadikot, and G. Pike: 2005, 'Flow-based fiber tracking with diffusion tensor q-ball data: Validation and comparison to principal diffusion direction techniques'. *NeuroImage* **27**(4), 725–736. [19, 95, 114, 207, 211, 246, 254, 261]
- Campbell, J. S.: 2004, 'Diffusion Imaging of White Matter Fiber Tracts'. Ph.D. thesis, McGill University. [36, 38, 241, 252]
- Campbell, J. S. W., P. Savadjiev, K. Siddiqi, and G. B. Pike: 2006, 'Validation and Regularization in Diffusion MRI Tractography'. In: *Third IEEE International Symposium on Biomedical Imaging (ISBI): from Nano to Macro*. Arlington, Virginia, USA, pp. 351–354. [207, 209]
- Carr, H. Y. and E. M. Purcell: 1954, 'Effects of Diffusion on Free Precession in Nuclear Magnetic Resonance Experiments'. *Physical Review* **94**(3), 630–638. [39]
- Catani, M., R. J. Howard, S. Pajevic, and D. K. Jones: 2002, 'Virtual in Vivo Interactive Dissection of White Matter Fasciculi in the Human Brain'. *NeuroImage* **17**, 77–94. [219]
- Cercignani, M. and M. Horsfield.: 2001, 'The physical basis of diffusion-weighted MRI'. *J. Neurol.* **186**(1), S11–S14. [41]
- Chan, T. and L. Vese: 2001, 'Active contours without edges'. *IEEE Transactions on Image Processing* **10**(2), 266–277. [146, 147, 150, 152]
- Chao, Y.-P., C.-Y. Yang, K.-H. Cho, C.-H. Yeh, K.-H. Chou, J.-H. Chen, and C.-P. Lin: 2007a, 'Probabilistic Anatomical Connection Derived from QBI with MFACT Approach'. In: *International Conference on Functional Biomedical Imaging*. Hangzhou, China. [209, 210]

- Chao, Y.-P., C.-H. Yeh, K.-H. Cho, J.-H. Chen, and C.-P. Lin: 2007b, ‘Multiple Streamline Tractography Approach with High Angular Resolution Diffusion Imaging Data’. In: *Proceedings of the International Society of Magnetic Resonance in Medicine*. Berlin, Germany, p. 1550. [208]
- Chefd’hotel, C., D. Tschumperlé, R. Deriche, and O. Faugeras: 2004, ‘Regularizing Flows for Constrained Matrix-Valued Images’. *Journal of Mathematical Imaging and Vision* **20**(1-2), 147–162. [99, 176]
- Chen, Y., W. Guo, Q. Zeng, G. He, B. Vemuri, and Y. Liu: 2004a, ‘Recovery of Intra-Voxel Structure From HARD DWI’. In: *ISBI*. pp. 1028–1031, IEEE. [55]
- Chen, Y., W. Guo, Q. Zeng, X. Yan, F. Huang, H. Zhang, G. He, B. Vemuri, and Y. Liu: 2004b, ‘Estimation, Smoothing, and Characterization of Apparent Diffusion Coefficient Profiles From High Angular Resolution DWI’. In: *Computer Vision and Pattern Recognition*, Vol. 1. pp. 588–593. [54, 75, 83]
- Chen, Y., W. Guo, Q. Zeng, X. Yan, M. Rao, and Y. Liu: 2005, ‘Apparent Diffusion Coefficient Approximation and Diffusion Anisotropy Characterization in DWI’. In: *Information Processing in Medical Imaging*. pp. 246–257. [84]
- Clayden, J. D., M. E. Bastin, and A. J. Storkey: 2006, ‘Improved segmentation reproducibility in group tractography using a quantitative tract similarity measure’. *NeuroImage* **33**, 482–492. [163]
- Clayden, J. D., A. J. Storkey, and M. E. Bastin: 2007, ‘A Probabilistic Model-based Approach to Consistent White Matter Tract Segmentation’. *IEEE Transactions on Medical Imaging* p. in press. [163]
- Cohen, Y. and Y. Assaf: 2002, ‘High b-value q-space analyzed diffusion-weighted MRS and MRI in neuronal tissues - a technical review’. *NMR in Biomedicine* **15**, 516–542. [51]
- Conturo, T., N. Lori, T. Cull, E. Akbudak, A. Snyder, J. Shimony, R. McKinstry, H. Burton, and M. Raichle: 1999, ‘Tracking neuronal fiber pathways in the living human brain’. *Proceedings of the National Academy of Sciences* **96**, 10422–10427. [208, 210]
- Cremers, D., M. Rousson, and R. Deriche: 2007, ‘Review of Statistical Approaches to Level Set Segmentation: Integrating Color, Texture, Motion and Shape’. *International Journal of Computer Vision* **72**(2), 195–215. [152]
- Dejerine, J.: 1901, *Anatomie des Centres Nerveux*. Paris, Rueff & Cie. [3, 11]
- Dell’Acqua, F., G. Rizzo, P. Scifo, R. Clarke, G. Scotti, and F. Fazio: 2007, ‘A Model-Based Deconvolution Approach to Solve Fiber Crossing in Diffusion-Weighted MR

