

Cerebral White Matter Analysis Using Diffusion Imaging

by

Lauren Jean O'Donnell

Submitted to the Harvard-MIT Division of Health Sciences and
Technology

in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

May 2006

© Massachusetts Institute of Technology 2006. All rights reserved.

Author
Harvard-MIT Division of Health Sciences and Technology
May 26, 2006

Certified by
W. Eric L. Grimson
Bernard Gordon Professor of Medical Engineering
Professor of Computer Science and Engineering, MIT
Thesis Supervisor

Certified by
Carl-Fredrik Westin
Associate Professor of Radiology, Harvard Medical School
Thesis Supervisor

Accepted by
Martha L. Gray, Ph.D.
Edward Hood Taplin Professor of Medical and Electrical Engineering
Director, Harvard-MIT Division of Health Sciences and Technology

Cerebral White Matter Analysis Using Diffusion Imaging

by

Lauren Jean O'Donnell

Submitted to the Harvard-MIT Division of Health Sciences and Technology
on May 26, 2006, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy

Abstract

In this thesis we address the whole-brain tractography segmentation problem. Diffusion magnetic resonance imaging can be used to create a representation of white matter tracts in the brain via a process called tractography. Whole brain tractography outputs thousands of trajectories that each approximate a white matter fiber pathway. Our method performs automatic organization, or segmentation, of these trajectories into anatomical regions and gives automatic region correspondence across subjects. Our method enables both the automatic group comparison of white matter anatomy and of its regional diffusion properties, and the creation of consistent white matter visualizations across subjects.

We learn a model of common white matter structures by analyzing many registered tractography datasets simultaneously. Each trajectory is represented as a point in a high-dimensional spectral embedding space, and common structures are found by clustering in this space. By annotating the clusters with anatomical labels, we create a model that we call a high-dimensional white matter atlas. Our atlas creation method discovers structures corresponding to expected white matter anatomy, such as the corpus callosum, uncinate fasciculus, cingulum bundles, arcuate fasciculus, etc. We show how to extend the spectral clustering solution, stored in the atlas, using the Nystrom method to perform automatic segmentation of tractography from novel subjects. This automatic tractography segmentation gives an automatic region correspondence across subjects when all subjects are labeled using the atlas. We show the resulting automatic region correspondences, demonstrate that our clustering method is reproducible, and show that the automatically segmented regions can be used for robust measurement of fractional anisotropy.

Thesis Supervisor: W. Eric L. Grimson

Title: Bernard Gordon Professor of Medical Engineering
Professor of Computer Science and Engineering, MIT

Thesis Supervisor: Carl-Fredrik Westin

Title: Associate Professor of Radiology, Harvard Medical School

Acknowledgments

I would like to thank my family, especially my mom for being wonderful and for knowing when not to ask about research. My father said he prefers not to be thanked until I get an Oscar but in case that doesn't work out I take this opportunity to thank him too. Thanks Dad. Thanks to my sister Adrienne and my brother Stephen for sticking with me through 24th grade.

Thanks to my great Boston friends (Eva, Zarmik, Jillian) and the members of our club de libros for encouragement about finishing, and lots of fun dinners. Thanks to my "old" friends David (for making me laugh instead of writing this thesis, and for the revolving door on his NY apartment) and Erin (for arriving randomly in Boston and exploring with me). Thanks to Heather and Steph and the chinatown bus for more fun NY trips, and to Omayra for introducing me to flamenco dancing which helped me make it through several years of grad school.

Thank you Raul for tirelessly maintaining our DTMRI module in Slicer and for taking me to the Grand Canyon (and back, which was the more important journey). Thanks Gordon for many helpful discussions about DTI as well as infinite criticism/encouragement about my research, and for bringing me snacks from 7-11 to speed the writing process. Thanks to Alex Golby for conversations about the brain and neurosurgery, and to Jim Fallon for lessons on neuroanatomy and images for this document. Thanks also to David Banks for rendering tract images with photon simulation. Thanks to Lilla Zollei for letting me use her fabulous congealing image registration, and to Lilla and Corey for living at Winter Street with me. Thanks to Ron Kikinis for not asking me for any Slicer demos for at least 3 years, for career advice, and for recent positive feedback on my results.

Thanks to all of my friends and colleagues who have helped proofread this document: Gordon Kindlmann, Wanmei Ou, Ulas Ziyen, Tomas Izo, and Corey Kemper, and to Louis-Philippe Morency for suggesting the all-important method flowchart. Finally, thanks to the members of my committee: C-F for tons of help with everything over the years, Eric for careful reading of the thesis and advice on MIT requirements

like the research qualifying exam, Polina for detailed proofreading and writing advice, and Steve for keeping my formulas honest.

Finally, I would like to acknowledge the fellowships that I have received during the time I have been at MIT: the Walter A. Rosenblith Fellowship, the NSF Fellowship, the HST Neuroimaging Training Grant, and the HST MEMP Fellowship. The grants which support our research group are the National Alliance for Medical Image Analysis (NIH NIBIB NAMED U54-EB005149), the Morphometry Biomedical Informatics Research Network (NIH NCRR mBIRN U24-RR021382), the Neuroimaging Analysis Center (NIH NCRR NAC P41-RR13218), the NIH NINDS R01 grant on Computational Modeling of Shape Distributions (NIH NINDS R01-NS051826), and the Engineering Research Center for Computer-Integrated Surgical Systems and Technology (NSF JHU ERC CISST).

Contents

1	Introduction	19
2	Background on White Matter Anatomy	27
2.1	Commissural Fibers	30
2.1.1	The Corpus Callosum	30
2.1.2	The Anterior and Posterior Commissures	31
2.2	Association Fibers	31
2.2.1	The Cingulum Bundles and the Fornix	32
2.2.2	The Arcuate Fasciculus And Superior Longitudinal Fasciculi	33
2.2.3	The Uncinate Fasciculus	33
2.2.4	The Inferior Occipitofrontal Fasciculus	34
2.2.5	The Inferior Longitudinal Fasciculus	34
2.3	Projection Fibers	35
2.3.1	Corona Radiata/Internal Capsule/Cerebral Peduncle Region	35
2.3.2	Cerebellar peduncles	36
2.4	Abbreviations	36
2.5	Summary	37
3	Background on Diffusion Tensor Imaging and Analysis	41
3.1	The Diffusion Coefficient	42
3.2	Anisotropic Diffusion in White Matter	43
3.3	Diffusion Tensor MRI	44
3.4	Other Diffusion Models	48

3.5	Scalar Invariants of the Diffusion Tensor	48
3.5.1	Measures of Diffusion Magnitude	48
3.5.2	Anisotropy Measures	49
3.6	DTI Visualization	51
3.7	Tractography	52
3.7.1	Streamline Tractography	56
3.7.2	Surface Evolution, Diffusion Simulation, and Optimization Tractography	58
3.7.3	Probabilistic Tractography	59
3.8	Tractography Grouping Methods	60
3.9	DTI Group Analysis Methods	62
3.9.1	Analysis Region: ROI, VBM, or Tractography	63
3.9.2	Data Analyzed: Scalar, Vector, Tensor, or Anatomy	63
3.9.3	Group Correspondence: Normalization, Tract Correspondence, or Tract Alignment	64
3.10	Classification of Thesis Work in Proposed Categories	65
4	Background on Spectral Clustering	69
4.1	Tractography Clustering Example	69
4.2	Chapter Overview	73
4.3	Spectral Embedding	75
4.3.1	Embedding Vectors and the Input Matrix	75
4.3.2	Embedding Vectors and Eigenvectors of the Input Matrix	78
4.4	Matrices Used for Spectral Embedding	78
4.4.1	The Gram Matrix	79
4.4.2	The Laplacian Matrix of a Graph	79
4.4.3	Markov Matrices	81
4.5	Spectral Embedding Methods	81
4.5.1	Multidimensional Scaling	82
4.5.2	Hall's r-Dimensional Quadratic Placement Algorithm	83

4.5.3	Normalized Cuts	85
4.6	A Note on the Diagonal of A	89
4.7	The Nystrom Method for Eigenvector Estimation	90
4.7.1	Eigenvector Calculation	90
4.7.2	Matrix Normalization	91
4.7.3	Performance	92
4.8	Clustering Example, Revisited	92
4.9	Summary	93
5	White Matter Fiber Tract Clustering	95
5.1	Related Work	98
5.2	Methods	99
5.2.1	Step 1: Tractography for Clustering	99
5.2.2	Step 2: Fiber Trajectory Distance and Matching Across Hemispheres	103
5.2.3	Step 3: Conversion to Affinity	107
5.2.4	Step 4: Clustering of Tractography: Single and Multiple Subjects	109
5.2.5	Step 5: Cluster Color Assignment	110
5.2.6	Step 6: Expert Labeling	111
5.3	Experimental Results	112
5.3.1	Anatomical Matching Across Hemispheres	112
5.3.2	Population Clustering	116
5.4	Discussion of Fiber Tract Clustering Method	120
5.4.1	Choice of Clustering Algorithm	120
5.4.2	Tractography as Input	121
5.4.3	Tractography for ROI Definition	122
5.4.4	Motivation for Atlas Generation	122
5.5	Summary	123
6	A White Matter Fiber Tract Atlas	125
6.1	Related Work	127

6.2	Methods	129
6.2.1	Atlas Generation	129
6.2.2	Atlas-Based Automatic Segmentation: Extending Spectral Clustering to New Data	131
6.2.3	Quantitative Measurement of Scalar Invariants	134
6.2.4	Creation of Voxel ROIs	135
6.3	Experimental Results: Atlas 1	135
6.3.1	Methods	135
6.3.2	Segmentation Results	136
6.3.3	Discussion	136
6.4	Experimental Results: Atlas 2	138
6.4.1	Methods	138
6.4.2	Segmentation Results	139
6.4.3	Automatic Corpus Callosum Subdivision	146
6.4.4	Measurement of Scalar Invariants in the Population	146
6.4.5	Discussion	147
6.5	Summary	149
7	Clinical Application: Neurosurgery	151
7.1	Introduction	151
7.1.1	Tumors, Fibers, and Eloquent Regions	151
7.1.2	DTI In Tumors	152
7.1.3	The Neurosurgical DTI Visualization Problem	153
7.2	Experimental Results	156
7.2.1	Methods	156
7.2.2	Case 1	157
7.2.3	Case 2	163
7.3	Summary	165
8	Reproducibility and Parameter Settings	167
8.1	Distance/Affinity Calculation	167

8.2	Spectral Embedding, Clustering and Atlas Generation	168
8.2.1	Reproducibility of Embedding Subspace	169
8.2.2	Number of Clusters	174
8.2.3	Segmentation Results From Repeat Scans	174
8.2.4	Consistency of Atlas Clustering and Atlas Labeling	175
8.3	Stability of Measured Scalar Invariants	176
8.3.1	Comparison of FA Values From Repeat Scans	176
8.4	Summary	176
9	Conclusion	179
A	Data Acquisition	181

List of Figures

1-1	Example fiber tracts from Gray's Anatomy	19
1-2	Example whole brain tractography	20
1-3	Overview of our method	21
1-4	White matter atlas creation and use for automatic segmentation	22
1-5	Automatic tractography segmentation result for three subjects	24
1-6	Automatic tractography segmentation result for deeper white matter structures	24
2-1	Illustration of a neuron	28
2-2	Lobes of the cerebral hemispheres	29
2-3	Commissural tracts	31
2-4	Illustration of selected association pathways	32
2-5	More association pathways	33
2-6	Axons enter and leave some fiber tracts along their course	34
2-7	Fornix	35
2-8	Projection fibers	36
2-9	Corona radiata/internal capsule/cerebral peduncle regions	37
2-10	Cerebellar peduncles	38
3-1	Illustration of anisotropic diffusion.	43
3-2	Six diffusion-weighted images	45
3-3	Illustration of the tensor eigensystem in a region of anisotropic diffusion	46
3-4	Voxel-based DTI visualization in 3D Slicer	54
3-5	Three-dimensional DTI visualization in 3D Slicer	55

3-6	The major eigenvector does not point in the direction of a fiber tract in the case of crossing fibers.	57
3-7	Isocontours of tensor-based distance map	60
4-1	Tractography clustering example	69
4-2	Weighted graph example	70
4-3	Example affinity matrix and embedding vectors	71
4-4	Example embedding vectors and affinity matrix, after clustering . . .	73
4-5	Example embedding vectors and trajectories, after clustering	74
4-6	Block diagonal affinity matrix example	76
4-7	MDS example	82
4-8	Laplacian embedding example	85
4-9	Normalized cuts embedding example	86
4-10	Diagram of the parts of the affinity matrix when using the Nystrom method.	91
4-11	Clustering example, revisited	93
5-1	Whole-brain tractography, spectral embedding, and example clusters	96
5-2	C_L thresholds can be used to identify white matter regions	100
5-3	One slice through the ROI used to seed tractography	102
5-4	Stopping tractography with C_L is comparable to using FA and a curvature constraint	103
5-5	Comparison of FA and C_L thresholds	104
5-6	Illustration of distance measure between two fibers	105
5-7	Soft color correspondences	107
5-8	Hard cluster correspondence	108
5-9	Diagram of the parts of the affinity (tract similarity) matrix for 5 subjects	110
5-10	Result of clustering with bilateral affinity measure	113
5-11	Selected clusters demonstrate anatomical matching across hemispheres	115
5-12	Anterior, right, and superior views of the five brains in the population	117
5-13	Anatomical correspondences: selected clusters, displayed in all 5 brains	118

6-1	Overview of our method	125
6-2	Atlas generation example	126
6-3	Diagram of the parts of the affinity matrix for atlas construction . . .	130
6-4	Diagram of the parts of the affinity matrix for extension of spectral embedding to new data	132
6-5	Atlas 1	137
6-6	Result of automatic segmentation of all subjects	140
6-7	Result of automatic segmentation of all subjects	141
6-8	Result of automatic segmentation of all subjects	142
6-9	Result of automatic segmentation of all subjects	143
6-10	Result of automatic segmentation of all subjects	144
6-11	Cingulum clusters in test subjects	145
6-12	Automatic subdivision of corpus callosum	146
6-13	Automatic subdivision of corpus callosum, converted to voxels	146
6-14	Mean FA in three regions for 15 subjects	147
6-15	Scalar invariants in all regions for 15 subjects	148
7-1	Slice-based DTI visualization of patient 1	158
7-2	Raw output of whole-brain clustering	159
7-3	Three-dimensional DTI visualization of patient 1	160
7-4	Spatial relationship of arcuate fasciculus cluster and tumor	161
7-5	Slice-based DTI visualization of patient 2.	163
7-6	Three-dimensional DTI visualization of patient 2	164
8-1	Effect of varying σ on affinity matrix	169
8-2	Effect of varying σ on embedding	170
8-3	Single-subject embedding reproducibility	172
8-4	Atlas embedding reproducibility	173
8-5	Fraction of inconsistently clustering pairs	175
8-6	Maximum cluster centroid distance sum versus number of clusters . .	176
8-7	Result of repeat automatic segmentation of novel subject	177

8-8	Result of repeat automatic segmentation of novel subject	177
8-9	FA measurements from three repeat scans of two subjects	178

List of Tables

2.1	Lobes of the brain and their main functions	29
2.2	Anatomical abbreviations	39
3.1	The scale of DTI and the brain: neuron sizes and quantities, and water diffusion times and distances.	44
3.2	Tractography methods	53
3.3	Categorization of DTI group analysis studies by method	66
5.1	Steps in our clustering method	99
5.2	Thresholds used in tractography	101
6.1	Contents of the High-Dimensional Atlas	130
7.1	Common Tumor Types in Adults (percentages from [1])	153
7.2	Neurosurgical planning questions potentially answerable by DTI and tractography.	154
7.3	Proposed methods to address the questions.	154
8.1	Parameters of the method	168

Chapter 1

Introduction

The goal of the methods described in this thesis is to automatically find and anatomically name regions of the white matter of the brain using diffusion MRI. The white matter contains pathways known as fiber tracts that connect functional areas of the brain. Example fiber tracts are shown in Figure 1-1. Previously it was not possible to automatically create patient-specific white matter models similar to these anatomical atlas diagrams of fiber tracts. However, methods are now available to estimate patient-specific white matter fiber trajectories using diffusion MRI. In this thesis we present a method for segmentation of the trajectories estimated from diffusion MRI by automatically grouping them into anatomical regions and labeling them with anatomical names.