- Imaging'. *Transactions in Biomedical Engineering* **54**(3), 462–472. [56, 167, 178, 184, 196]
- Deriche, R. and M. Descoteaux: 2007, 'Splitting Tracking Through Crossing Fibers: Multidirectional Q-Ball Tracking'. In: *4th IEEE International Symposium on Biomedical Imaging: From Nano to Macro (ISBI'07)*. Arlington, Virginia, USA, pp. 756–759. [226]
- Dervieux, A. and F. Thomasset: 1979, 'A finite element method for the simulation of Rayleigh-Taylor instability'. *Lecture Notes in Mathematics* **771**, 145–159. [150]
- Dervieux, A. and F. Thomasset: 1981, 'Multifluid Incompressible Flows by a Finite Element Method'. *Lecture Notes in Physics* **11**, 158–163. [150]
- Descoteaux, M., E. Angelino, S. Fitzgibbons, and R. Deriche: 2005a, 'Apparent Diffusion Coefficients from High Angular Resolution Diffusion Images: Estimation and Applications'. Technical Report 5681, INRIA Sophia Antipolis. [66, 84, 100]
- Descoteaux, M., E. Angelino, S. Fitzgibbons, and R. Deriche: 2005b, 'A Linear and Regularized ODF Estimation Algorithm to Recover Multiple Fibers in Q-Ball Imaging'. Technical Report 5768, INRIA Sophia Antipolis. [107, 125, 136]
- Descoteaux, M., E. Angelino, S. Fitzgibbons, and R. Deriche: 2006a, 'Apparent Diffusion Coefficients from High Angular Resolution Diffusion Imaging: Estimation and Applications'. *Magnetic Resonance in Medicine* **56**, 395–410. [75, 100]
- Descoteaux, M., E. Angelino, S. Fitzgibbons, and R. Deriche: 2006b, 'Apparent Diffusion Coefficients from High Angular Resolution Diffusion Imaging: Estimation and Applications'. In: *SPIE Medical Imaging*. San Diego. [100, 259]
- Descoteaux, M., E. Angelino, S. Fitzgibbons, and R. Deriche: 2007a, 'Regularized, Fast, and Robust Analytical Q-Ball Imaging'. *Magnetic Resonance in Medicine* **58**(3), 497–510. [57, 107, 136, 174]
- Descoteaux, M., E. Angelino, S. Fitzgibbons, and R. Deriche: 2006c, 'A Fast and Robust ODF Estimation Algorithm in Q-Ball Imaging'. In: *Third IEEE International Symposium on Biomedical Imaging: from Nano to Macro*. Arlington, Virginia, USA, pp. 81–84. [107, 136]
- Descoteaux, M. and R. Deriche: 2007a, 'Q-Ball Images Segmentation Using Region-Based Statistical Surface Evolution'. Technical Report 6257, INRIA Sophia Antipolis. [164]
- Descoteaux, M. and R. Deriche: 2007b, 'Segmentation of Q-Ball Images Using Statistical Surface Evolution'. In: Springer (ed.): *Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, Vol. LNCS 4792. Brisbane, Australia, pp. 769–776. [164]

- Descoteaux, M. and R. Deriche: 2007c, ‘Sharpening Improves Clinically Feasible Q-Ball Imaging Reconstructions’. In: *Proceedings of the International Society of Magnetic Resonance in Medicine*. Berlin, Germany, p. 906. [197]
- Descoteaux, M., R. Deriche, and A. Anwander: 2007b, ‘Deterministic and Probabilistic Q-Ball Tractography: from Diffusion to Sharp Fiber Distributions’. Technical Report 6273, INRIA Sophia Antipolis. [197]
- Descoteaux, M., R. Deriche, and C. Lenglet: 2007c, ‘Diffusion Tensor Sharpening Improves White Matter Tractography’. In: *SPIE Medical Imaging*. San Diego, California, USA. [197, 208, 212, 224, 226, 236]
- Descoteaux, M., P. Savadjiev, J. Campbell, B. Pike, K. Siddiqi, and R. Deriche: 2007d, ‘Validation and Comparison of Analytical Q-Ball Imaging Methods’. In: *4th IEEE International Symposium on Biomedical Imaging: From Nano to Macro (ISBI'07)*. Arlington, Virginia, USA, pp. 1084–1087. [136]
- Dong, Q., R. Welsh, T. Chenevert, R. Carlos, P. Maly-Sundgren, D. Gomez-Hassan, and S. Mukherji: 2004, ‘Clinical Applications of Diffusion Tensor Imaging’. *Journal of Magnetic Resonance Imaging* **19**, 6–18. [4]
- Douek, P., R. Turner, J. Pekar, N. Patronas, and D. LeBihan: 1991, ‘MR color mapping of myelin fiber orientation’. *Journal of Computer Assisted Tomography* **15**, 923–929. [43]
- Doug, Q., R. C. Welsh, T. L. Chenevert, R. C. Carlos, P. Maly-Sundgren, D. M. Gomez-Hassan, and S. K. Mukherji: 2004, ‘Clinical applications of diffusion tensor imaging’. *Journal of Magnetic Resonance Imaging* **19**, 6–18. [75]
- Einstein, A.: 1956, *Investigations on the Theory of the Brownian Movement*. Dover Pubns. [37]
- Evans, L.: 1998, *Partial Differential Equations*, Vol. 19 of *Graduate Studies in Mathematics*. Proceedings of the American Mathematical Society. [62]
- Feddern, C., J. Weickert, and B. Burgeth: 2003, ‘Level-Set Methods for Tensor-Valued Images’. In: *Proceedings of the Second IEEE Workshop on Geometric and Level Set Methods in Computer Vision*. pp. 65–72. [145, 146]
- Feddern, C., J. Weickert, B. Burgeth, and M. Welk: 2004, ‘Curvature-driven PDE methods for matrix-valued images’. Technical Report 104, Department of Mathematics, Saarland University, Saarbrücken, Germany. [146]
- Ferrers, N. M.: 1969, *An elementary treatise on spherical harmonics and subjects connected with them*. The Cornell University Library Digital Collections. [69]
- Fick, A.: 1855, ‘On liquid diffusion’. *Philos. Mag. J. Sci.* **10**, 31–39. [37]

- Fillard, P., V. Arsigny, X. Pennec, and N. Ayache: 2007a, ‘Clinical DT-MRI Estimation, Smoothing and Fiber Tracking with Log-Euclidean Metrics’. *IEEE Transactions on Medical Imaging* **26**(11), 1472–1482. [45, 245, 254]
- Fillard, P., V. Arsigny, X. Pennec, K. M. Hayashi, P. M. Thompson, and N. Ayache: 2007b, ‘Measuring brain variability by extrapolating sparse tensor fields measured on sulcal lines’. *NeuroImage* **34**, 639–650. [45]
- Fonteijn, H. M. J., F. A. J. Verstraten, and D. G. Norris: 2007, ‘Probabilistic Inference on Q-Ball Imaging Data’. *IEEE Transactions in Medical Imaging* **26**(11), 1515–1524. [167, 245, 254]
- Frank, L.: 2002, ‘Characterization of Anisotropy in High Angular Resolution Diffusion-Weighted MRI’. *Magnetic Resonance in Medicine* **47**(6), 1083–1099. [53, 54, 68, 69, 75, 83, 87]
- Friman, O., G. Farneback, and C.-F. Westin: 2006, ‘A Bayesian Approach for Stochastic White Matter Tractography’. *IEEE Transactions in Medical Imaging* **25**(8). [209]
- Friman, O. and C.-F. Westin: 2005, ‘Uncertainty in Fiber Tractography’. In: *Eighth International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI’05)*. Palm Springs, CA, USA, pp. 107–114. [224, 245, 254]
- FSL: 2006, ‘FMRIB Software Library’. *University of Oxford*. [262]
- Gia, Q. T. L.: 2003, ‘Approximation of linear partial differential equations on spheres’. Ph.D. thesis, Texas A&M University. [103]
- Gray, H.: 1918, *Gray’s Anatomy of the Human Body*. LEA & FEBIGER. [3, 24, 26, 28, 31, 32, 33]
- Guo, W., Q. Zeng, Y. Chen, and Y. Liu: 2006, ‘Using Multiple Tensor Deflection to Reconstruct White Matter Fiber Traces with Branching’. In: *Third IEEE International Symposium on Biomedical Imaging: from Nano to Macro*. Arlington, Virginia, USA, pp. 69–72. [207, 209]
- Hagmann, P., L. Jonasson, T. Deffieux, R. Meuli, J.-P. Thiran, and V. J. Wedeen: 2006a, ‘Fibertract segmentation in position orientation space from high angular resolution diffusion MRI’. *NeuroImage* **32**, 665–675. [145, 147, 167]
- Hagmann, P., L. Jonasson, P. Maeder, J.-P. Thiran, V. J. Wedeen, and R. Meuli: 2006b, ‘Understanding Diffusion MR Imaging Techniques: From Scalar Diffusion-Weighted Imaging to Diffusion Tensor Imaging and Beyond’. *RadioGraphics* **26**, S205–S223. [36, 42, 49, 50, 51, 211]

- Hagmann, P., T. G. Reese, W.-Y. I. Tseng, R. Meuli, J.-P. Thiran, and V. J. Wedeen: 2004, 'Diffusion Spectrum Imaging Tractography in Complex Cerebral White Matter: an Investigation of the Centrum Semiovale'. In: *Proceedings of the International Society of Magnetic Resonance in Medicine*. p. 623. [207, 208]
- Hahn, E.: 1950, 'Spin Echoes'. *Physical Review* **80**, 580–594. [39]
- Hansen, P.: 2001, 'The L-curve and its use in the numerical treatment of inverse problems'. In: *Computational Inverse Problems in Electrocardiology*. pp. 119–142. [86, 115, 175]
- Haroon, H. A. and G. J. Parker: 2007, 'Using the Wild Bootstrap to Quantify Uncertainty in Fibre Orientations from Q-Ball Analysis'. In: *Proceedings of the International Society of Magnetic Resonance in Medicine*. Berlin, Germany, p. 903. [167, 207, 209, 225, 245, 254]
- Healy, D. M., H. Hendriks, and P. T. Kim: 1998, 'Spherical Deconvolution'. *Journal of Multivariate Analysis* **67**, 1–22. [170, 173]
- Hess, C., P. Mukherjee, E. Han, D. Xu, and D. Vigneron: 2006, 'Q-Ball Reconstruction of Multimodal Fiber Orientations Using The Spherical Harmonic Basis'. *Magnetic Resonance in Medicine* **56**, 104–117. [viii, 57, 69, 79, 107, 108, 109, 110, 114, 117, 118, 120, 121, 123, 124, 126, 135, 136, 174, 259]
- Hirsch, J. G., S. M. Schwenk, C. Rossmanith, M. G. Hennerici, and A. Gass: 2003, 'Deviations from the diffusion tensor model as revealed by contour plot visualization using high angular resolution diffusion-weighted imaging (HARDI)'. *Magnetic Resonance Materials in Physics, Biology and Medicine* **16**(2), 93–102. [54]
- Hlawitschka, M. and G. Scheuermann: 2005, 'HOT-Lines: Tracking Lines in Higher Order Tensor Fields'. In: *IEEE Visualization 2005*. Los Alamitos, CA, USA, pp. 27–34. [116]
- Hofer, S. and J. Frahm: 2006, 'Topography of the human corpus callosum revisited- Comprehensive fiber tractography using diffusion tensor magnetic resonance imaging'. *NeuroImage* **32**(3), 989–994. [231, 235]
- Horsfield, M. and D. K. Jones: 2002, 'Applications of diffusion-weighted and diffusion tensor MRI to white matter diseases - a review'. *NMR in Biomedicine* **15**, 570–577. [12, 46, 208]
- Hosey, T., G. Williams, and R. Ansorge: 2005, 'Inference of Multiple Fiber Orientation in High Angular Resolution Diffusion Imaging'. *Magnetic Resonance in Medicine* **54**, 1480–1489. [55]

- Jansons, K. M. and D. C. Alexander: 2003, 'Persistent angular structure: new insights from diffusion magnetic resonance imaging data'. *Inverse Problems* **19**, 1031–1046. [59, 107, 116]
- Jbabdi, S.: 2006, 'Modélisation de la connectivité anatomique cérébrale et simulation de la croissance des gliomes de bas grade: apport de l'IRM de diffusion'. Ph.D. thesis, École Centrale Paris. [24]
- Jbabdi, S., P. Bellec, G. Marrelec, V. Perlbarg, and H. Benali: 2004, 'A level set method for building anatomical connectivity paths between brain areas using DTT'. In: *International Symposium on Biomedical Imaging*. Washinton DC, USA, pp. 1024–1027. [209]
- Jbabdi, S., P. Bellec, R. Toro, J. Daunizeau, M. Pelegrini-Issac, and H. Benali: 2007a, 'Accurate anisotropic fast marching for diffusion-based geodesic tractography'. *International Journal of Biomedical Imaging* **in press**. [209]
- Jbabdi, S., M. Woolrich, J. Andersson, and T. Behrens: 2007b, 'A Bayesian framework for global tractography'. *NeuroImage* **37**, 116–129. [209]
- Jenkinson, M., P. Bannister, M. Brady, and S. Smith: 2002, 'Improved Optimization for the Robust and Accurate Linear Registration and Motion Correction of Brain Images'. *NeuroImage* **17**(2), 825–841. [262]
- Jian, B. and B. C. Vemuri: 2007a, 'Multi-Fiber Reconstruction from Diffusion MRI Using Mixture of Wisharts and Sparse Deconvolution'. In: *International Conference on Information Processing in Medical Imaging (IPMI)*. [56]
- Jian, B. and B. C. Vemuri: 2007b, 'A Unified Computational Framework for Deconvolution to Reconstruct Multiple Fibers from Diffusion Weighted MRI'. *IEEE Transactions on Medical Imaging* **26**(11), 1464–1471. [56, 99, 167, 174, 175, 176, 178, 196]
- Jian, B., B. C. Vemuri, E. Ozarslan, P. R. Carney, and T. H. Mareci: 2007, 'A novel tensor distribution model for the diffusion-weighted MR signal'. *NeuroImage* **37**, 164–176. [56]
- Jonasson, L.: 2006, 'Segmentation of diffusion weighted MRI using the level set framework'. Ph.D. thesis, Ecole Polytechnique federale de Lausanne. [156]
- Jonasson, L., X. Bresson, P. Hagmann, O. Cuisenaire, R. Meuli, and J.-P. Thiran: 2005a, 'White matter fiber tract segmentation in DT-MRI using geometric flows'. *Medical Image Analysis* **9**, 223–236. [145, 146, 147]
- Jonasson, L., X. Bresson, J.-P. Thiran, V. J. Wedeen, and P. Hagmann: 2007, 'Representing Diffusion MRI in 5-D Simplifies Regularization and Segmentation of White