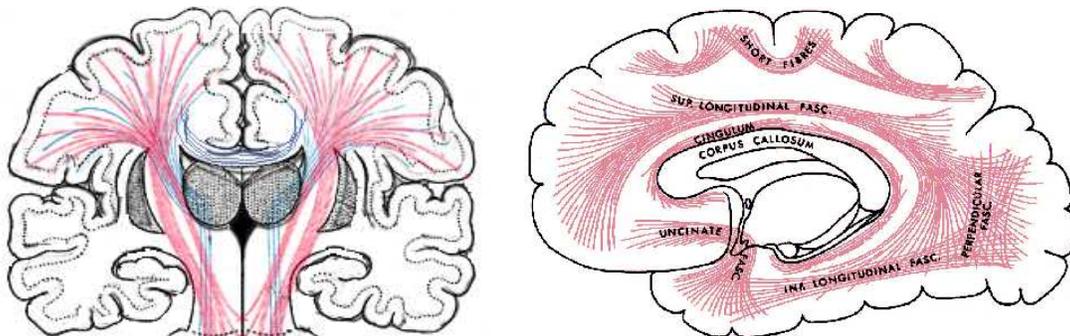


Figure 1-1: Example fiber tracts viewed from the front (coronal) and from the left (sagittal). In this thesis we aim to create patient-specific white matter models similar to these anatomical atlas diagrams which were modified from Gray's Anatomy [46].

The white matter fiber tracts are actually large bundles of axons that interconnect the gray matter processing areas both within and across hemispheres. (The brain consists of two main parts: the white matter, containing axons that transmit information; and the gray matter, containing cell bodies of neurons that integrate information.) Diffusion MRI allows estimation of the trajectories of the white matter fiber tracts via a process called tractography (Figure 1-2), but there has been relatively little work in automatic localization of individual anatomical fiber tracts. There are three

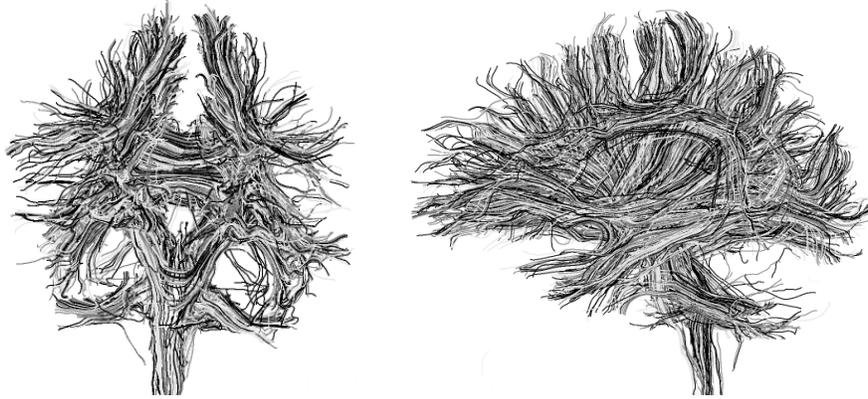


Figure 1-2: Example whole brain tractography viewed from the front and from the side. This is the input to our method. For visual differentiation of individual estimated trajectories, the trajectories have been randomly colored using shades of gray.

main goals which should be satisfied by an automatic tractography segmentation algorithm: automatic grouping of like trajectories into regions, region correspondence across subjects, and anatomical labeling of regions. Our method is the first to achieve these three goals. The method presented in this thesis learns a white matter fiber tract model (a tractography “atlas”) that we then use to segment and label the fiber tracts of a new brain. Our ability to perform automatic, patient-specific definition of the white matter fiber tracts has applications in neuroanatomical visualization, neurosurgical planning, and neuroscientific studies of white matter integrity, structure, and variability.

Because we have access to input tractography data from multiple subjects, we have designed our method to learn common white matter structures that are present in the

data. The structures are found by simultaneous clustering¹ of registered (spatially aligned) tractography from a group of subjects. Each structure discovered by the method is expertly labeled with an anatomical name. We save information about the structures and their labels in a new type of anatomical atlas (a “high-dimensional atlas”) which we describe in this thesis. The information stored in the atlas allows automatic segmentation and anatomical labeling of tractography from new subjects. A flowchart describing our atlas generation and automatic segmentation methods is shown in Figure 1-3. Figure 1-4 displays an example atlas and segmentation result.

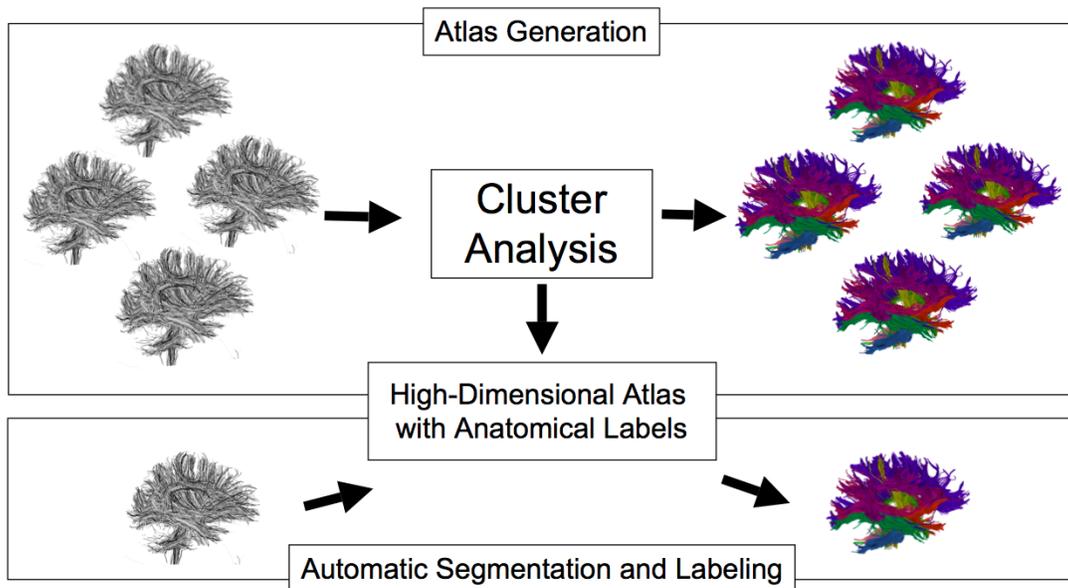


Figure 1-3: Overview of our method. In the first part of the method, an atlas is generated. Clustering of tractography from multiple subjects is used to discover common white matter structures that are anatomically labeled and stored in a high-dimensional atlas. In the second part of the method, tractography from a new subject is segmented and anatomically labeled using the atlas.

Our method can be applied to tractography from any type of diffusion MRI, however the tractography data analyzed in this thesis is derived from diffusion *tensor* MRI or DTI. Diffusion MRI measurements describe a 3D water diffusion profile at each voxel² and in voxels with one fiber tract, the direction of fastest diffusion is

¹Clustering is a general name for techniques which automatically group like items, in this case similar trajectories.

²Analogous to a pixel, or picture element, a voxel is a volume element, usually a region about

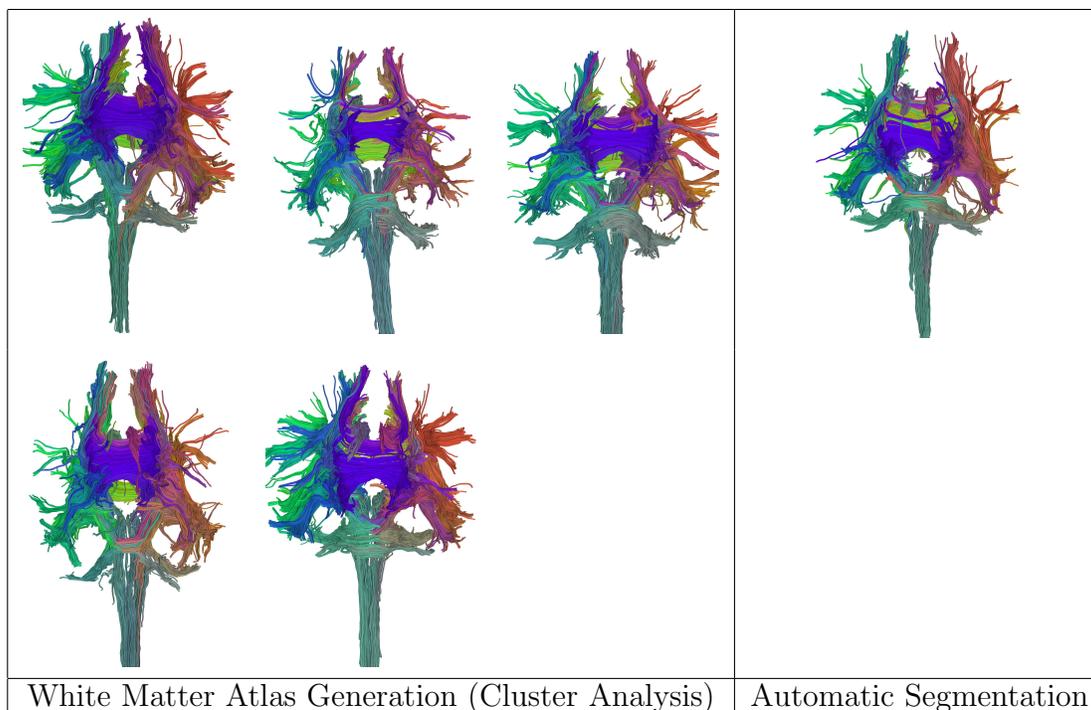


Figure 1-4: White matter atlas creation and use for automatic segmentation. Using the five brains on the left, we automatically found common regions in the white matter: each is shown with a unique color. The information we learned about the regions was stored in our atlas. We then applied the information in the atlas to label a new brain (right) using the same regions and colors. These brains are viewed from the front.

associated with the fiber tract direction. Often a tensor model is chosen to represent the diffusion as Gaussian [3]. This means that we assume that the probability of diffusion of water molecules can be represented by concentric ellipsoids, where each ellipsoid has equal diffusion probability everywhere on its surface. The shape of the ellipsoids is significant: a cigar-like ellipsoid is found in strongly oriented white matter regions where diffusion is faster along the tracts, and a spherical ellipsoid is found in the cerebrospinal fluid where diffusion is the same in all directions. Mathematically the ellipsoid is represented as a 3×3 symmetric matrix, the diffusion tensor, which is proportional to the covariance matrix of the Gaussian distribution.

Compared to standard MRI data that consists of one number per voxel, diffusion tensor MRI is more complicated to analyze, in part because each voxel has a diffu-

2.5 by 2.5 by 2.5 mm for diffusion MRI.

sion tensor measurement containing six unique numbers. The tractography analysis method aims to trace the fiber pathways by associating the major eigenvector of the tensor (the direction of fastest diffusion) with the orientation of a fiber tract, then stepping repeatedly in this direction to estimate the trajectory of the tract. This outputs one trajectory, or path, from each start point. The estimated fiber trajectories produced by tractography (see Figure 1-2) are the input to our method. Though the method of tractography has been studied since the late 1990’s, before now it has not been possible to achieve an automatic higher-level organization of the trajectories into white matter structures with anatomical significance. Interactive “virtual dissection” methods exist for selecting trajectories based on anatomical region [24], however they require sophisticated three-dimensional neuroanatomical knowledge. It is extremely time-consuming to interactively organize all of the possible trajectories in a whole brain (well over 10,000 at current image resolutions) though this type of manual organization has been used to create a white matter atlas [80].

The methods presented in this thesis automatically learn a model of the anatomical organization of white matter pathways from a population of subjects, using an unsupervised spectral clustering method. By simultaneously clustering tractography from multiple subjects, our method is able to learn a representation of regions corresponding to expected white matter fiber tract anatomy. In addition we propose a method for bilateral clustering of anatomical structures, where corresponding tractography is automatically matched across hemispheres in the clustering process.

The anatomical organization learned by the method is represented in a high-dimensional embedding space which is created during clustering. We then introduce expert anatomical knowledge to complement the lower-level cluster information, thereby forming what we refer to as a high-dimensional white matter atlas. We demonstrate the application of the atlas to automatically segment and anatomically label a new subject by embedding its tractography in the high-dimensional embedding space. Our atlas differs from traditional digital (voxel-based) atlases because it represents long-range connections from tractography rather than local voxel-scale information.

Example results of our method are presented in Figures 1-4, 1-5, and 1-6. Figure 1-4 provides an example of the method where a five-subject tractography dataset is used for atlas creation, and a sixth subject is labeled using the atlas. In Figure 1-5 an automatic atlas-based tractography segmentation is demonstrated for three different subjects. In Figure 1-6, selected deeper white matter structures are displayed in an automatically segmented subject.

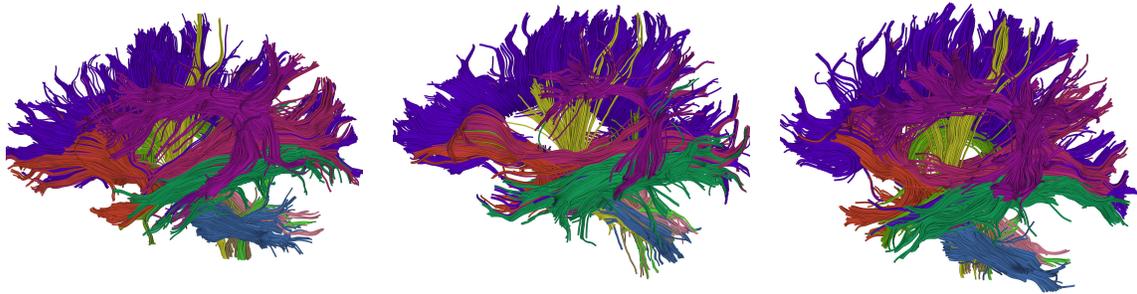


Figure 1-5: Automatic tractography segmentation result for three different subjects (view from left). Colors indicate regions as follows: navy blue, corpus callosum; yellow, corticospinal fibers; purple, arcuate fasciculus/SLF region; orange, uncinate fasciculus; green, inferior longitudinal fasciculus; sky blue, middle cerebellar peduncle; light pink, superior cerebellar peduncle; hot pink, occipitofrontal fasciculus.

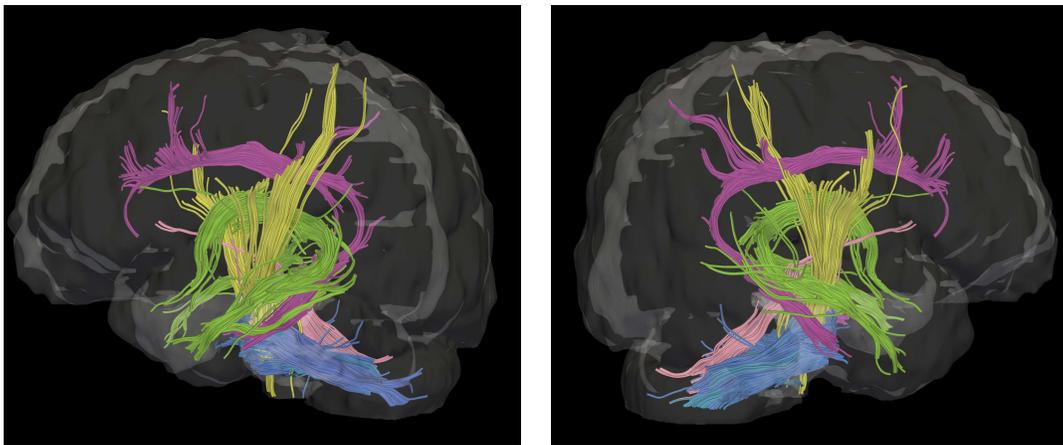


Figure 1-6: Automatic tractography segmentation result for deeper white matter structures (viewed from both left and right sides). Colors indicate regions as follows: yellow, corticospinal fibers; sky blue, middle cerebellar peduncle; light pink, superior cerebellar peduncle; purple, cingulum bundles; lime green, fornix. The tract matching across hemispheres was performed automatically by our method.

To our knowledge, our work is the first automatic method for whole brain tractography segmentation, and the first method to employ multiple-subject clustering to learn an anatomical model. We produce a quantitative description of white matter architecture in the form of a cluster, and a quantitative model of white matter architecture in the group. Our work enables statistical white matter analysis using tractography. We find corresponding white matter regions across subjects and across hemispheres, allowing neuroscientific hypotheses to be tested regarding group differences and also questions of symmetry.