- Matter Tracts'. *IEEE Transactions on Medical Imaging* **26**(11), 1547–1554. [145, 147, 167]
- Jonasson, L., P. Hagmann, X. Bresson, J.-P. Thiran, and V. J. Wedeen: 2005b, 'Representing Diffusion MRI in 5D for Segmentation of White Matter Tracts with a Level Set Method'. In: Springer (ed.): *Information Processing in Medical Imaging*, Vol. Lecture Notes in Computer Science. pp. 311–320. [145]
- Jones, D. K.: 2004, 'The Effect of Gradient Sampling Schemes on Measures Derived From Diffusion Tensor MRI: A Monte Carlo Study'. *Magnetic Resonance in Medicine* **51**, 807–815. [245]
- Jones, D. K. and P. J. Basser: 2004, 'Squashing Peanuts and Smashing Pumpkins: How Noise Distorts Diffusion-Weighted MR Data'. *Magnetic Resonance in Medicine* **52**, 979–993. [245]
- Jones, D. K., M. A. Horsfield, and A. Simmons: 1999, 'Optimal strategies for measuring diffusion in anisotropic systems by magnetic resonance imaging'. *Magnetic Resonance in Medicine* **42**, 515–525. [261, 262]
- Jones, D. K. and C. Pierpaoli: 2005, 'Confidence Mapping in Diffusion Tensor Magnetic Resonance Imaging Tractography Using a Bootstrap Approach'. *Magnetic Resonance in Medicine* **53**, 1143–1149. [209]
- Jones, D. K., S. C. R. Williams, D. Gasston, M. A. Horsfield, A. Simmons, and R. Howard: 2002, 'Isotropic resolution diffusion tensor imaging with whole brain acquisition in a clinically acceptable time'. *Human Brain Mapping* **15**, 216–230. [262]
- Kaden, E., T. R. Knosche, and A. Anwander: 2007, 'Parametric spherical deconvolution: Inferring anatomical connectivity using diffusion MR imaging'. *NeuroImage* **37**, 474–488. [56, 167, 208, 209, 225]
- Khachaturian, M. H., J. J. Wisco, and D. S. Tuch: 2007, 'Boosting the Sampling Efficiency of q-Ball Imaging Using Multiple Wavevector Fusion'. *Magnetic Resonance in Medicine* **57**, 289–296. [48, 58, 167, 245]
- Kindlmann, G., D. B. Ennis, R. T. Whitaker, and C.-F. Westin: 2007a, 'Diffusion Tensor Analysis with Invariant Gradients and Rotation Tangents'. *IEEE Transactions on Medical Imaging* p. in press. [245]
- Kindlmann, G., R. Estepar, M. Niethammer, S. Haker, and C. Westin: 2007b, 'Geodesic-loxodromes for diffusion tensor interpolation and difference measurement'. In: *10th International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI'07)*, Vol. LNCS 4791. pp. 1–9. [245]
- King, M. D., J. Houseman, D. G. Gadian, and A. Connelly: 1997, 'Localized q-space imaging of the mouse brain'. *Magnetic Resonance in Medicine* **38**(6), 930–937. [51]

- King, M. D., J. Houseman, S. A. Roussel, N. V. Bruggen, S. R. Williams, and D. G. Gadian: 1994, 'q-Space imaging of the brain'. *Magnetic Resonance in Medicine* **32**(6), 707–713. [51]
- Koch, M., D. Norris, and M. Hund-Georgiadis: 2002, 'An investigation of functional and anatomical connectivity using magnetic resonance imaging'. *NeuroImage* **16**, 241–250. [208, 209, 212, 213, 216, 224]
- Kouby, V. E., Y. Cointepas, C. Poupon, D. Rivière, N. Golestani, J.-B. Poline, and J.-F. Mangin: 2005, 'MR diffusion-based inference of a fiber bundle model from a population of subjects'. In: *8th International Conference on Medical Image Computing and Computer Assisted Intervention (MICCAI'05)*. Palmsprings, California, USA, pp. 196–204. [246, 254]
- Kreher, B. W., J. F. Schneider, J. Mader, E. Martin, H. J, and K. Il'yasov: 2005, 'Multi-tensor Approach for Analysis and Tracking of Complex Fiber Configurations'. *Magnetic Resonance in Medicine* **54**, 1216–1225. [207, 208]
- Lauterbur, P.: 1973, 'Image formation by induced local interactions: examples employing nuclear magnetic resonance'. *Nature* **242**, 190–191. [39]
- Lazar, M. and A. L. Alexander: 2005, 'Bootstrap white matter tractography (BOOT-TRACT)'. *NeuroImage* **24**, 524–532. [209]
- Lazar, M., D. Weinstein, J. Tsuruda, K. Hasan, K. Arfanakis, M. Meyerand, B. Badie, H. Rowley, V. Haughton, A. Field, and A. Alexander: 2003, 'White Matter Tractography Using Diffusion Tensor Deflection'. In: *Human Brain Mapping*, Vol. 18. pp. 306–321. [208, 224]
- LeBihan, D.: 2003, 'Looking into the functional architecture of the brain with diffusion MRI'. *Nature Reviews Neuroscience* **4**, 469–480. [36]
- LeBihan, D. and E. Breton: 1985, 'Imagerie de diffusion in vivo par résonance magnétique nucléaire'. *C. R. Acad. Sci. Paris* **301 Série II**, 1109–1112. [3, 12, 15, 42]
- LeBihan, D., E. Breton, D. Lallemand, P. Grenier, E. Cabanis, and M. L. Jeantet: 1986, 'MR imaging of intravoxel incoherent motions: application to diffusion and perfusion in neurologic disorders'. *Radiology* **61**, 401–407. [42]
- LeBihan, D., J. Mangin, C. Poupon, C. Clark, S. Pappata, N. Molko, and H. Chabriat: 2001, 'Diffusion Tensor Imaging: Concepts and Applications'. *Journal of Magnetic Resonance Imaging* **13**, 534–546. [4]
- LeBihan, D., C. Poupon, A. Amadon, and F. Lethimonnier: 2006, 'Artifacts and Pitfalls in Diffusion MRI'. *Journal of Magnetic Resonance Imaging* **24**, 478–488. [223]

- Leemans, A.: 2006, 'Modeling and Processing of Diffusion Tensor Magnetic Resonance Images for Improved Analysis of Brain Connectivity'. Ph.D. thesis, Universiteit Antwerpen. [231, 232, 235, 246, 254]
- Leemans, A., J. Sijbers, and P. Parizel: 2005, 'A Graphical Toolbox for Exploratory Diffusion Tensor Imaging and Fiber Tractography'. In: *Section for Magnetic Resonance Technologists (SMRT) - 14th Annual Meeting in Miami Beach*. [231]
- Lenglet, C.: 2006, 'Geometric and Variational Methods for Diffusion Tensor MRI Processing'. Ph.D. thesis, Université de Nice-Sophia Antipolis. [24, 36, 45, 60, 209]
- Lenglet, C., M. Rousson, and R. Deriche: 2006a, 'DTI Segmentation by Statistical Surface Evolution'. *IEEE Transactions in Medical Imaging* **25**(6), 685–700. [145, 146, 147, 149, 151, 152, 153, 154, 155, 156, 157, 158, 159, 211]
- Lenglet, C., M. Rousson, R. Deriche, and O. Faugeras: 2006b, 'Statistics on the Manifold of Multivariate Normal Distributions: Theory and Application to Diffusion Tensor MRI Processing'. *Journal of Mathematical Imaging and Vision* **25**(3), 423–444. [99, 176]
- Lim, K. O. and J. A. Helpert: 2002, 'Neuropsychiatric applications of DTI - a review'. *NMR in Biomedicine* **15**, 587–593. [12, 46]
- Lin, C., V. Wedeen, J. Chen, C. Yao, and W. I. Tseng: 2003, 'Validation of diffusion spectrum magnetic resonance imaging with manganese-enhanced rat optic tracts and ex vivo phantoms'. *NeuroImage* **19**, 482–495. [51, 246, 254]
- Liu, C., R. Bammer, B. Acar, and M. E. Moseley: 2004, 'Characterizing Non-Gaussian Diffusion by Using Generalized Diffusion Tensors'. *Magnetic Resonance in Medicine* **51**, 924–937. [54]
- Maier, S. E., S. Vajapeyam, H. Mamata, C.-F. Westin, F. A. Jolesz, and R. V. Mulkern: 2004, 'Biexponential Diffusion Tensor Analysis of Human Brain Diffusion Data'. *Magnetic Resonance in Medicine* **51**, 321–330. [55]
- Mansfield, P.: 1977, 'Multi-Planar Image Formation using NMR Spin Echoes'. *Journal of Physics C* **10**, 55–58. [39]
- McGraw, T., B. Vemuri, B. Yeziarski, and T. Mareci: 2006a, 'Von Mises-Fisher Mixture Model of the Diffusion ODF'. In: *3rd IEEE International Symposium on Biomedical Imaging (ISBI): Macro to Nano*. [57]
- McGraw, T., B. Vemuri, R. Yeziarski, and T. Mareci: 2006b, 'Segmentation of High Angular Resolution Diffusion MRI Modeled as a Field of von Mises-Fisher Mixtures'. In: *European Conference on Computer Vision (ECCV)*, Vol. 3953. pp. 463–475. [58, 145, 147]

- Merboldt, K., W. Hanicke, and J. Frahm: 1985, 'Self-diffusion NMR Imaging Using Stimulated Echoes'. *J. Magn. Reson.* **64**, 479–486. [3, 12]
- Mori, S., B. Crain, V. Chacko, and P. V. Zijl: 1999, 'Three-Dimensional Tracking of Axonal Projections in the Brain by Magnetic Resonance Imaging'. *Annals of Neurology* **45**(2), 265–269. [208, 210, 226]
- Mori, S. and P. C. M. van Zijl: 2002, 'Fiber tracking: principles and strategies - a technical review'. *NMR in Biomedicine* **15**, 468–480. [12, 25, 46, 207, 208, 210, 211]
- Mori, S., S. Wakana, L. M. Nagae-Poetscher, and P. C. M. van Zijl: 2005, *MRI Atlas of Human White Matter*. Elsevier. [216]
- Moseley, M., Y. Cohen, J. Mintorovitch, J. Kucharczyk, J. Tsuruda, P. Weinstein, and D. Norman: 1990, 'Evidence of Anisotropic Self-Diffusion'. *Radiology* **176**, 439–445. [3, 12, 15, 43]
- Neil, J., J. Miller, P. Mukherjee, and P. S. Huppi: 2002, 'Diffusion tensor imaging of normal and injured developing human brain - a technical review'. *NMR in Biomedicine* **15**, 543–552. [12, 46]
- Neuman, C. H.: 1974, 'Spin echo of spins diffusing in a bounded medium'. *Journal of Chemical Physics* **60**(11), 4508–4511. [56]
- Nieuwenhuys, R., J. Voogd, and C. van Huijzen: 1981, *The Human Central Nervous System. A Synopsis and Atlas*. Springer. [29]
- O'Donnell, L. J.: 2006, 'Cerebral White Matter Analysis Using Diffusion Imaging'. Ph.D. thesis, Massachusetts Institute of Technology. [28, 246, 254]
- O'Donnell, L. J. and C.-F. Westin: 2007, 'Automatic Tractography Segmentation Using a High-Dimensional White Matter Tracts'. *IEEE Transactions on Medical Imaging* **26**(11), 1562–1575. [246, 254]
- Osher, S. and J. Sethian: 1988, 'Fronts Propagating with Curvature-Dependent Speed: Algorithms Based on Hamilton–Jacobi Formulations'. *Journal of Computational Physics* **79**(1), 12–49. [150]
- Osment, P., K. Packer, M. Taylor, J. J. Attard, T. A. Carpenter, L. D. Hall, S. J. Doran, and N. J. Herrod: 1990, 'NMR Imaging of Fluids in Porous Solids'. *Phil. Trans. Roy. Soc.* **333**, 441–452. [3, 12]
- Ozarslan, E. and T. Mareci: 2003a, 'Generalized Diffusion Tensor Imaging and Analytical Relationships Between Diffusion Tensor Imaging and High Angular Resolution Imaging'. *Magnetic Resonance in Medicine* **50**, 955–965. [53, 54, 75, 77, 79, 80, 82, 83, 87, 90, 109]