This thesis begins with background chapters on neuroanatomy (Chapter 2), diffusion imaging (Chapter 3), and spectral clustering (Chapter 4). The contributions of the thesis are then presented in the following chapters. Chapter 5 describes the tractography clustering method. Chapter 6 demonstrates the creation of the tractography atlas and its use for automatic segmentation, as well as quantitative measurement of diffusion values in tracts. Chapter 7 then presents a preliminary neurosurgical application of our method, and Chapter 8 discusses parameter settings and reproducibility. Finally, we conclude the thesis in Chapter 9.

Chapter 2

Background on White Matter

Anatomy

We begin with an overview of neuroanatomy, especially white matter fiber tract anatomy, to enable the reader to understand the white matter segmentation problem and the results presented in later chapters. We describe the major white matter fiber tracts which are seen using DTI tractography. The information in this chapter is primarily directed towards readers with a computer science background, however we also include references for diffusion imaging neuroscience studies that investigate some of the white matter tracts. For more information we recommend [84] for neuroanatomy and function, and [50, 80, 116] for DTI-specific neuroanatomy.

At the cellular level, neural information processing is handled by neurons, the cells which transmit electrochemical signals in the brain. A diagram of an example neuron is shown in Figure 2-1. Two main parts of a neuron have important functions: the cell body integrates information from other neurons, while the axon transmits information. When a neuron receives input from other neurons, it may depolarize (fire), producing an electrical action potential which travels away from the neuronal cell body, down the axon, and finally arrives at synapses with other neurons. Neurotransmitter is released at the synapses, and if there is sufficient quantity of neurotransmitter in a synaptic cleft, the postsynaptic neuron may also fire. The conduction velocity of the action potential is increased by a insulating process called

myelination. Supporting glial cells wrap layers of myelin around the axon, periodically leaving open small regions called nodes of Ranvier which allow charged molecules to cross the cell membrane as the action potential travels.

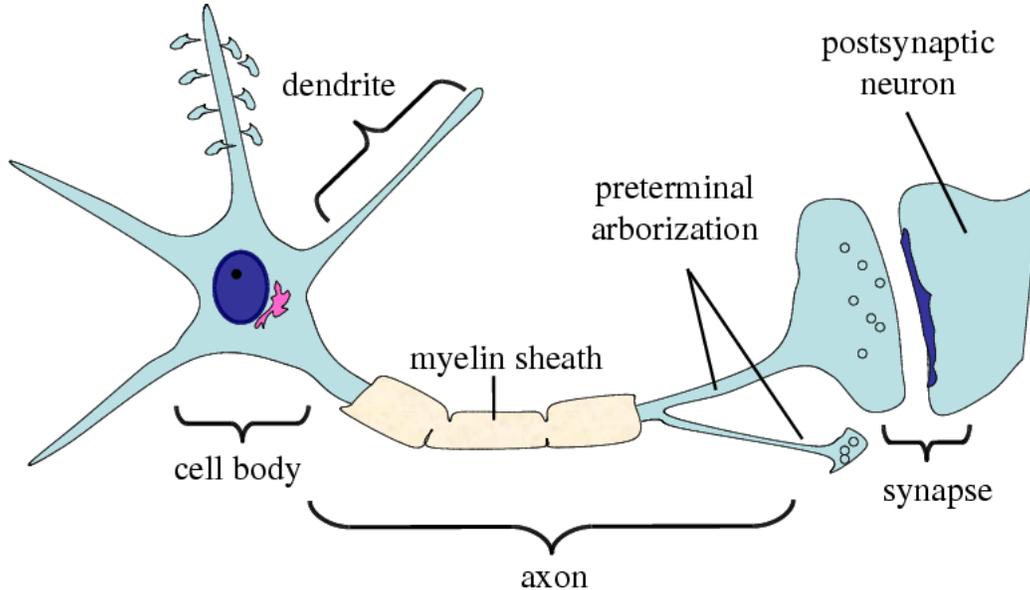


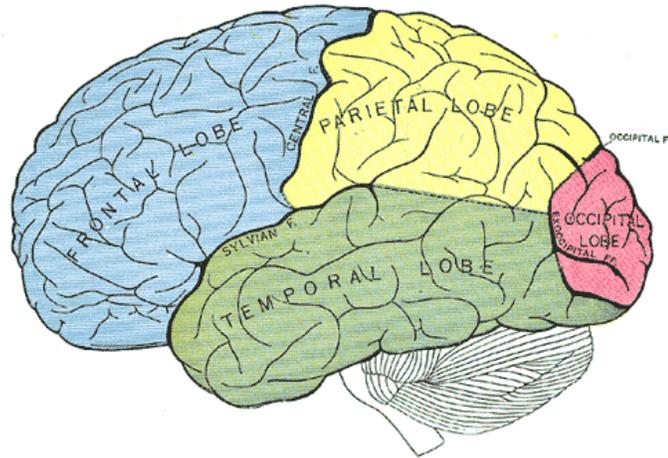
Figure 2-1: Illustration of a neuron. Courtesy of James Fallon, Professor of Anatomy and Neurobiology at the University of California, Irvine.

Myelin is white in color, and the tissue containing the cell bodies is gray in color, giving a natural division of the brain into two main parts: the white matter and the gray matter. The gray matter primarily contains the cell bodies of neurons and their dendrites, while the white matter contains their axons (though there are small interneurons that exist entirely in the gray matter also). The gray matter is located around the outside of the brain in the highly convoluted cortex (with “ridges” called gyri and “valleys” called sulci) and in internal brain structures such as the basal ganglia and thalamus. The white matter contains the axonal pathways that interconnect the brain: bundles of axons travel together and are called white matter fiber tracts. In addition to the white and gray matter regions, there are two important fluid systems in the brain: cerebrospinal fluid fills the ventricles and spaces (cisternae) around the brain, and a complex vasculature system provides blood flow.

The brain has two hemispheres, the right and left, whose functions are complimen-

tary. For example each performs visual processing of one half of the visual field, and the left hemisphere produces language while the right handles the emotional content or prosody of speech. The brain is further divided into five lobes in each hemisphere, the frontal, parietal, temporal, occipital, and limbic lobes (see Figure 2-2). Important functions or regions of each lobe are listed in Table 2.1.

S



A

Figure 2-2: Lobes of the cerebral hemispheres. Frontal is blue, parietal yellow, temporal green, and occipital red. The limbic lobe is internal and can't be seen (it forms a "C" shape that wraps around above and below the corpus callosum.) Image from Gray's Anatomy [46].

Lobe	Functions/Functional Regions
frontal	primary motor, premotor and supplementary motor, Broca's area (speech production), and prefrontal cortex (personality, insight, and foresight)
parietal	primary somatosensory cortex, language comprehension, spatial orientation and perception
temporal	primary auditory cortex, Wernicke's area (language comprehension), visual processing, memory
occipital	primary visual cortex, visual association cortex
limbic	emotional responses, drive-related behavior, memory

Table 2.1: Lobes of the brain and their main functions, from [84].

Processing streams in the brain's functional regions are interconnected by white matter fiber tracts. The white matter contains three types of fiber tract: commissural,

association, and projection. A commissure is a crossing site for fibers which connect similar areas [84], so commissural tracts connect related regions of the two cerebral hemispheres, coordinating their activity. On the other hand, association fibers connect regions in the same hemisphere. Association fibers come in various sizes: the smallest fibers are completely within the cortex, the medium ones are called u-fibers or arcuate fibers and connect one gyrus to the next, and the longest association bundles connect different lobes [84]. Finally, projection fibers connect the cortex and subcortical structures such as the thalamus, basal ganglia, and spinal cord. The connections to and from the cerebellum are also called projection tracts [46]. Fiber tracts may be named according to start and end points (e.g. mammillothalamic tract) or may have different names along their course (corona radiata, internal capsule, and cerebral peduncle). Bundles named for a region (e.g. cingulum) may have axons that enter and leave, similar to entrances and exits on a highway. The preceding details complicate the naming of fiber regions which can be calculated from DTI.

The rest of this chapter describes the major white matter fiber tracts, especially those that are relevant to this thesis, i.e. the ones that can be found with DTI tractography and clustering methods. We begin with commissural tracts, followed by association tracts and finally projection tracts. Images of tract location are provided along with general descriptions of tract connections. To clarify orientation, each image is annotated with the letter “A” for anterior and the letter “S” for superior (or the letter “L” for left if it is a superior view).

2.1 Commissural Fibers

2.1.1 The Corpus Callosum

The largest fiber bundle in the human brain, the corpus callosum (Figure 2-3) is made up of more than 300 million axons [84]. This commissural fiber bundle mainly connects cortical areas to matching contralateral cortical areas, though some fibers connect related areas (for example primary visual cortex connects to contralateral

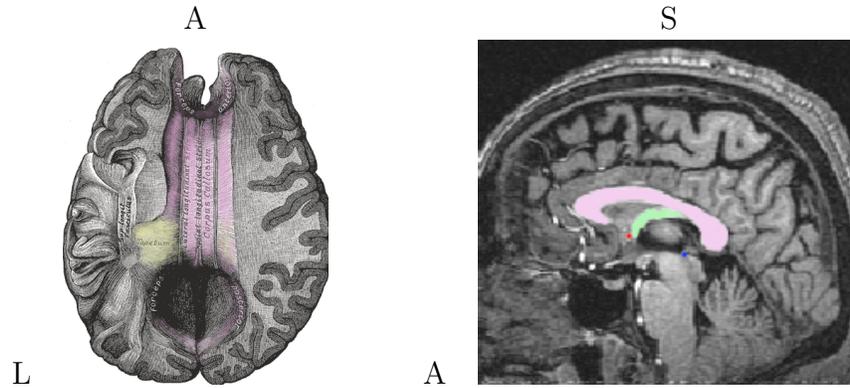


Figure 2-3: Commissural tracts. Left: corpus callosum (superior view) adapted from Gray’s Anatomy [46]. Color was added to show corpus (pink) and tapetum of corpus (yellow). Right: Anterior (red dot) and posterior (blue dot) commissures. Corpus callosum (pink) and fornix (not a commissural tract) in green.

visual association cortex) [84]. The anterior part is called the genu, the posterior the splenium, and the part that connects temporal regions is called the tapetum [46]. The corpus callosum has been studied using DTI, and changes in tract integrity (changes in diffusion anisotropy) have been found for example in Alzheimer’s [99], schizophrenia [66], and alcohol dependence [64].

2.1.2 The Anterior and Posterior Commissures

The anterior commissure (Figure 2-3) connects the temporal cortex and also olfactory regions [84]. It can be seen with DTI (though not satisfactorily traced [133]). The posterior commissure (Figure 2-3) is a smaller structure and lies above the superior colliculi.

2.2 Association Fibers

Major association pathways are shown in Figures 2-4 and 2-5. The association fibers primarily have anterior-posterior trajectories [116]. The structure of association bundles is like superhighways with entrances and exits, rather than isolated trajectories. According to Nolte ([84] page 533), “None of these association bundles should be

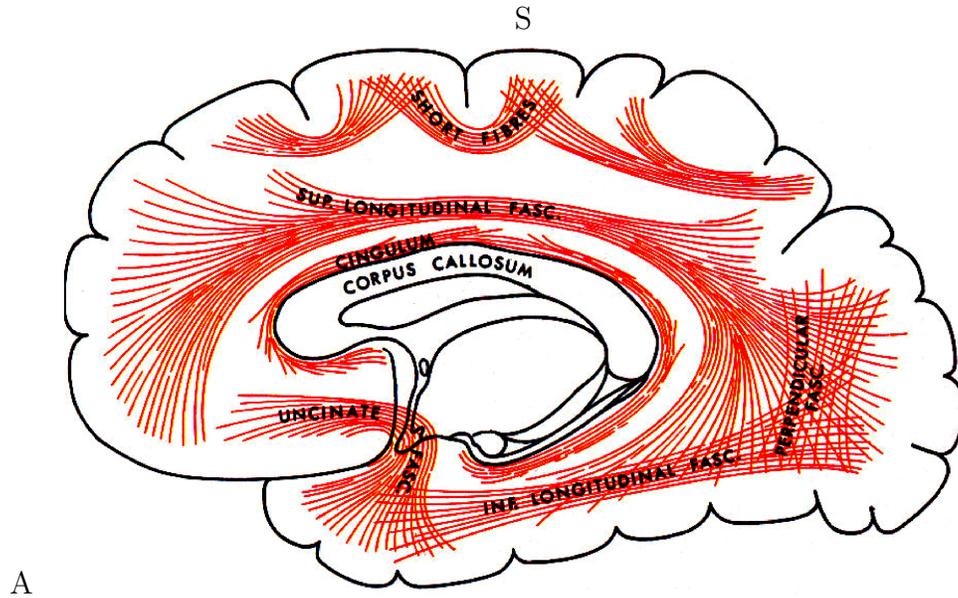


Figure 2-4: Illustration of selected association pathways. From Gray's Anatomy [46].

thought of as a discrete, point-to-point pathway from one place to another; rather fibers enter and leave each all along its course.” An example of axons entering and leaving a bundle is shown in the cingulum diagram, Figure 2-6. In this section we give an overview of the main association bundles which are visible on DTI and can be traced using tractography.

2.2.1 The Cingulum Bundles and the Fornix

The cingulum (Figures 2-4 and 2-6) and fornix (Figure 2-3 and 2-7) are prominent C-shaped fiber tracts in the limbic lobe. The limbic system handles functions related to memory and emotion and its contents include the cingulate and parahippocampal gyri, the hippocampus, and the amygdala [84]. The cingulum bundle connects the cingulate gyrus to other regions in the limbic lobe, especially the hippocampus, which is involved in memory formation [84]. The fornix is an output tract from the hippocampus, connecting primarily to the mammillary bodies [84]. The fornix and cingulum are involved in the classic Papez circuit of the limbic system [84]. DTI has been used to study these limbic system tracts, and reduced white matter integrity has

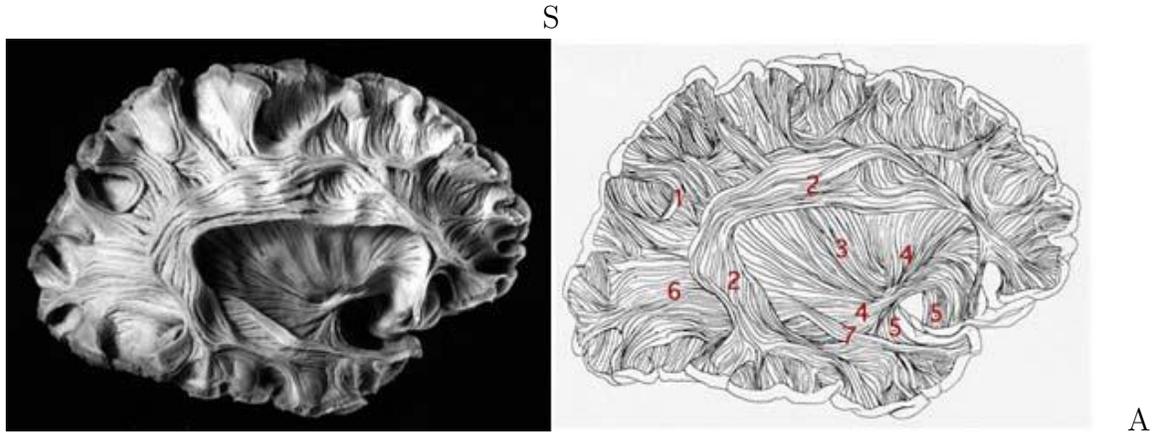


Figure 2-5: More association pathways. 1. Short arcuate bundles (u-fibers). 2. Superior longitudinal fasciculus (part is known as the arcuate fasciculus). 3. External capsule. 4. Inferior occipitofrontal fasciculus. 5. Uncinate fasciculus. 6. Sagittal stratum. 7. Inferior longitudinal fasciculus. From Virtual Hospital [128].

been found in schizophrenia [65] and temporal lobe epilepsy [28]. Using tractography, the limbic system's diffusion properties have been characterized in healthy subjects [29].

2.2.2 The Arcuate Fasciculus And Superior Longitudinal Fasciculi

The superior longitudinal fasciculus (Figure 2-5) is also known as the arcuate fasciculus. It runs from the frontal lobe, above the insula, to the parietal, occipital, and temporal lobes [84]. Classically it connects Broca's and Wernicke's language areas (see Table 2.1). DTI studies have reported reduced tract integrity in the region of the arcuate fasciculus in schizophrenia [66] and asymmetry of this tract across hemispheres in normal subjects [67] (thought to be due to the localization of language production to one hemisphere, usually the left).