- Ozarslan, E. and T. H. Mareci: 2003b, 'Anisotropy as a Certainty Measure in terms of Entropy'. In: *Proceedings of the International Society of Magnetic Resonance in Medicine*. p. 249. [84]
- Ozarslan, E., T. Shepherd, B. Vemuri, S. Blackband, and T. Mareci: 2005a, 'Fast Orientation Mapping from HARDI'. In: *MICCAI*, Vol. 3749. pp. 156–163. [69, 107]
- Ozarslan, E., T. Shepherd, B. Vemuri, S. Blackband, and T. Mareci: 2006, 'Resolution of Complex Tissue Microarchitecture Using the Diffusion Orientation Transform (DOT)'. *NeuroImage* **31**(3), 1086–1103. [59, 60]
- Ozarslan, E., B. Vemuri, and T. Mareci: 2005b, 'Generalized Scalar Measures for Diffusion MRI Using Trace, Variance and Entropy'. *Magnetic Resonance in Medicine* **53**(4), 866–876. [77, 84, 85, 90, 91, 93, 260]
- Ozarslan, E., B. C. Vemuri, and T. Mareci: 2004a, 'Fiber Orientation Mapping Using Generalized Diffusion Tensor Imaging'. In: *ISBI*. pp. 1036–1039. [59]
- Ozarslan, E., B. C. Vemuri, and T. Mareci: 2004b, 'Fiber Orientation Mapping Using Generalized Diffusion Tensor Imaging'. In: *ISBI*. pp. 1036–1038. [260]
- Pajevic, S. and C. Pierpaoli: 1999, 'Color schemes to represent the orientation of anisotropic tissues from diffusion tensor data: application to white matter fiber tract mapping in the human brain'. *Magnetic Resonance in Medicine* **42**, 526–540. [46]
- Paragios, N. and R. Deriche: 2002, 'Geodesic active regions: a new paradigm to deal with frame partition problems in computer vision'. *Journal of Visual Communication and Image Representation, Special Issue on Partial Differential Equations in Image Processing, Computer Vision and Computer Graphics* **13**(1/2), 249–268. [146, 147, 150]
- Parker, G. and D. Alexander: 2003, 'Probabilistic Monte Carlo Based Mapping of Cerebral Connections Utilising Whole-Brain Crossing Fibre Information'. In: *IPMI*. pp. 684–695. [208, 209, 224, 245, 254]
- Parker, G. J. M. and D. C. Alexander: 2005, 'Probabilistic Anatomical Connectivity Derived from the Microscopic Persistent Angular Structure of Cerebral Tissue'. *Philosophical Transactions of the Royal Society, Series B* **360**, 893–902. [207, 209, 223]
- Peleda, S., O. Friman, F. Jolesz, and C.-F. Westin: 2006, 'Geometrically constrained two-tensor model for crossing tracts in DWI'. *Magnetic Resonance Imaging* **24**, 1263–1270. [55]

- Perrin, M.: 2006, 'Imagerie de diffusion à haute résolution angulaire : étude du modèle q-ball par couplage simulations - fantôme et applications au suivi de fibres et à la parcellisation du cortex'. Ph.D. thesis, Université Paris-Sud XI. [24, 29, 31, 241, 246, 252, 254]
- Perrin, M., C. Poupon, Y. Cointepas, B. Rieul, N. Golestani, C. Pallier, D. Riviere, A. Constantinesco, D. L. Bihan, and J.-F. Mangin: 2005, 'Fiber tracking in q-ball fields using regularized particle trajectories'. In: *Information Processing in Medical Imaging*. pp. 52–63. [167, 207, 209]
- Pierpaoli, C. and P. J. Basser: 1996, 'Toward a quantitative assessment of diffusion anisotropy'. *Magnetic Resonance in Medicine* **36**, 893–906. [44, 46]
- Pierpaoli, C., P. Jezzard, P. Basser, A. Barnett, and G. D. Chiro: 1996, 'Diffusion Tensor MR Imaging of Human Brain'. *Radiology* **201**, 637–648. [4, 12, 44]
- Poupon, C.: 1999, 'Détection des faisceaux de fibres de la substance blanche pour l'étude de la connectivité anatomique cérébrale.'. Ph.D. thesis, Ecole Nationale Supérieure des Télécommunications. [4, 24, 45, 48, 208]
- Poupon, C., F. Poupon, L. Allirol, and J.-F. Mangin: 2006, 'A database dedicated to anatomo-functional study of human brain connectivity'. In: *Twelfth Annual Meeting of the Organization for Human Brain Mapping (HBM)*. [18, 145, 151, 152, 153, 159, 160, 161, 263]
- Poupon, C., F. Poupon, A. Roche, Y. Cointepas, J. Dubois, and J.-F. Mangin: 2007, 'Real-Time MR Diffusion Tensor and Q-Ball Imaging Using Kalman Filtering'. In: *10th International Conference on Medical Image Computing and Computer Assisted Intervention (MICCAI'07)*. Brisbane, Australia, pp. 27–35. [245, 253]
- Pribam, K. and P. MacLean: 1953, 'Neuronographic Analysis of Medial and Basal Cerebral Cortex'. *J. of Neurophysiology* **16**, 324–340. [3, 11]
- Price, G., M. Cercignani, G. J. Parker, D. R. Altmann, T. R. Barnes, G. J. Barker, E. M. Joyce, and M. A. Ron: 2007, 'Abnormal brain connectivity in first-episode psychosis: A diffusion MRI tractography study of the corpus callosum'. *NeuroImage* **35**(2), 458–466. [232]
- Purcell, E., H. Torrey, and R. Pound: 1946, 'Resonance Absorption by Nuclear Magnetic Moments in a Solid'. *Physical Review* **69**, 37–38. [39]
- Ramirez-Manzanares, A. and M. Rivera: 2006, 'Basis Tensor Decomposition for Restoring Intra-Voxel Structure and Stochastic Walks for Inferring Brain Connectivity in DT-MRI'. *International Journal of Computer Vision* **69**(1), 77 – 92. [207, 209]