2.2.3 The Uncinate Fasciculus

The uncinate fasciculus (Figures 2-4 and 2-5) connects anterior temporal and inferior frontal lobes. Changes in diffusion anisotropy in the region of the uncinate fasciculus

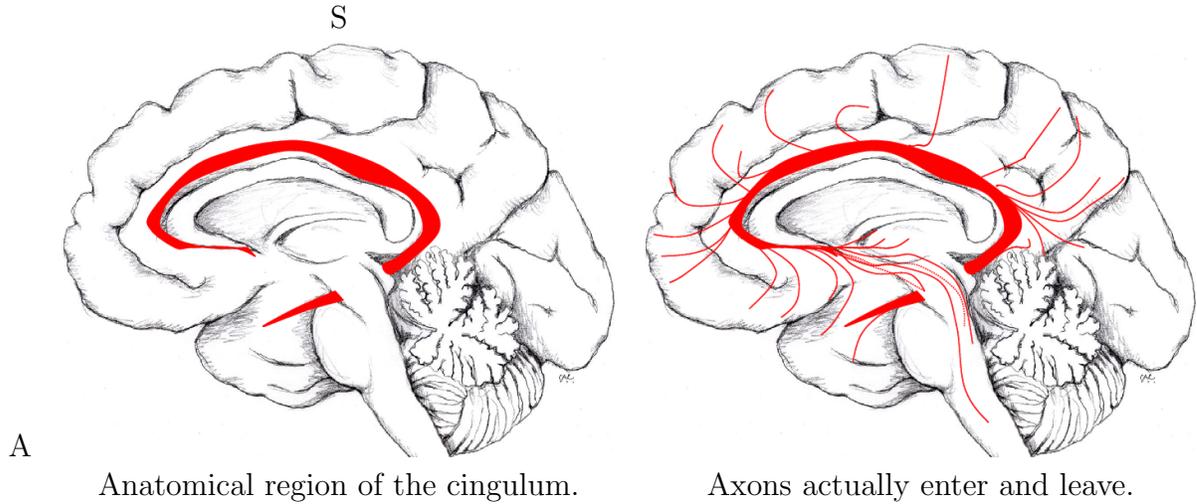


Figure 2-6: Axons enter and leave some fiber tracts along their course, as seen in this example of the cingulum bundle. Images courtesy of James Fallon, Professor of Anatomy and Neurobiology at the University of California, Irvine.

have been reported in schizophrenia [63], and asymmetry of the uncinate fasciculus across hemispheres was found in a postmortem dissection study [48].

2.2.4 The Inferior Occipitofrontal Fasciculus

The inferior occipitofrontal fasciculus (Figure 2-5) connects occipital and frontal lobes. It runs from the frontal lobe, through the temporal lobe (passing below the insula), to the occipital lobe [84].

2.2.5 The Inferior Longitudinal Fasciculus

The inferior longitudinal fasciculus (Figures 2-4 and 2-5) connects occipital and temporal lobes. It is involved in the ventral stream for visual processing, the object recognition pathway [25]. Its connections have been investigated using DTI tractography [25].

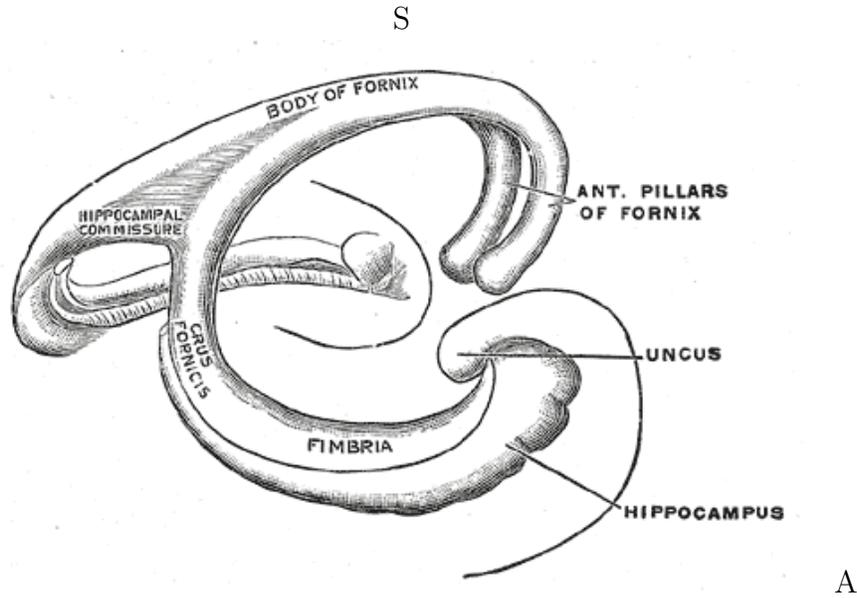


Figure 2-7: The fornix. From Gray's Anatomy [46].

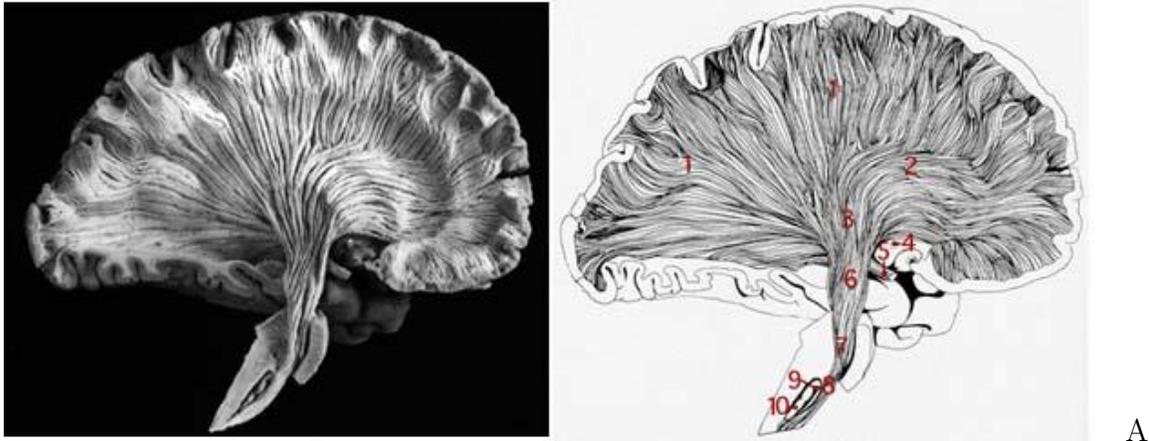
2.3 Projection Fibers

The projection fibers mainly travel along the superior-inferior axis [116]. The main projection fibers visible with DTI tractography are those which travel in the corona radiata/internal capsule/cerebral peduncle regions, and those which enter and leave the cerebellum.

2.3.1 Corona Radiata/Internal Capsule/Cerebral Peduncle Region

A dissection of this region is shown in Figure 2-8, and Figure 2-9 gives an illustration of the anatomical names of each part of this region. In many DTI papers these regions are referred to collectively as the corona radiata [116] because fiber tracing produces a continuous structure. Tracts contained within the corona radiata and internal capsule are named by regions they connect such as the cortico-spinal, cortico-thalamic, and cortico-pontine fibers, and their function varies depending on the regions connected.

S



A

Figure 2-8: Projection fibers (view from the right). 1. Corona radiata 2. Anterior thalamic radiation 3. Internal capsule 4. Anterior commissure 5. Optic tract 6. Cerebral peduncle 7. Longitudinal pontine fibers (corticospinal and corticonuclear tracts) 8. Pyramidal tract of medulla oblongata 9. Hilus of olivary nucleus 10. Olivary nucleus. From Virtual Hospital [128].

2.3.2 Cerebellar peduncles

The cerebellar peduncles (Figure 2-10) are inflow and outflow projection tracts to and from the cerebellum, which is responsible for sensory and motor processing such as control of posture and on-the-fly correction of limb motions [84]. The middle cerebellar peduncle carries input to the cerebellum from the contralateral side of the pons (the synapse point for many neurons from the motor and sensory cortices) [84]. The inferior cerebellar peduncle contains cerebellar inputs from the spinal cord and brainstem, while the superior cerebellar peduncle mainly contains cerebellar outputs which travel to the red nuclei and the thalamus [84].

2.4 Abbreviations

Table 2.2 gives anatomical abbreviations that are used elsewhere in the thesis.

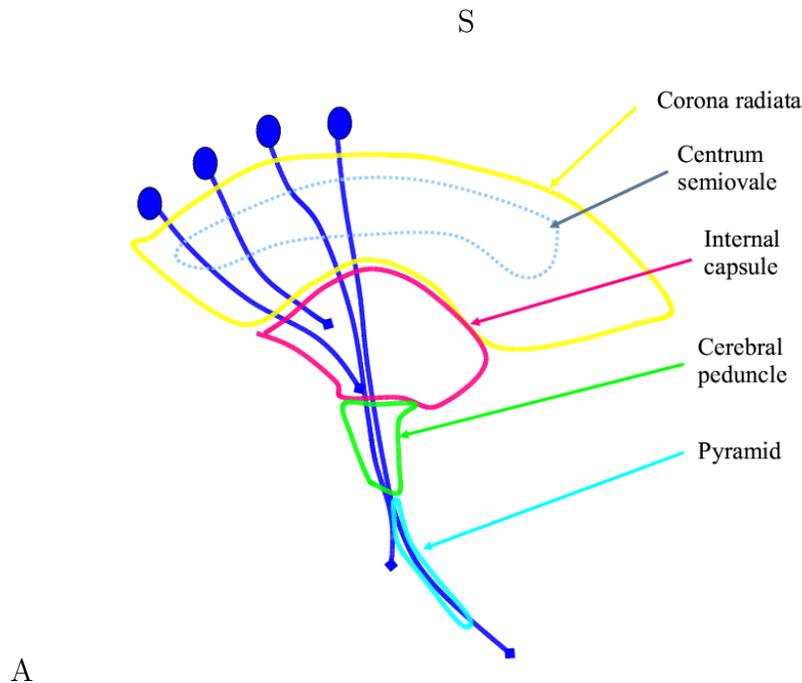


Figure 2-9: Corona radiata/internal capsule/cerebral peduncle region names (view from the left). Courtesy of James Fallon, Professor of Anatomy and Neurobiology at the University of California, Irvine.

2.5 Summary

In this chapter we have presented a brief overview of white matter neuroanatomy which is helpful for understanding the results presented later in the thesis. First of all, the brain is divided into gray matter (containing neuronal cell bodies and dendrites) and white matter (containing insulated, or myelinated, axons which are white)¹. The gray matter is responsible for the integration of information and the white matter handles the transmission of information.

The axons in the white matter form many bundles which are called white matter fiber tracts. These tracts are divided into three types: commissural (which cross hemispheres), association (which connect regions of the same hemisphere), and projection (which project to/from the cortex or cerebellum). Major tracts of each of the three types have been described in this chapter and shown in images. Later in

¹However there are also other types of cells in the gray and white matter (in addition to neurons), and some neurons are completely in the gray matter (so their axons do not enter the white matter).

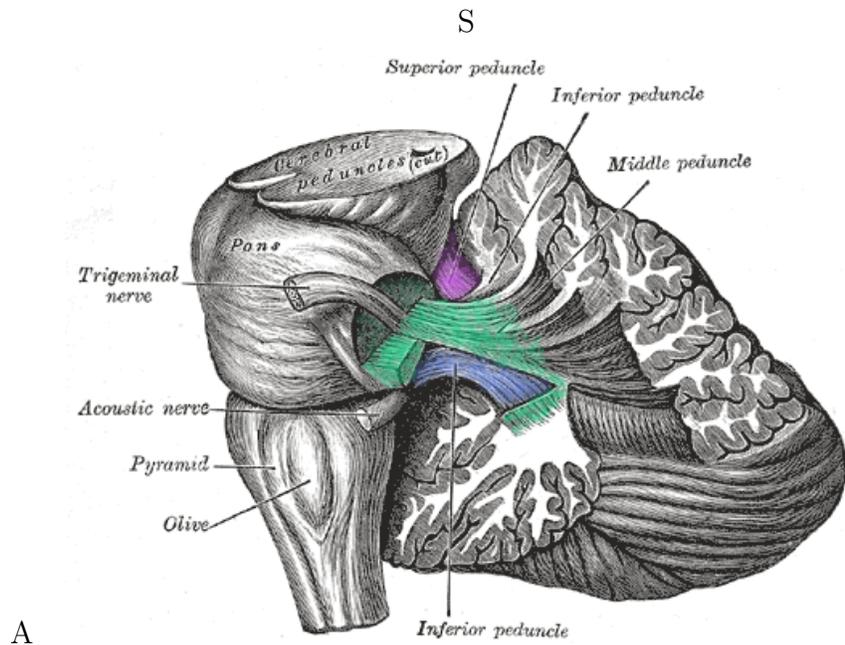


Figure 2-10: Cerebellar peduncles. Color was added to show the three peduncles: blue for inferior, green for middle, and purple for superior. The cerebellum is the “brain-like” structure on the right side of the image. Adapted from Gray’s Anatomy [46].

the thesis we present segmentations of tractography that include most of the regions presented in this chapter.

Neuroscientific studies of the white matter tracts have been performed using diffusion imaging, and specific examples were given in this chapter. Differences in tract integrity have been found in various disease states such as Alzheimer’s, schizophrenia, and epilepsy. In some tracts such as the uncinate and arcuate fasciculi, asymmetry has been demonstrated across hemispheres. The automatic tract segmentation and anatomical labeling methods presented in this thesis enable quantitative neuroscience studies of disease and symmetry in specific white matter tracts.

AF	arcuate fasciculus
CB	cingulum bundle
CC	corpus callosum
CST	corticospinal tract
FOR	fornix
IFO	inferior fronto-occipital fasciculus
ILF	inferior longitudinal fasciculus
OR	optic radiation
SLF	superior longitudinal fasciculus
UF	uncinate fasciculus
WM	white matter

Table 2.2: Anatomical abbreviations

Chapter 3

Background on Diffusion Tensor Imaging and Analysis

The power of diffusion MRI lies in the fact that the diffusion of water molecules probes tissue structure at very small scales, much smaller than the voxel resolution. This allows measurement of the voxel-averaged effects of collisions of water molecules with cellular membranes. In tissues which have a fibrous structure, such as white matter and muscle, diffusion is anisotropic (varies with direction), and the direction of fastest diffusion parallels the fibrous structure. Diffusion MRI is the only method that presently allows measurement of white matter fiber orientation in the human brain in vivo.

Diffusion MRI is different from standard structural MR imaging which measures quantities related to tissue composition (the T1 and T2 relaxation times depend on water and fat content [21]). In diffusion MRI, the quantity measured also relates to the three-dimensional organization of the tissue. Consequently, it is not always possible to transfer image visualization or analysis methods directly from standard structural imaging to diffusion imaging. In this chapter we give background information including an introduction to diffusion imaging and visualization, and a description of the current state of the art in diffusion imaging analysis methods.

3.1 The Diffusion Coefficient

The macroscopic process of diffusion is described by Fick’s first law, derived by Adolf Fick in 1855 [126]. It relates a concentration difference to a flux (a flow across a unit area), stating that the flux, j , is proportional to the gradient of the concentration, ∇u . The proportionality constant d is the diffusion coefficient.

$$j = -d\nabla u. \tag{3.1}$$

For an anisotropic material, in general the flux vector field does not follow the concentration gradient directly, since the material properties also affect diffusion. Consequently, the diffusion tensor, \mathbf{D} , is introduced to model the material.

$$j = -\mathbf{D}\nabla u. \tag{3.2}$$

Microscopically, the process of diffusion consists of the random thermal (Brownian) motion of particles, as observed by Robert Brown in 1827 [125]. The preceding macroscopic equations refer to the diffusion of one substance in another (for example ink in water), but when there is only one type of molecule, it still experiences random motion though there is no concentration gradient. The self-diffusion coefficient is defined as the diffusion coefficient of a species in the absence of a chemical potential gradient [75]. It is the self-diffusion coefficient of water which is measured in diffusion MRI.

However, the self-diffusion coefficient measured in diffusion MRI is not the true self-diffusion coefficient of water (unless the imaged object is pure water), because diffusion is restricted by the local structure of the tissue. This is why diffusion imaging is useful. The “apparent self-diffusion coefficient” or ADC is a function of the diffusion time and the geometry of the local volume, and in addition can be affected by non-diffusive processes including tissue perfusion and fluid flow [71].

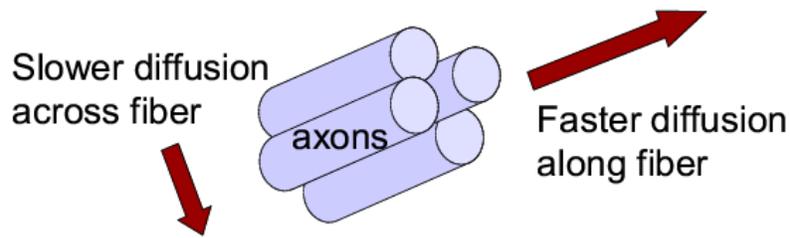


Figure 3-1: Illustration of anisotropic diffusion.

3.2 Anisotropic Diffusion in White Matter

Figure 3-1 contains a diagram of the effect of fibrous tissues on water diffusion, in the ideal case of a fibrous tissue which is uniformly oriented. Faster diffusion is measured along the fiber direction. As in this example, when the diffusion coefficient is not the same in all directions the diffusion is called anisotropic. The measured macroscopic anisotropy is due to microscopic tissue heterogeneity [7]. In the white matter of the brain, diffusion anisotropy is primarily caused by cellular membranes, with some contribution from myelination and the packing of the axons [9, 103, 94]. A model of axon packing and diffusion has been used to support the theory that the myelin sheath restricts diffusion in the axon, causing the overall mean diffusion coefficient to be determined primarily by the water in the extraaxonal compartment [103].

The diffusion effects are averaged over a voxel, complicating the biophysical interpretation of the ADC values. To give an idea of the complexity of the human brain and the size/time scales of the diffusion imaging experiment, Table 3.1 lists relevant quantities such as the number of neurons in the brain (10^{11}) and the distance over which water diffuses during an imaging experiment (1-15 μm , a distance similar to the diameter of an axon).

To measure anisotropic diffusion using MRI, magnetic field gradients are employed to sensitize the image to diffusion in a particular direction. Due to loss of phase coherence, the magnetization of the tissue decreases in locations where diffusion takes place along that particular direction. This loss of magnetization (signal) creates darker vox-

Quantity	Measurement	Reference
axon packing density (pyramidal tract)	60,000 – 70,000/ mm^2	[94]
axon packing density (corpus callosum)	338,000/ mm^2	[94]
axon diameter (pyramidal tract)	26 μm	[94]
axon diameters in central nervous system	0.2 to 20 μm	[57]
neuron cell body diameter	50 μm or more	[57]
voxel size in diffusion MRI	$2.5 \times 2.5 \times 2.5$ mm	
diffusion time in MRI	30–100ms	[69, 9]
mean water diffusion distance	1-15 μm (in 50-100 ms)	[69]
number of neurons in human brain	100 billion (10^{11})	[57]
synaptic connections per axon	up to 1,000	[57]

Table 3.1: The scale of DTI and the brain: neuron sizes and quantities, and water diffusion times and distances.

els. This means that white matter fiber tracts parallel to the gradient direction will appear dark in the diffusion-weighted image. Example diffusion-weighted images are shown in Figure 3-2. By repeating the process of diffusion weighting in multiple directions, and possibly for various diffusion times and magnetic field gradient strengths, at each voxel a three-dimensional diffusion pattern can be estimated and related to the shape of the underlying anatomy.

3.3 Diffusion Tensor MRI

Diffusion MRI is a generic name referring to any MRI measurement of water diffusion. More specific names exist to describe the ways in which diffusion information may be reconstructed and represented at each voxel. The simplest representation of the shape of diffusion in 3D is a 3x3 symmetric, positive-definite matrix called the diffusion tensor, originally proposed for use in diffusion MRI by Peter Basser in 1994 [6, 7]. When diffusion is represented using this tensor model, the imaging modality is referred to as diffusion tensor MRI (DTI or DT-MRI). DTI models the diffusion profile as

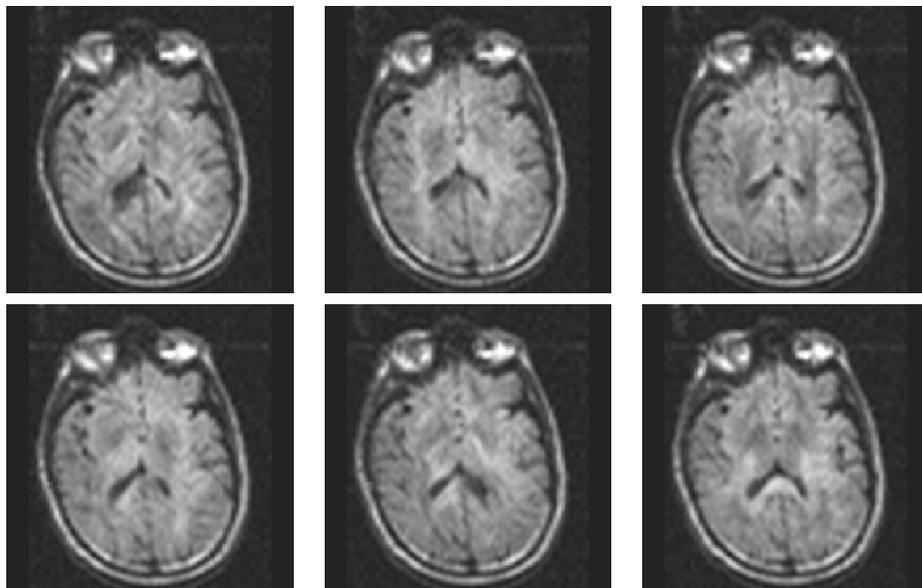


Figure 3-2: Six diffusion-weighted images (the minimum for tensor calculation). In diffusion MRI, magnetic field gradients are employed to sensitize the image to diffusion in a particular direction. The direction is different for each image, resulting in a different pattern of signal loss (dark areas) due to anisotropic diffusion.

Gaussian¹ as described in [7]:

$$p(x|x_0, \tau) = \frac{1}{\sqrt{|\mathbf{D}_{\text{eff}}(\tau)|(4\pi\tau)^3}} \exp \left[\frac{-(x - x_0)^T \mathbf{D}_{\text{eff}}^{-1}(\tau)(x - x_0)}{4\tau} \right]. \quad (3.3)$$

Here \mathbf{D}_{eff} is the effective self-diffusion tensor (the meaning of “effective” is analogous to the meaning of “apparent” in the term apparent self-diffusion coefficient), τ is the diffusion time, x is the final position of the diffusing water molecule, and x_0 is its initial position. For simplicity, from now on we will refer to the effective self-diffusion tensor, $\mathbf{D}_{\text{eff}}(\tau)$, as the diffusion tensor, \mathbf{D} .

It can be seen from (3.3) that the diffusion tensor is proportional to the covariance matrix of water molecule displacements during the diffusion time. Comparing (3.3) to the standard equation for a multivariate normal distribution with covariance Σ ,

¹The equation for a standard multivariate normal distribution (Gaussian) in n dimensions is

$$p(x) = \frac{1}{\sqrt{|\Sigma|(2\pi)^n}} \exp \left(-\frac{1}{2}(x - \mu)^T \Sigma^{-1}(x - \mu) \right).$$

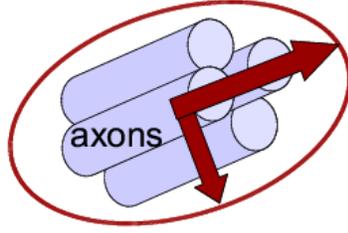


Figure 3-3: Illustration of the tensor eigensystem in a region of anisotropic diffusion. The ellipse represents diffusion in two dimensions, while three-dimensional diffusion would be described by an ellipsoid.

we obtain

$$\Sigma = 2\tau\mathbf{D} \quad (3.4)$$

so the constant of proportionality is two times the diffusion time. The eigenvalues (λ_i) of the tensor are similarly related to displacements: the mean effective diffusion distance in the i th principal direction is $\sqrt{2\tau\lambda_i}$ [7]. The eigenvalues are the apparent diffusion coefficients in the principal directions of diffusion, with units of $\frac{mm^2}{s}$ (millimeters squared per second). The apparent diffusion coefficient (d_g) in any direction \hat{g} can be calculated from the tensor model using the formula

$$d_g = \hat{g}^T \mathbf{D} \hat{g}. \quad (3.5)$$

Figure 3-3 is a schematic representation of the diffusion tensor in a region of anisotropic diffusion. The major eigenvector of the diffusion tensor points in the principal diffusion direction (“pdd,” or direction of the fastest diffusion). In anisotropic fibrous tissues the major eigenvector defines the fiber tract axis of the tissue [7], and thus the three orthogonal eigenvectors can be thought of as a local fiber coordinate system. In three dimensions an ellipsoid can be drawn to represent an isosurface of (Gaussian) diffusion probability: the axes of the ellipsoid are aligned with the eigenvectors and their lengths are $\sqrt{2\tau\lambda_i}$ [7].

In simplified terms, diffusion imaging works by introducing extra gradient pulses

whose effect “cancels out” for stationary water molecules, and causes a random phase shift for molecules which diffuse. Due to their random phase, signal from diffusing molecules is lost, and this decreased signal (S_k) is compared to the original signal (S_0) to calculate the diffusion coefficient or tensor. In DTI, the diffusion tensor field is calculated from a set of diffusion-weighted images by solving the Stejskal-Tanner equation (3.6) [2]. This equation describes how the signal intensity at each voxel decreases in the presence of diffusion:

$$S_k = S_0 e^{-b \hat{g}_k^T \mathbf{D} \hat{g}_k}. \quad (3.6)$$

Here S_0 is the image intensity at the voxel (measured with no diffusion-sensitizing gradient) and S_k is the intensity measured after the application of the k th diffusion-sensitizing gradient. \hat{g}_k is a unit vector representing the direction of this diffusion-sensitizing magnetic field gradient. \mathbf{D} is the diffusion tensor, and the product $\hat{g}_k^T \mathbf{D} \hat{g}_k$ represents the diffusion coefficient in direction \hat{g}_k . In addition, b is LeBihan’s factor describing the pulse sequence, gradient strength, and physical constants. For rectangular gradient pulses the b-factor is defined by $b = \gamma^2 \delta^2 (\Delta - \frac{\delta}{3}) |g|^2$, where γ is the proton gyromagnetic ratio (42 MHz/Tesla), $|g|$ is the strength of the diffusion sensitizing gradient pulses, δ is the duration of the diffusion gradient pulses, and Δ is the time between diffusion gradient RF pulses [123]. The b-factor is near 0 for the image S_0 which is T2-weighted, and the b-factor is near 1,000 for the other images S_k in DTI. Equation 3.6 can be solved via the least squares method at each voxel. In order to calculate the 6 independent numbers in the 3×3 symmetric matrix \mathbf{D} , at least 7 images are needed: 6 diffusion-weighted images from 6 gradient directions (giving S_k) plus one baseline image (giving S_0). For more information on diffusion imaging and the tensor calculation process, see for example [6, 123].

3.4 Other Diffusion Models

It is known that the Gaussian model of diffusion employed in DTI cannot represent complicated geometry in a voxel such as crossing or branching tracts. To overcome this issue, high angular resolution diffusion (HARD) imaging with multiple tensor [113] and q-ball [114] reconstructions has been proposed. High angular resolution data uses the fiber orientation distribution function (ODF), a probability density function on the unit sphere, to describe diffusion at each voxel [112, 114].

The method presented in this thesis does not depend on the type of diffusion imaging, since its input consists of estimates of white matter fiber paths (tractography). Tractography can be produced in multiple ways, using various types of diffusion data. In this introduction our focus is diffusion tensor imaging due to the fact that DTI is the type of data we have analyzed to date, and because DTI is the current standard in research and clinical use.

3.5 Scalar Invariants of the Diffusion Tensor

A scalar invariant is a function of the tensor that outputs a scalar which is the same regardless of the coordinate frame of the tensor. It is invariant to rotation of the tissue in the MRI magnet because it can be expressed as a function of the eigenvalues of the tensor, which are invariant to rotation. In this section we will describe commonly used scalar invariant quantities, which can be divided into two categories: diffusion magnitude measures and anisotropy measures. We will use $\lambda_1 \geq \lambda_2 \geq \lambda_3 \geq 0$ to refer to the eigenvalues of the symmetric, positive-definite diffusion tensor \mathbf{D} .

3.5.1 Measures of Diffusion Magnitude

The simplest and most clinically useful scalar invariant is the tensor trace.

$$\text{trace}(\mathbf{D}) = \sum_{i=1}^3 \mathbf{D}_{ii} = \sum_{i=1}^3 \lambda_i \quad (3.7)$$

The average of the eigenvalues ($\text{trace}(\mathbf{D})/3$) is also used. This average is referred to as the mean diffusivity, or MD [70]; the bulk mean diffusivity, or $\langle D \rangle$ [5]; or the ADC map (often output by DTI software). Note that the multiple uses of the term “ADC” are confusing because ADC originally meant the apparent diffusion coefficient in a particular direction. For example, in early diffusion imaging the parallel and perpendicular ADC were measured with gradient directions parallel to and perpendicular to fibers in a sample [9]. Additionally, in clinical imaging “ADC maps” may be calculated with fewer gradients than needed for the tensor. Finally, the trace may also be referred to as the “trace ADC” as in [28].

The trace and mean diffusivity relate to the total amount of diffusion in a voxel, which is related to the amount of water in the extracellular space. The trace is clinically useful in early stroke detection because it is sensitive to the initial cellular swelling (cytotoxic edema) which restricts diffusion [100]. The trace is high in cerebrospinal fluid, around $9.6 \times 10^{-3} \frac{\text{mm}^2}{\text{s}}$, and relatively constant in normal brain parenchyma (white and gray matter), between $1.95 \times 10^{-3} \frac{\text{mm}^2}{\text{s}}$ and $2.2 \times 10^{-3} \frac{\text{mm}^2}{\text{s}}$ [94]. For comparison, the self-diffusion coefficient of water at body temperature of 37°C is $3 \times 10^{-3} \frac{\text{mm}^2}{\text{s}}$ [69], which would give a trace of $9 \times 10^{-3} \frac{\text{mm}^2}{\text{s}}$. ADCs measured in ventricles and edema can be higher than in water due to fluid flow or enhanced perfusion, respectively [71].

3.5.2 Anisotropy Measures

Tensor anisotropy measures are ratios of the eigenvalues that are used to quantify diffusion shape. They are useful for describing the amount of tissue organization and for locating voxels likely to contain a single tract. The following measures are normalized and all range from 0 to 1, except for the mode which ranges from -1 to +1.

Linear, planar, and spherical measures

The most intuitive measures are C_L , C_P , and C_S : the linear, planar, and spherical shape measures [123, 124]. They describe whether the shape of diffusion is like a cigar (linear), pancake (planar), or sphere (spherical).

$$C_L = \frac{\lambda_1 - \lambda_2}{\lambda_1} \quad (3.8)$$

$$C_P = \frac{\lambda_2 - \lambda_3}{\lambda_1} \quad (3.9)$$

$$C_S = \frac{\lambda_3}{\lambda_1} \quad (3.10)$$

In voxels with high planar or spherical measure, the principal eigenvector will not always match an underlying fiber tract direction (where tracts cross the eigenvector may point to neither one). But if the largest eigenvalue is much larger than the other two eigenvalues, the linear measure will be large, giving evidence for the presence of a single fiber tract. Note that these measures can be normalized by λ_1 , by the trace, or by $\sqrt{\lambda_1^2 + \lambda_2^2 + \lambda_3^2}$.

FA and mode

The fractional anisotropy, or FA [4], is the most widely used anisotropy measure². Its name comes from the fact that it measures the *fraction* of \mathbf{D} which is *anisotropic*. Two possible expressions for FA follow ($\hat{\lambda}$ is the mean diffusivity).

$$\text{FA} = \frac{1}{\sqrt{2}} \frac{\sqrt{(\lambda_1 - \hat{\lambda})^2 + (\lambda_2 - \hat{\lambda})^2 + (\lambda_3 - \hat{\lambda})^2}}{\sqrt{\lambda_1^2 + \lambda_2^2 + \lambda_3^2}} = \frac{\sqrt{3}}{\sqrt{2}} \frac{\text{norm}(\mathbf{D} - \frac{1}{3}\text{trace}(\mathbf{D})\mathbf{I})}{\text{norm}(\mathbf{D})} \quad (3.11)$$

While FA basically measures how far the tensor is from a sphere, another measure is needed for the other axis, between linear and planar anisotropy. This is given by the mode, a recently introduced scalar invariant which is orthogonal to the FA measure and relates to the skewness of the eigenvalues [38]. Its definition includes

²FA is the standard measurement for studies of white matter integrity as demonstrated in Table 3.3.

the determinant and Frobenius norm of the anisotropic part of the tensor³.

$$\text{mode} = \frac{(-\lambda_1 - \lambda_2 + 2\lambda_3)(2\lambda_1 - \lambda_2 - \lambda_3)(-\lambda_1 + 2\lambda_2 - \lambda_3)}{2(\lambda_1^2 + \lambda_2^2 + \lambda_3^2 - \lambda_1\lambda_2 - \lambda_1\lambda_3 - \lambda_2\lambda_3)^{3/2}} \quad (3.12)$$

$$= (3\sqrt{6})\det\left(\frac{\mathbf{D} - \frac{1}{3}\text{trace}(\mathbf{D})\mathbf{I}}{\text{norm}(\mathbf{D} - \frac{1}{3}\text{trace}(\mathbf{D})\mathbf{I})}\right) \quad (3.13)$$

3.6 DTI Visualization

Unlike conventional scalar MRI images, DTI is fundamentally three-dimensional, in that the quantity measured at each voxel is diffusion information in 3D. This poses a visualization (and analysis) challenge. In this section we briefly describe visualization methods; for a thorough explanation of the anatomy that can be seen with DTI, the reader is referred to [50, 80, 116].