- Ramirez-Manzanares, A., M. Rivera, B. Vemuri, P. Carney, and T. Mareci: 2007, 'Diffusion Basis Functions Decomposition for Estimating White Matter Intra-Voxel Fiber Geometry'. *IEEE Transactions on Medical Imaging* p. in press. [55, 56, 167]
- Rao, M., Y. Chen, B. Vemuri, and F. Wang: 2004, 'Cumulative Residual Entropy: A New Measure of Information'. *IEEE Transactions on Information Theory* **50**, 1220–1228. [84]
- Rousson, M.: 2004, 'Cue Integration and Front Evolution in Image Segmentation'. Ph.D. thesis, Universite de Nice - Sophia Antipolis. [151, 152, 156, 163]
- Rousson, M., C. Lenglet, and R. Deriche: 2004, 'Level Set and Region Based Surface Propagation for Diffusion Tensor MRI Segmentation'. In: *Computer Vision Approaches to Medical Image Analysis (CVAMIA) and Mathematical Methods in Biomedical Image Analysis (MMBIA) Workshop*. Prague. [145, 146, 147, 149, 150, 152, 153, 156, 157]
- Sakaie, K. E. and M. J. Lowe: 2007, 'An Objective Method For Regularization of Fiber Orientation Distributions Derived From Diffusion-Weighted MRI'. *NeuroImage* **34**, 169–176. [8, 18, 79, 86, 111, 115, 117, 124, 166, 167, 168, 173, 174, 175, 184, 196, 197]
- Savadjiev, P., J. S. W. Campbell, M. Descoteaux, R. Deriche, G. B. Pike, and K. Siddiqi: 2007, 'Disambiguation of Complex Subvoxel Fibre Configurations in High Angular Resolution Fibre Tractography'. In: *Joint Annual Meeting ISMRM-ESMRMB*. Berlin, Germany, p. 1477. [209, 210, 223, 224, 242, 252]
- Savadjiev, P., J. S. W. Campbell, M. Descoteaux, R. Deriche, G. B. Pike, and K. Siddiqi: 2008, 'Labeling of ambiguous sub-voxel fibre bundle configurations in high angular resolution diffusion MRI'. *Submitted to NeuroImage*. [243]
- Savadjiev, P., J. S. W. Campbell, G. B. Pike, and K. Siddiqi: 2006, '3D curve inference for diffusion MRI regularization and fibre tractography'. *Medical Image Analysis* **10**, 799–813. [130, 131, 224, 243, 252]
- Selden, N., D. Gitelman, N. Salamon-Murayama, T. Parrish, and M. Mesulam: 1998, 'Trajectories of Cholinergic Pathways within the Cerebral Hemispheres of the Human Brain'. *Brain* **121**, 2249–2257. [3, 11]
- Seunarine, K. K. and D. C. Alexander: 2006, 'Linear Persistent Angular Structure MRI and non-linear Spherical Deconvolution for Diffusion MRI'. In: *International Society for Magnetic Resonance in Medicine*. p. 2726. [59]
- Seunarine, K. K., P. A. Cook, K. Embleton, G. J. M. Parker, and D. C. Alexander: 2006, 'A General Framework for Multiple-Fibre PICO Tractography'. In: *Medical Image Understanding and Analysis*. [207, 209]

- Seunarine, K. K., P. A. Cook, M. G. Hall, K. V. Embleton, G. J. M. Parker, and D. C. Alexander: 2007, 'Exploiting peak anisotropy for tracking through complex structures'. In: *Mathematical Methods in Biomedical Image Analysis (MMBIA 2007)*. [207, 209, 210, 224, 242, 245, 254]
- Sijbers, J., A. J. den Dekker, J. V. Audekerke, M. Verhoye, and D. V. Dyck: 1998, 'Estimation of the noise in magnitude MR images'. *Magnetic Resonance Imaging* **16**(1), 87–90. [260]
- Soderman, O. and B. Jonsson: 1995, 'Restricted Diffusion in Cylindrical Geometry'. *Journal Of Magnetic Resonance A* **117**, 94–97. [260]
- Sotak, C.: 2002, 'The Role of Diffusion Tensor Imaging (DTI) in the Evaluation of Ischemic Brain Injury'. *NMR in Biomedicine* **15**, 561–569. [12, 46]
- Stanisz, G., A. Szafer, G. Wright, and R. Henkelman: 1997, 'An Analytical Model of Restricted Diffusion in Bovine Optic Nerve'. *Magnetic Resonance in Medicine* **37**, 103–111. [24, 25]
- Stejskal, E. and J. Tanner: 1965, 'Spin diffusion measurements: spin echoes in the presence of a time-dependent field gradient'. *Journal of Chemical Physics* **42**, 288–292. [15, 39, 40, 41]
- Talairach, J. and P. Tournoux: 1988, *Co-Planar Stereotaxic Atlas of the Human Brain: 3-Dimensional Proportional System : An Approach to Cerebral Imaging*. Thieme Medical Publishers. [233, 262]
- Taylor, D. and M. Bushell: 1985, 'The spatial mapping of translational diffusion coefficients by the NMR imaging technique'. *Physics in Medicine and Biology* **30**(4), 345–349. [3, 12, 42]
- Thiebaut, M. S., M. H. Ezquerro, E. Bardinnet, C. Poupon, H. Benali, J. Yelnick, and S. Lehericy: 2007, 'Three dimensional probabilistic atlas of the human striatal territories based on diffusion imaging'. In: *Proceedings of the International Society of Magnetic Resonance in Medicine*. [244]
- Tournier, J.-D., F. Calamante, and A. Connelly: 2007, 'Robust determination of the fibre orientation distribution in diffusion MRI: Non-negativity constrained super-resolved spherical deconvolution'. *NeuroImage* **35**(4), 1459–1472. [8, 18, 99, 166, 167, 168, 173, 175, 176, 177, 178, 184, 187, 196, 197, 221]
- Tournier, J.-D., F. Calamante, D. Gadian, and A. Connelly: 2003, 'Diffusion-weighted magnetic resonance imaging fibre tracking using a front evolution algorithm'. *NeuroImage* **20**, 276–288. [208, 224]