DTI may be visualized in a slice plane (a section through the data) or in three dimensions, depending on the subset of the data that is presented. Planar visualization methods are voxel-based, meaning an image is generated to display information from the tensor that is in each voxel in one slice plane. For example, images may be displayed of any anisotropy measure, or of the trace. Another type of image can represent the major eigenvector field using a mapping to colors. The color scheme commonly used to represent the orientation of the major eigenvector works as follows: blue is superior-inferior, red is left-right, and green is anterior-posterior [90]. The brightness of the color is controlled by tensor anisotropy (FA). Another voxel-based visualization method uses small three-dimensional objects called glyphs to display information from each tensor eigensystem. Example glyphs include “sticks” representing the orientation of the major eigenvector, ellipsoids related to the diffusion isoprobability surfaces [7], and superquadric tensor glyphs [58]. The dominant method for three-dimensional visualization of DTI is tractography, a very commonly employed method which estimates the trajectories of major fiber tracts in the white matter [3]. It is closely related to an earlier method for visualization of tensor fields known as hyper-

³Thanks to Gordon Kindlmann for the fully simplified formula expressed as a function of the eigenvalues.

streamlines [34]. Volume rendering strategies for three-dimensional visualization of tensor fields have also been proposed [59].

Examples of two- and three-dimensional visualization are provided in Figures 3-4 and 3-5. Figure 3-4 demonstrates voxel-based slice visualization, including simple glyphs and an RGB image. Note the glyphs “poke through” the slice plane. Figure 3-5 is an elaborate three-dimensional tractography visualization which was created by a neuroanatomist. These and other DTI images in this document were generated with the DTMRI module in 3D Slicer, software written by the author in conjunction with the Laboratory for Mathematics in Imaging (LMI) at BWH and Harvard Medical School. (Additional matlab code was written by the author and used to perform the automatic segmentation of tractography shown elsewhere in this document.)

3.7 Tractography

This section summarizes current methods for DTI-based estimation of connected regions in the white matter. The central theme is tracing paths by following probable tract orientations, in order to reconstruct an estimate of the underlying white matter fiber structure. Many methods have been proposed in the literature for addressing this problem, and most produce output which corresponds well to known anatomy in regions where the data is not made ambiguous by crossing fibers. In this section we summarize existing tractography methods. Table 3.2 lists the methods covered in this section.

First, however, we give a clarification of terminology. A single path from tractography is often somewhat inaccurately referred to as a “tract” or a “fiber” in the literature [3]. Though these names are motivated by anatomy they are incorrect for several reasons, which is of concern because the anatomical names may inspire misplaced confidence in novices to DTI. At current voxel resolutions, each path is an estimate of (part of) the trajectory of some larger fiber tract (mm diameter), and has no direct correspondence to smaller features like individual axons (μm diameter). In addition, many paths are needed to fill the volume of a region that could be consid-

Method	Description	References
streamline	Follows principal diffusion direction (pdd)	[3, 30, 79, 122]
FMT	Surface evolution based on pdd	[92]
flow-based	Surface evolution based on pdd or ODF	[23]
diffusion sim.	Connectivity based on arrival time of front	[8]
diffusion sim.	Steady-state flow field estimated	[85]
geodesic path	Connection is shortest path using \mathbf{D}^{-1} metric	[85]
optimal path	Simulated annealing to find most probable path	[112]
probabilistic	Bayesian approaches with various fiber models	[12, 16, 41]

Table 3.2: Tractography methods.

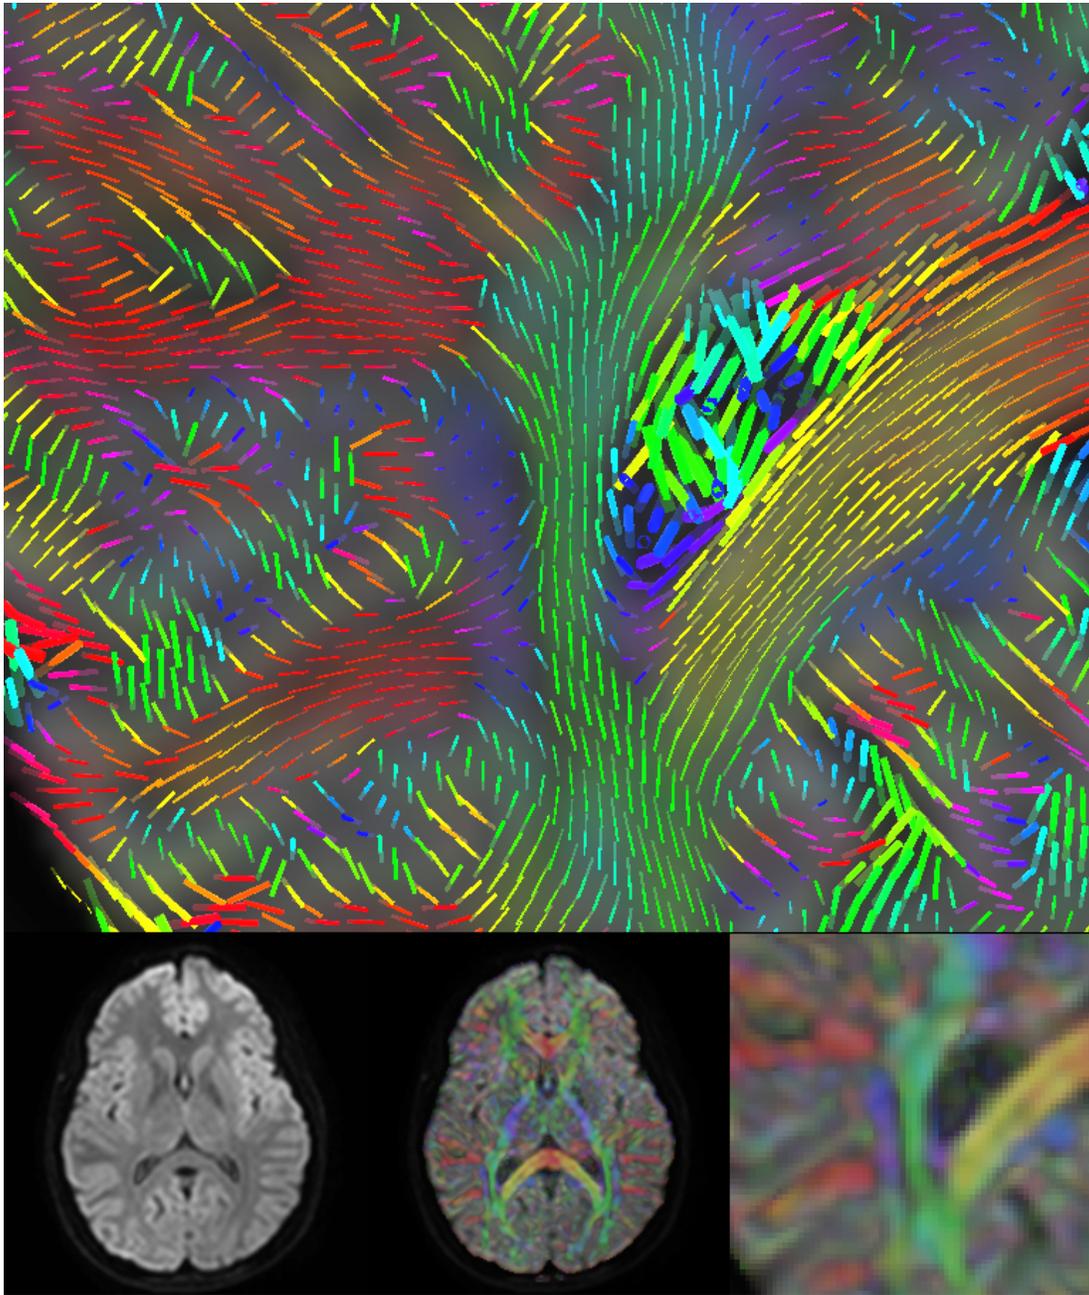


Figure 3-4: Voxel-based DTI visualization in 3D Slicer. Note that the color red indicates right-left orientation, blue is superior-inferior, and green is anterior-posterior. The top image contains glyphs representing major eigenvector orientations in the region of the corpus callosum (yellow and red) and right lateral ventricle. The cingulum can be seen in blue, and the posterior limb of the internal capsule in green. The bottom images are the average diffusion-weighted image (left), the average diffusion-weighted image with the RGB orientation map as a semi-transparent overlay (middle), and the zoomed in region from the top image, without the glyphs (right).



Figure 3-5: Three-dimensional DTI visualization in 3D Slicer. Tractography seed locations, colors, and transparencies were manually chosen by a neuroanatomist. Courtesy of James Fallon, professor of Anatomy and Neurobiology at UCSD.

ered an estimate of a fiber tract. So basically, the volume occupied by one path (say it has some diameter) is smaller than a tract but larger than an axon, and probably larger than a fiber. When referring to one path output from tractography, we will attempt in this thesis to use the terms “path,” “tractographic path,” or “trajectory,” but the more common and descriptive “fiber” will also be used. When referring to a collection of paths that have similar trajectories and approximate the course of a white matter fiber tract, we will use the word “tract.” Now we move on to the overview of tractography methods.

3.7.1 Streamline Tractography

This is the most common approach and is generally what is meant by the word “tractography.” Streamline tractography for DTI [3, 30, 79, 122] has been a topic of research for more than half a decade. The first step in performing streamline tractography is to associate the major eigenvector with the tangent to a curve (the putative fiber path). Then the curve may be estimated by stepping repeatedly in the direction of the tangent. An important consideration is that an eigenvector has an orientation but not a direction (plus or minus the eigenvector is still an eigenvector) so consistency of orientations must be checked on each step. This can be done simply by taking the dot product of the current and previous tangent, and switching the sign of the current tangent if the dot product was negative. Another important consideration is calculation of the tangent at an arbitrary location, which may be done for example by interpolating the tensor there (component-wise) and then computing the major eigenvector.

To estimate the path from its tangents, standard numerical solution methods for differential equations can be used. These include Euler’s method (following the tangent for a fixed step size), second order Runge-Kutta (also known as the midpoint method, where the tangent is followed for half a step, then a new tangent is calculated at the midpoint of the interval and used to take the full step), and fourth order Runge-Kutta (where the weighted average of four estimated tangents to the curve is used when taking each step) [95]. A fixed step size of one millimeter or less (smaller than a

voxel) is generally used for DTI data. The application of the Euler and Runge-Kutta methods to white matter tractography was explored in [3, 30]. Another popular method, called FACT, takes only one step per voxel, so the step size varies and each voxel’s tensor is used directly without interpolation [79]. Some related methods attempt to introduce “inertia” when tracking through regions of planar anisotropy. These methods modulate the incoming tangent direction by the tensor instead of directly using the major eigenvector of the tensor [118, 68, 122, 123], an approach which is related to the power method of computing eigenvectors. A similar method, generalized streamline tractography, was created for tracking through ODF (fiber orientation distribution function from high angular resolution diffusion data) fields and selects the highest peak within 15 degrees of the incoming trajectory [112].

The standard tractographic curve estimation approach has one main drawback, which is that all decisions are made locally. Thus errors can accumulate and the tracing can be confounded by regions of crossing fibers (with high planar or spherical indices). This is demonstrated schematically in Figure 3-6. To ameliorate this problem, one may restrict fiber tracking to areas of high linear anisotropy (the approach taken in this thesis), employ a tensor deflection algorithm (with many parameter settings), or design a method that can branch to follow multiple paths. Various approaches that attempt to allow branching are covered in the following sections.

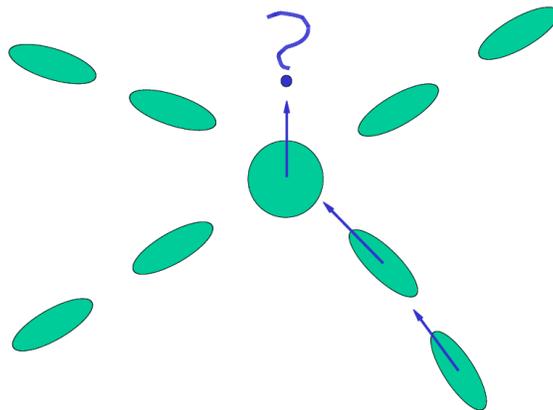


Figure 3-6: The major eigenvector does not point in the direction of a fiber tract in the case of crossing fibers.

3.7.2 Surface Evolution, Diffusion Simulation, and Optimization Tractography

Surface evolution, diffusion simulation, and optimization algorithms all attempt to allow branching, effectively by following all possible paths out from a start point. However it is not always clear how to interpret the results, and perhaps for this reason these methods have not yet found widespread use. We differentiate these methods from Bayesian approaches which also allow branching, because there is no probabilistic model proposed for tract orientation at each voxel; rather the algorithms are motivated by physical analogies or optimality of solutions. Some of the methods use the major eigenvector only, and some use the full tensor or a fiber orientation distribution function (ODF), a probability density function on the unit sphere. In general, approaches which use only the major eigenvector have the advantage that the solutions generated are far sharper, and consequently look more like what is expected anatomically.

Surface Evolution

The fast marching tractography (FMT) method evolves a surface using the fast marching method, where the speed function is dependent on the principal eigenvector field [92]. This seems to be the only method in Section 3.7.2 which has been applied to group analysis of DTI. Another method evolves a surface in a field of vectors created perpendicular to the major eigenvector field, so that the surfaces will tend to enclose the tracts [22]. The same group later presented a method for flow-based fiber tracking, which they categorize as an extension of the FMT technique and apply to both DTI and ODF data [23].

Diffusion Simulation

An iterative diffusion simulation approach used the full tensor and quantified connection strengths based on a probabilistic interpretation of the arrival time of the diffusion front [8]. Another diffusion simulation approach using the full tensor (by

the author of this thesis along with Steven Haker and Carl-Fredrik Westin) placed sources (releasing “solute”) and sinks (absorbing “solute”) in the tensor field and solved for the steady-state concentration of the solute and the flow vector field [85].

Optimization

As an alternative to local approaches (which include basically all other tractography methods), these algorithms try to fit a globally optimal path between points.

In prior work which is not otherwise covered in this thesis, the author (with Steven Haker and Carl-Fredrik Westin) investigated a geodesic (shortest) path approach to tractography. The connection (“fiber path”) between pairs of points was defined as the geodesic connecting those points. The inverse of the diffusion tensor was used as a metric tensor to calculate distances (so high diffusion was low distance) as described in [85]. To restrict paths to the white matter a soft C_L mask was used to multiply the tensors, encouraging paths in highly anisotropic regions. The paths produced were optimal in the geodesic sense, however a path was calculated between any two input points even if there was no anatomical connection there, complicating analysis of the results. An example distance map, resulting from computing distances from the seed point to all other points, is displayed in Figure 3-7.

Another tractography method used simulated annealing to find the optimal path between two points of interest through an ODF field, where path probability was quantified using the ODF and curvature at each step [112].

3.7.3 Probabilistic Tractography

Instead of using either only the major eigenvector, or the full tensor, these methods place a probability model on the fiber orientation at each voxel. Rather than producing one path from each seed point, a distribution of paths is produced by sampling. Like the previous methods, these methods allow branching (in the sense that many paths are produced so the total output can have a branching structure) and have the advantage that their interpretation is more straightforward due to a Bayesian

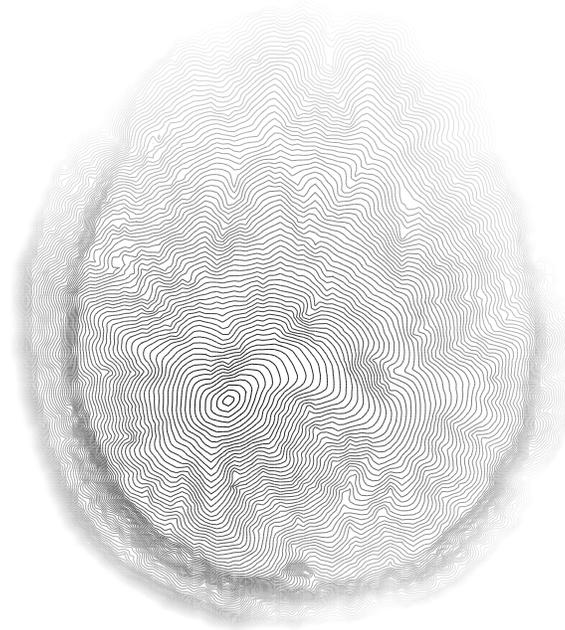


Figure 3-7: Isocontours of tensor-based distance map

framework. In the case of models with one fiber direction, the uncertainty about this direction is captured in the width of the distribution about the mean. The most successful of these methods, by Behrens et al., was used to segment the thalamus based on its cortical connections (the whole thalamus could be segmented because probabilistic tractography can extend further into gray matter than streamline tractography) [11, 12, 10]. Other probabilistic tractography methods include [16, 42, 41]. The disadvantage of these methods in comparison with streamline tractography is the explosion of tractographic path data to view or analyze. For one seed point, streamline tractography outputs a single highly probable path while probabilistic tractography can produce an infinite number of paths, though usually on the order of 1,000 are generated per seed point, and this can take minutes or hours [41].

3.8 Tractography Grouping Methods

If (streamline) tractography is seeded in all voxels in the white matter, from 10,000 to 100,000 paths may be produced depending on the white matter mask and the voxel

size. Even when seeding in smaller regions, organization of the paths is necessary if the application requires anatomical structures to be isolated from each other. For example, Figure 3-5 is an elaborate tractography image created by a neuroanatomist using 3D Slicer software. The tractography was organized by hand into anatomical regions and given transparency to allow viewing of deeper structures. Automatically achieving such organization is one of the goals of this thesis. Manual and automatic methods for organizing the results of tractography are described in this section.

Multiple ROI (Interactive)

Manual interactive grouping of tractographic paths using multiple selection regions of interest (ROIs) [30] (also known as “virtual dissection” [24]) has been performed to create a fiber tract atlas [80] and in several clinical studies ([28, 54, 93]). In fact, Partridge et al. found tractography-based definitions of a pyramidal tract ROI to be more reproducible than manual ROI drawing [93].

Clustering (Automatic)

There are three main goals which should be satisfied by an automatic tractography segmentation algorithm: automatic grouping of like trajectories, region/cluster correspondence across subjects, and anatomical labeling of regions/clusters. Our method is the first to achieve these three goals. Most related work in automatic clustering has focused on the first goal, automatic grouping. Several methods have been proposed for clustering of tractography. Here we give a brief summary; more details and additional related work will be presented in Chapter 5. An early ad-hoc algorithm was proposed to subdivide existing tractography into bundles, and was applied to paths passing through a small region of the internal capsule [35]. Automated tractographic path grouping using clustering algorithms has been proposed by Brun et al. [18], Gerig et al. [43], O’Donnell et al. [87], and Corouge et al. [33]. Automated path grouping via atlas-based labeling of tractography was described by Maddah et al. [73] who manually created a tractography atlas and gave a method for transferring its labels to a novel subject. The preceding methods all use streamline tractography

as input, however one method by Johansen–Berg et al. uses probabilistic tractography and connectivity matrix reordering to find the boundary between the supplementary motor area (SMA) and the presupplementary motor area (preSMA) [51].

3.9 DTI Group Analysis Methods

Here we develop a framework for organizing the many current approaches to group analysis of DTI. We then situate the thesis contribution in this milieu.

Standard structural image analysis is based on comparison of scalar values in one of two ways, region of interest-based (ROI-based) or voxel-based morphometry (VBM). The first method uses ROIs which are manually or automatically defined (e.g. regions in the image such as the thalamus or parts of the cortex), and then compares features (volume, average scalar, etc.) within those regions. This requires an a priori hypothesis that specific regions are involved in the functionality or pathology which is being studied. The second method performs group registration (usually by spatial normalization, or registration to a template image [110]) and compares some function of the scalar values at each voxel. All voxels are analyzed simultaneously, and no a priori hypothesis is required, however sensitivity could be reduced due to blurring or misalignment.

Diffusion imaging presents a challenge to these established methods. First, the quantity at each voxel is no longer a scalar. To address this, many clinical studies have used scalar values such as FA which can be derived from the tensor and have a physiological interpretation related to white matter integrity. Second, the nature of the data itself is different from standard scalar images, in that there is a geometric quantity (white matter fibers) underlying the discrete samples of diffusion. Some analysis methods have attempted to reconstruct this underlying geometry using tractography. Finally, group registration via spatial normalization is more complicated with tensor data because the tensors must be re-oriented, and because there is no standard tensor template brain in common use, so scalar values are generally registered.

We propose a division of current DTI analysis into categories based on analysis

region (ROI, VBM, or tractography), data analyzed (scalar, vector, tensor, or tract anatomy), and group correspondence method employed (spatial normalization, tract correspondences, or tract alignment). Many of the methods in these categories are novel to DTI. In the rest of this section we describe each category, providing example techniques and/or studies in each. Finally in Table 3.3 a large sample of DTI group analysis techniques from the current literature is organized using these categories.

3.9.1 Analysis Region: ROI, VBM, or Tractography

The analysis region refers to the spatial part of the data which is studied. ROI and VBM methods are complimentary: VBM methods may be used in hypothesis development, followed by ROI-based methods to confirm that VBM results are not artifactual. With tensor data, the ROI and VBM approaches function similarly standard scalar analysis, but a decision must be made about what data to analyze (scalar/vector/tensor) and what data to register (for VBM, generally registration uses FA or the T2-weighted image with $B=0$, though one study used tensors [62]).

An interesting novel VBM approach called tract-based spatial statistics (TBSS) finds a group average FA skeleton, then locally aligns it with maximal FA values in the direction perpendicular to the skeleton. These values are associated with the skeleton points (voxels) and VBM is performed on the FA values [105].

DTI analysis adds a new method for defining ROIs based on tractography (either whole tract ROIs or cross-sectional ROIs within a tract). For delineating ROIs based on tractography, current studies have all used the manual interactive selection method, though one methods paper by Corouge et al. [33] advocates clustering to remove outliers from bundles.

3.9.2 Data Analyzed: Scalar, Vector, Tensor, or Anatomy

With DTI there are two new types of data to analyze, tensors and tract anatomy. Most studies measure the FA, and some also use the trace of the tensor or the eigenvalues. It appears that no studies to date have compared the full tensor, however

some statistical foundation has been laid for analysis of major eigenvector data [102]. While most studies have measured FA, interestingly in [35] the measurement of ADCs parallel to and perpendicular to fiber bundles was demonstrated (presumably the average of the two lesser eigenvalues was used for the perpendicular ADC).

Some novel analyses have focused on anatomical questions, where the data analyzed was anatomical connection information from a tractography method [51, 25, 26]. Catani et. al. have used “virtual dissection” (multiple ROI guided) tractography in group-averaged DTI data to perform anatomical studies. Their findings support the existence of an inferior longitudinal fasciculus connecting occipital and temporal lobes [25] and suggest a novel indirect language pathway lateral to the traditional anatomical c-shaped arcuate fasciculus [26]. As well as tractography in group-averaged DTI they demonstrate similar results in single subjects. In addition, as mentioned previously, Behrens et al. have used probabilistic tractography to segment the thalamus according to cortical connection regions [11], and to define the border between pre-SMA and SMA (the latter was then group mapped to investigate the location of the border in multiple subjects) [51].

3.9.3 Group Correspondence: Normalization, Tract Correspondence, or Tract Alignment

Two new population correspondence methods have been introduced in the literature as alternatives to group alignment (spatial normalization). The motivation for this is twofold. First, the sensitivity to errors in spatial normalization is likely to be increased with tensors relative to scalars, due to the fact that the tensor has to be reoriented as well as transformed (so errors are possible in position and in local orientation), and the fact that spatial smoothing is less appropriate with tensors (since componentwise averaging is appropriate for studies of trace but will decrease FA). Secondly, since the underlying anatomical unit is a fiber tract, not a voxel, it makes sense to measure group differences in regions which comprise tracts.

The first population correspondence method is correspondence of tract ROIs (i.e.

obtaining the region of the entire corpus callosum, or the splenium of the corpus callosum, or the cingulum, etc., in all subjects). The second correspondence method is finer-grained: once tract correspondences are known, a spatial tract alignment may be calculated. Current approaches entail finding point correspondences along tracts, which so far have been defined relative to a reference point that is chosen to be anatomically reasonable for the tract of interest [33, 45].

Another variant of the tract correspondence is to actually use spatial normalization (calculated from FA or T2 data) to put images derived from tractography into a common coordinate system, and to then perform group mapping of the images derived from tractography. In one study, voxel probability maps were generated from tractography in the corpus callosum of ten subjects [132]. Here, the voxel probability was defined as the percentage of subjects having a path from tractography at a particular voxel location. Other related methods are listed in Table 3.3.

3.10 Classification of Thesis Work in Proposed Categories

The work presented in this thesis provides an automatic way to group and label (segment) paths from tractography. It fits in the following way into the proposed DTI group analysis categories of region analyzed, alignment, and data analyzed. The region analyzed can be either tractography ROIs from the automatic segmentation, or voxel ROIs defined from the segmented tractography (we demonstrate voxel ROI generation for the example of the midsagittal plane of the corpus callosum in Chapter 6). We employ two types of alignment. Our method for creation of an atlas that models the anatomy of a population uses group registration for alignment. However the clustering algorithm, which is then applied to the multi-subject paths from tractography, uses additional fiber path similarity information which can be robust to slight misalignments. To label additional subjects, each is spatially aligned to the atlas, and then the tractography is labeled, giving tract correspondences across subjects

Method	Region	Align	Data	Analysis	Ref
	T R V r O B a I M c t	N T T o C A r o l m r g r n	S V A c e n a c a l t t		
FMT	X X	X X	X	group-mapped CC, OR, CST	[92, 27]
Catani	X	X X	X	AF subregions, ILF in human	[25, 26]
Xu	X	X X	X	10 subject map of 5 CC regions	[132]
Behrens	X	X X	X	thalamic segmentation	[11]
Kubicki	X	X	X	FA and MTR in schizophrenia	[62]
Park	X	X	X	FA symmetry analysis	[91]
TBSS	X	X	X	mean FA skeleton VBM	[105]
Tuch	X X	X	X	FA & visual reaction time	[115]
Salat	X X	X	X	FA drop w/ age in WM ROIs	[98]
Schwart.	X	X	X X	pdd statistics (dyslexia)	[102]
Lu	X		X	peritumoral FA, MD	[106]
Nakam.	X		X	FA vs. personality traits and symptoms in schizotypal personality disorder	[81]
Gong	X	X X	X	FA in CB (parametric tract model)	[45]
Corouge	X	X X	X	FA, curvature on aligned CC region	[33]
Concha	X	X	X	FA, Tr, λ in FOR and CB in epilepsy	[28]
Jones	X	X	X	FA in CB, UF, IFO, SLF (age effect)	[54]
Partridge	X X	X	X	FA, MD, λ in tract voxels; tract ROIs reproducible	[93]
Kanaan	X	X	X	FA in genu of CC	[56]
O'Donnell	X X	X X	X	automatic tract labeling and correspondence across subjects, voxel ROI generation	[88]

Table 3.3: Categorization of DTI group analysis studies by method. From top to bottom: anatomy studies, VBM studies, ROI studies, and tractography-based studies/methods. Columns from left to right are: Region analyzed (tract, region of interest, VBM or all voxels), Alignment method (normalization, tract correspondence, tract alignment), Data analyzed (scalars, vectors, tract anatomy), Analysis (more information about the region, quantity, or disease studied), Ref (references). Anatomical abbreviations are as in Table 2.2.

(the second alignment category we have defined). Finally, the data which has been analyzed with our method has so far been limited to scalars.

Chapter 4

Background on Spectral Clustering

Spectral clustering groups data points using a mapping, or embedding, of the points that is constructed from the eigenvectors of a data-dependent matrix. In this chapter we give technical details of the clustering method we employ. To make things more concrete, however, we begin with a simple clustering example using paths from tractography.

4.1 Tractography Clustering Example

Figure 4-1 shows example input paths on the left and output clusters on the right. The arrow in the center represents the clustering process.

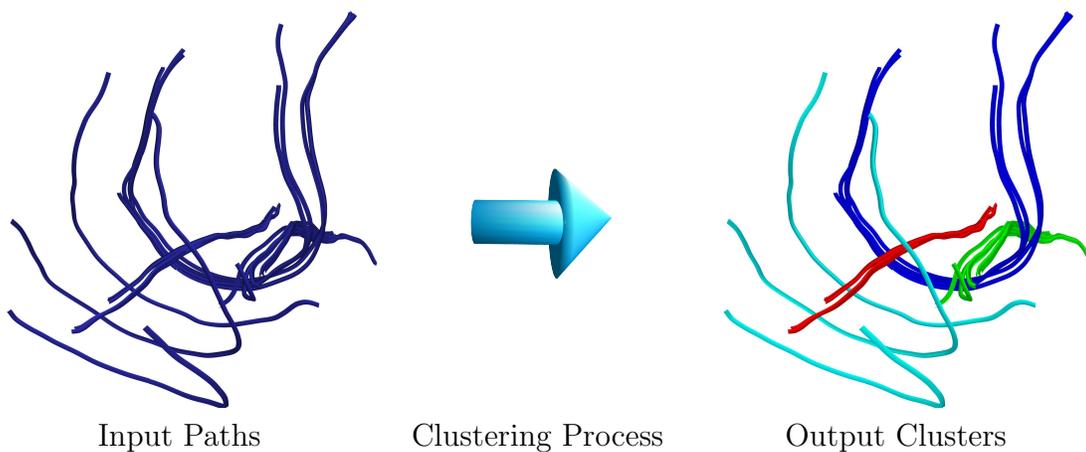


Figure 4-1: Tractography clustering example for illustration of spectral clustering.

One way to think of the input to the clustering process is as a graph, where each input path from tractography is a node in the graph, and the nodes are connected by edges. For each edge a number, or edge weight, is calculated to describe the similarity or affinity of the nodes connected by the edge. The clustering process is then a way to isolate groups of nodes in the graph by cutting edges which represent low affinities. In order to represent the paths from tractography using a weighted graph, we first have to decide how to measure path similarity. Then, once we are armed with an affinity function¹ that can measure the pairwise similarity between any two paths, many clustering algorithms become available for our use. Figure 4-2 contains a possible weighted graph for the clustering example from Figure 4-1.

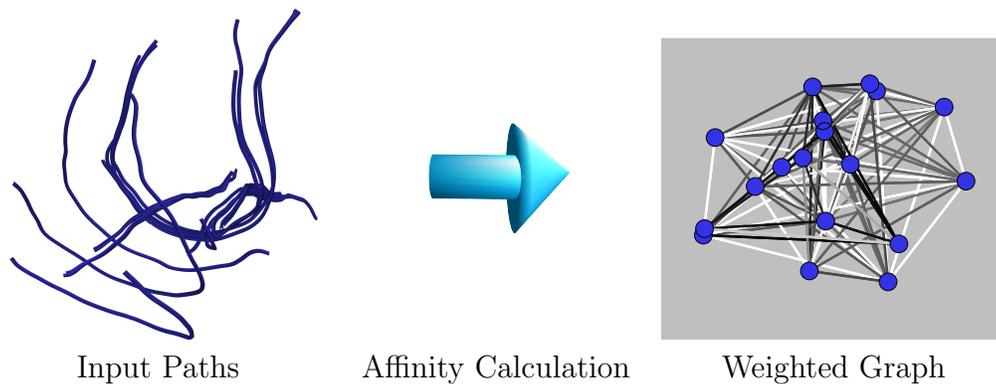


Figure 4-2: Weighted graph example. Each path (left) can be thought of as a node in a graph (right). The edges in the graph are colored according to the affinity between the nodes they connect (the highest affinity is colored white). This coloring represents the actual affinity values calculated for the input trajectories on the left (using our method which is detailed further in Chapter 5).