- Tournier, J.-D., F. Calamante, D. Gadian, and A. Connelly: 2004, 'Direct estimation of the fiber orientation density function from diffusion-weighted MRI data using spherical deconvolution'. *NeuroImage* **23**, 1176–1185. [8, 18, 56, 69, 79, 107, 110, 116, 118, 133, 166, 167, 168, 170, 173, 175, 178, 179, 184, 196, 197, 198, 221, 260]
- Toussaint, N., M. Sermesant, and P. Fillard: 2007a, 'vtkINRIA3D: A VTK Extension for Spatiotemporal Data Synchronization, Visualization and Management'. In: *MICCAI'07 Workshop on Open Source and Open Data*. Brisbane, Australia. [245, 253]
- Toussaint, N., J. Souplet, and P. Fillard: 2007b, 'MedINRIA: Medical Image Navigation and Research Tool by INRIA'. In: *MICCAI'07 Workshop on Interaction in medical image analysis and visualization*. Brisbane, Australia. [245, 253]
- Tschumperlé, D. and R. Deriche: 2001, 'Diffusion Tensor Regularization with Constraints Preservation'. In: *IEEE Computer Society Conference on Computer Vision and Pattern Recognition*. Kauai Marriott, Hawaii, pp. 948–953. [99, 176]
- Tuch, D.: 2002, 'Diffusion MRI of Complex Tissue Structure'. Ph.D. thesis, Harvard University and Massachusetts Institute of Technology. [36, 48, 51, 53, 55, 57, 107, 108, 109, 118, 167, 207, 208, 209, 252, 260]
- Tuch, D.: 2004, 'Q-Ball Imaging'. *Magnetic Resonance in Medicine* **52**(6), 1358–1372. [57, 79, 98, 108, 109, 111, 114, 115, 116, 122, 130, 131, 133, 135, 136, 138, 142, 167, 209, 250, 259, 260]
- Tuch, D., J. Belliveau, T. Reese, and V. Wedeen: 1999, 'High Angular Resolution Imaging of the Human Brain'. In: *Proceedings of the International Society for the Magnetic Resonance in Medicine*. Philadelphia, p. 321. [48, 53, 107]
- Tuch, D., T. Reese, M. Wiegell, N. Makris, J. Belliveau, and V. Wedeen: 2002, 'High Angular Resolution Diffusion Imaging Reveals Intravoxel White Matter Fiber Heterogeneity'. *Magnetic Resonance in Medicine* **48**(4), 577–582. [53, 54, 55]
- Tuch, D., T. Reese, M. Wiegell, and V. Wedeen: 2003, 'Diffusion MRI of Complex Neural Architecture'. *Neuron* **40**, 885–895. [241]
- Tuch, D., J. Wisco, M. Khachaturian, L. Ekstrom, R. Kotter, and W. Vanduffel: 2005, 'Q-ball Imaging of Macaque White Matter Architecture'. *Philosophical Transactions of the Royal Society B* **360**, 869–879. [133]
- Vese, L. and T. Chan: 2002, 'A multiphase level set framework for image segmentation using the Mumford and Shah model'. *The International Journal of Computer Vision* **50**(3), 271–293. [152]
- von dem Hagen, E. and R. Henkelman: 2002, 'Orientational Diffusion Reflects Fiber Structure Within a Voxel'. *Magnetic Resonance in Medicine* **48**, 454–459. [54, 260]

- Wahba, G.: 1990, *Spline Models for Observational Data*. Society for Industrial and Applied Mathematics (SIAM). [86, 115, 175]
- Wang, Z., B. Vemuri, Y. Chen, and T. Mareci: 2004, 'A constrained variational principle for direct estimation and smoothing of the diffusion tensor field from complex DWI'. *IEEE Transactions on Medical Imaging* **23**(8), 930–939. [99]
- Wang, Z. and B. C. Vemuri: 2004, 'Tensor Field Segmentation Using Region Based Active Contour Model'. In: Springer (ed.): *European Conference on Computer Vision (ECCV)*, Vol. Lecture Notes in Computer Science. pp. 304–315. [145, 146]
- Wang, Z. and B. C. Vemuri: 2005, 'DTI Segmentation Using an Information Theoretic Tensor Dissimilarity Measure'. *IEEE Transactions in Medical Imaging* **24**(10), 1267–1277. [145, 146, 149]
- Wassermann, D., M. Descoteaux, and R. Deriche: 2007a, 'Diffusion Maps Clustering for Magnetic Resonance Q-Ball Imaging Segmentation'. *International Journal on Biomedical Imaging, special issue on Recent Advances in Neuroimaging Methodology* p. in press. [247]
- Wassermann, D., M. Descoteaux, and R. Deriche: 2007b, 'Recovering cerebral white matter structures with Spectral Clustering of Diffusion MRI Data'. Technical Report 6351, INRIA Sophia Antipolis - Méditerranée. [247]
- Wedeen, V., T. Reese, D. Tuch, M. Wiegell, J.-G. Dou, R. Weiskoff, and D. Chessler: 2000, 'Mapping fiber orientation spectra in cerebral white matter with Fourier-transform diffusion MRI'. In: *Proceedings of the International Society of Magnetic Resonance in Medicine*. p. 82. [48, 51, 53, 107, 109]
- Wesbey, G., M. Moseley, and R. Ehman: 1984a, 'Translational molecular self-diffusion in magnetic resonance imaging. I. Effects on observed spin-spin relaxation'. *Investigative Radiology* **19**, 484–490. [42, 44]
- Wesbey, G. E., M. E. Moseley, and R. H. Ehman: 1984b, 'Translational molecular self-diffusion in magnetic resonance image. II. Measurement of the self-diffusion coefficient.'. *Investigative Radiology* **19**, 491–498. [42, 44]
- Westin, C., S. Maier, H. Mamata, A. Nabavi, F. Jolesz, and R. Kikinis: 2002, 'Processing and Visualization for Diffusion Tensor MRI'. *Medical Image Analysis* **6**(2), 93–108. [46, 208]
- Wiegell, M. R., D. S. Tuch, H. B. Larsson, and V. J. Wedeen: 2003, 'Automatic segmentation of thalamic nuclei from diffusion tensor magnetic resonance imaging'. *NeuroImage* **19**, 391–401. [145, 146]
- Williams, T., N. Gluhbegovic, and J. Jew: 1997, 'The human brain: dissections of the real brain'. Virtual Hospital. [3, 24, 25, 31, 32, 33]

- Young, M., G. Burns, and J. Scannell: 1995, *The Analysis of Cortical Connectivity*. Landes Bioscience. [3, 11]
- Zhan, W., H. Gu, S. Xu, D. A. Silbersweig, E. Stern, and Y. Yang: 2003, 'Circular Spectrum Mapping for Intravoxel Fiber Structures Based on High Angular Resolution Apparent Diffusion Coefficients'. *Magnetic Resonance in Medicine* **49**, 1077–1088. [54, 69, 75]
- Zhan, W., E. A. Stein, and Y. Yang: 2004, 'Mapping the orientation of intravoxel crossing fibers based on the phase information of diffusion circular spectrum'. *NeuroImage* **23**, 1358–1369. [75]
- Zhan, W., E. A. Stein, and Y. Yang: 2006, 'A rotation-invariant spherical harmonic decomposition method for mapping intravoxel multiple fiber structures'. *NeuroImage* **29**, 1242–1223. [75]
- Zhang, H., P. A. Yushkevich, and J. C. Gee: 2004, 'Registration of Diffusion Tensor Images'. In: *IEEE Conference on Computer Vision and Pattern Recognition*. pp. 842–847. [149]
- Zhao, H.-K., T. Chan, B. Merriman, and S. Osher: 1996, 'A Variational Level Set Approach to Multiphase Motion'. *Journal of Computational Physics* **127**(1), 179–195. [151]
- Zhukov, L., K. Museth, D. Breen, R. Whitakert, and A. H. Barr: 2003, 'Level Set Modeling and Segmentation of DT-MRI Brain Data'. *Journal of Electronic Imaging* **12**, 125–133. [145, 146, 156]