The clustering algorithm we employ is called spectral clustering, and it uses only these pairwise affinity values (edge weights in the graph) as input. The affinity values are placed in a matrix, and spectral clustering analyzes this matrix, finding the most important affinity information by calculating the eigenvectors of the matrix. The affinity matrix corresponding to the graph in Figure 4-2 is displayed in Figure 4-3. Each row and each column corresponds to a path (a node), and each entry in the matrix corresponds to an edge. The most important thing about this matrix is that

¹The specific affinity function we use is discussed in Chapter 5.

the affinity values are high (white) where paths are similar.

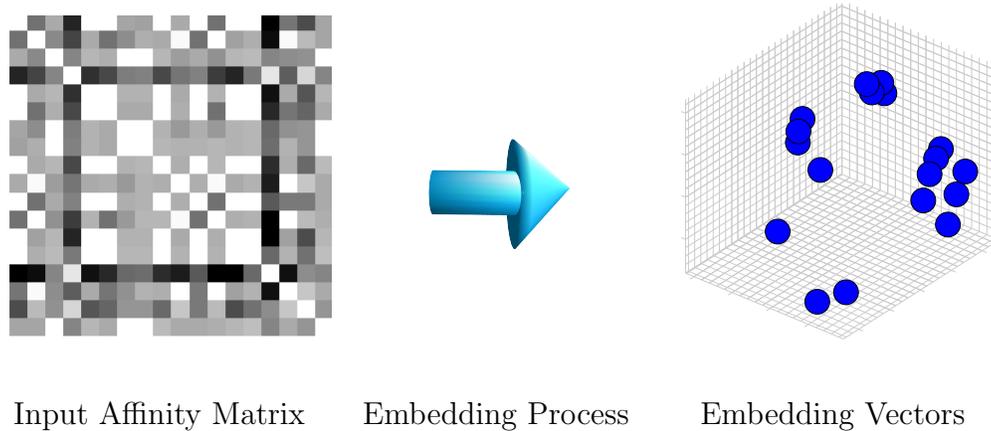


Figure 4-3: Example affinity matrix and embedding vectors. Each row of the affinity matrix (left) contains information about one input path from Figure 4-1. So the path is represented in this matrix by its affinities to all other paths. High affinities are white and low affinities are black. The arrow in the center represents the embedding process that uses the eigenvectors of the affinity matrix. On the right the embedding vector corresponding to each path is plotted. These embedding vectors give a three-dimensional representation of the affinity relationships between the pairs of input paths. In this particular example, the cluster information is represented well in three dimensions; however when there is a higher number of clusters inherent in the data, then more dimensions are needed to separate the clusters.

The eigenvectors corresponding to the largest eigenvalues of a matrix contain “most” of the information about the matrix, so these eigenvectors can be used to approximate the matrix. Because the highest eigenvectors of the affinity matrix can be used to approximate the matrix, they can also be used to create a more compact representation² of the affinity information that was calculated for each input path from tractography. In our example, the input information about one path is 18 affinity values, and this information can be represented more compactly with 3 values by using the eigenvectors of the affinity matrix. These 3 values form a three-dimensional vector representing the path, and once the three-dimensional vectors are calculated for all paths, they can all be represented in the three-dimensional space shown in Figure 4-3. In spectral clustering, this space is called the embedding space, and the basis vectors of the space are (scaled versions of) the eigenvectors of the affinity matrix.

²Specific equations for performing spectral embedding are provided later in this chapter.

The vectors which represent each path in this space are called embedding vectors. In Figure 4-3 the embedding vector is displayed for each input path.

The meaning of the embedding vectors is relatively intuitive: they were created using only the affinity information between pairs of paths, so paths with relatively similar affinity information must have similar embedding vectors. Paths which have similar affinity information tend to be paths which are similar to each other, so this means that paths which are similar to each other will have similar embedding vectors. Simply put, if paths are alike they should map near each other in the embedding space.

After performing the embedding, the affinity relationships between paths are represented spatially, but another step is still needed to segment clusters in the data. This can be done using a standard clustering method in the embedding space. One method which is commonly used is called k-means clustering, an iterative method that represents each cluster based on the mean (“centroid”) of the points in the cluster, and essentially moves these centroids around until the clusters are all compact (in the sense that the total sum of distances from centroids to points in their cluster is minimized). Once the clusters have been labeled, the rows and columns of the affinity matrix can be reordered according to cluster, to visually inspect the within-cluster and between-cluster affinities. The reordering puts the within-cluster affinity information close to the diagonal, forming blocks in the example. The clustering solution for the example, using $k = 4$ clusters, is shown in Figure 4-4.

The fact that similar paths are embedded near each other can be seen in our simple example after clustering has been performed. In Figure 4-5, one can see that the red points (embedding vectors) are very close to each other, because those paths are very similar in the tractography image in Figure 4-5. On the other hand, the cyan dots are more spread out in the embedding space, because those paths are relatively further apart and thus less similar (according to our definition of similarity which has not been discussed yet and is covered in Chapter 5). One advantage of spectral clustering is that it only depends on pairwise affinity values; therefore it can find clusters with arbitrary shapes in input space (for example points on concentric circles) as long as the affinity values are created such that those clusters map to

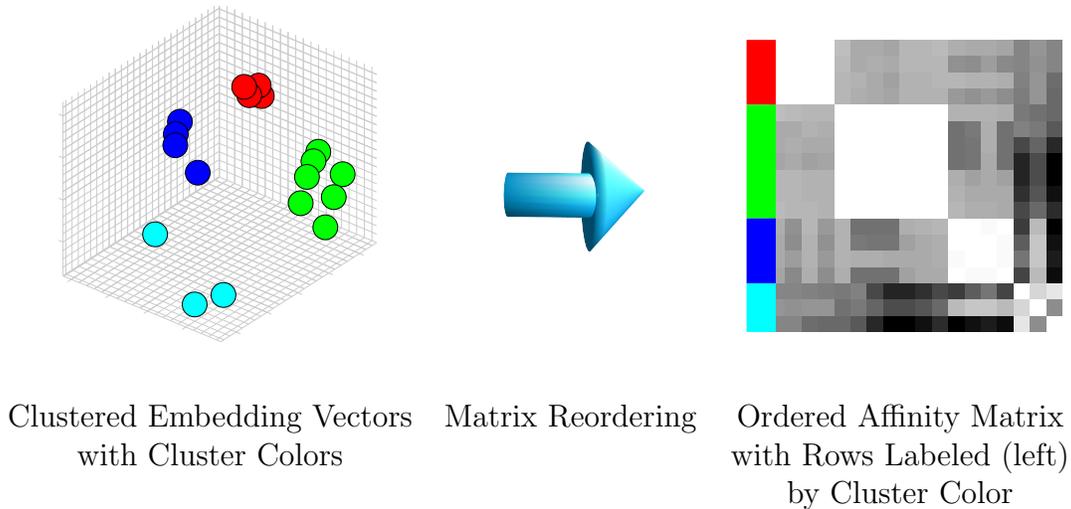
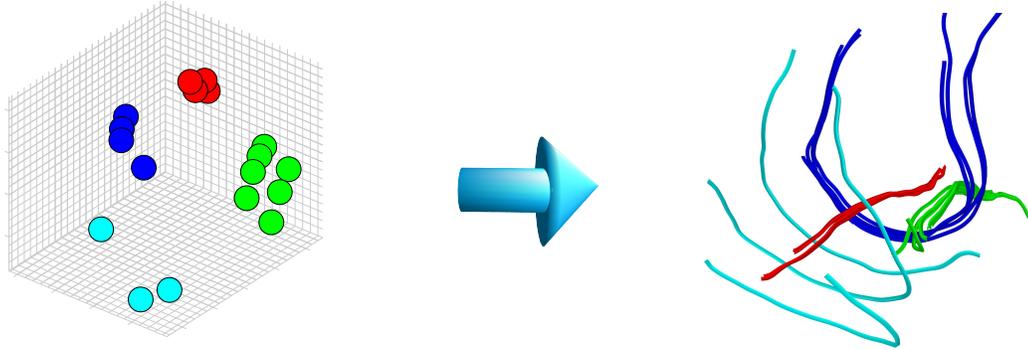


Figure 4-4: Example embedding vectors and affinity matrix, after clustering. The clusters (left) are found in embedding space using k-means. Then the matrix can be re-ordered (permuting rows and columns) so that rows and columns are ordered by cluster (right). In this example the matrix becomes nearly block diagonal and each block represents affinities within a cluster. (Permutation of rows and columns of a matrix does not change the eigenvectors of the matrix so the matrix reordering is a valid operation.) Each row of the affinity matrix (right) still contains affinities for one input path from Figure 4-1. The colorbar to the left of the matrix indicates the cluster membership of the embedding vector corresponding to each matrix row. So the white square in the upper left (next to the red part of the colorbar) represents the high affinity values within the red cluster and the larger white square toward the middle of the matrix contains affinities from the green cluster. In the lower right part of the matrix, the affinities for the cyan-colored cluster are not as high.

compact, easily-segmented “blobs” in embedding space. Thus the performance of the clustering method is determined by the performance of the affinity measure for identifying cluster structure in the data.

4.2 Chapter Overview

In the rest of this chapter we mainly focus on explaining the embedding process that was represented by the arrow in Figure 4-3. We first give a more precise explanation of spectral embedding, followed by background on matrices whose eigenvectors are used for spectral embedding. Then we present selected historical and recent spectral embedding methods: multidimensional scaling (MDS), Hall’s quadratic placement



Clustered Embedding Vectors Trajectory Coloring Clustered Trajectories

Figure 4-5: Example embedding vectors and trajectories, after clustering. The colors indicate cluster membership. It is clear that similar trajectories (right) map near each other when embedded as points (left) in the embedding space.

algorithm, and the Normalized Cuts spectral clustering method (which we use in this thesis to group paths from tractography). We then provide specific details on the Nystrom approximation to Normalized Cuts, which we employ to allow spectral clustering of large datasets. At the very end of the chapter we return to our simple tractography clustering example and summarize it in the context of the Normalized Cuts method.

The formulae in this chapter use boldface to denote vectors and matrices, and calligraphic font (such as \mathcal{A}) to denote normalized versions of the matrices. The eigenvectors themselves are normalized to have length of one and thus have no inherent size, however for convenience we use the informal terminology and refer to the eigenvector associated with the smallest eigenvalue as the “smallest eigenvector”, and that associated with the largest eigenvalue as the “largest eigenvector.” In general, the input data to clustering is referred to as “data points” so for generality we use this terminology instead of referring to paths from tractography. In addition, when providing embedding examples to illustrate the methods presented in this chapter, the input data will be points in two dimensions, rather than tractography.

4.3 Spectral Embedding

“Spectral” refers to the use of eigenvectors, and “embedding” means that the mapping is one-to-one (an injection) [120]. The goal in spectral embedding is to find some useful representation of the input data in an embedding space. There are a variety of algorithms that create embeddings whose properties satisfy different goals. For example in principal components analysis (PCA) the goal is to use the embedding to represent the variance in the input with as few dimensions as possible. In multidimensional scaling (MDS) the goal is to preserve known distances between input points. For spectral clustering the goal is to find an embedding that groups points that satisfy some notion of what forms a good cluster. A mathematical formulation of the particular application-specific goal (generally in the form of an affinity measure) gives rise to a matrix whose eigenvectors can be used to perform embedding. The embedding of each input data point is calculated using one of two closely-related methods: projection onto the eigenvectors of the input matrix, or direct use of the rows of the eigenvector matrix. Each algorithm may also apply some type of scaling to produce a final embedding.

For spectral clustering, the embedding process is usually explained and analyzed in terms of a block-diagonal affinity matrix (see for example [119, 82]) where each cluster has perfect affinity of one to all points in the cluster, and affinity of zero to other points. In these analyses, it is shown that the leading eigenvectors “pick out” the clusters, as we have demonstrated in the toy example in Figure 4-6. In the rest of this section we give two simple mathematical explanations regarding the relationships of the embedding vectors to the original input matrix and to its eigenvectors. In our application, properties of the embedding space are of interest, and these explanations are helpful in addition to the usual block-diagonal affinity matrix analysis.

4.3.1 Embedding Vectors and the Input Matrix

An embedding method takes as input a real symmetric matrix, \mathbf{A} , where each row of \mathbf{A} corresponds to one input data point, and the entries in the row describe its

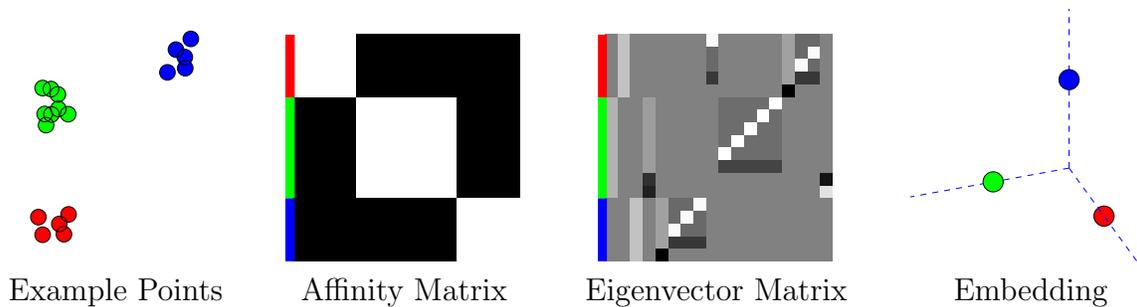


Figure 4-6: Block diagonal affinity matrix example. Example data points that could give rise to this affinity matrix are shown on the left. The matrix has already been ordered by cluster to illustrate the block diagonal concept. The colorbars to the left of the matrices indicate cluster membership of the point corresponding to each matrix row. Inside the affinity matrix, white is affinity of 1 and black represents 0. The “largest” three eigenvectors (leftmost columns of the eigenvector matrix) describe the cluster structure, and the rest are not useful for clustering. (The number of connected components of this graph is related to the number of nonzero eigenvalues, and in this case the eigenvalues are $[8, 5, 5, 0, \dots, 0]$. The eigenvectors corresponding to eigenvalues of 0 are not useful for clustering. Note that in a more realistic example there may be only one connected component in the graph and the affinity values are not all 1’s, so there are multiple nonzero eigenvalues.) Embedding using the three useful eigenvectors produces the coordinates on the right.

relationship to the other data points. (These relationships may for example be defined as affinity values.) Each row of \mathbf{A} is then associated with an embedding vector, \mathbf{e} , which is the result of mapping the original data point into the embedding space. Without considering the specific optimization problem used to construct \mathbf{A} , we look at the standard equation for diagonalization of a symmetric matrix:

$$\mathbf{A} = \mathbf{U}\mathbf{\Lambda}\mathbf{U}^T. \tag{4.1}$$

The columns of \mathbf{U} contain the orthonormal eigenvectors of \mathbf{A} , and $\mathbf{\Lambda}$ is a diagonal matrix containing \mathbf{A} ’s real eigenvalues. By interpreting rows of \mathbf{U} (and columns of \mathbf{U}^T) as embedding vectors, \mathbf{e} , the equation becomes an inner product

$$\mathbf{A}_{ij} = \mathbf{e}_i^T \mathbf{\Lambda} \mathbf{e}_j \tag{4.2}$$

where \mathbf{e}_i and \mathbf{e}_j are the embedding vectors³ in rows i and j of the eigenvector matrix \mathbf{U} . This shows that \mathbf{A}_{ij} is the dot product between those embedding vectors, where $\mathbf{\Lambda}$ is the diagonal metric tensor. (A metric tensor has the same distance scaling function as the covariance matrix does when it is used in computing Mahalanobis distances. The expression for the dot product in a space with a Riemannian metric is $\langle \mathbf{x}_1, \mathbf{x}_2 \rangle_{\mathbf{M}} = \mathbf{x}_1^T \mathbf{M} \mathbf{x}_2$ where \mathbf{M} is the metric tensor.) The squared lengths of the embedding vectors are $\langle \mathbf{e}_i, \mathbf{e}_i \rangle_{\mathbf{\Lambda}} = \mathbf{e}_i^T \mathbf{\Lambda} \mathbf{e}_i$ which are the values on the diagonal of \mathbf{A} .

The matrix \mathbf{A} need not be positive definite to perform embedding using kernel PCA⁴, spectral clustering, or multidimensional scaling [101, 14, 76] which means that $\mathbf{\Lambda}$ could contain negative eigenvalues. In practice those negative eigenvalues and eigenvectors are generally not used in creating the embedding vectors. When only positive eigenvalues are used, the selected subspace of the embedding space has a positive definite metric. Further discussion of this issue related to kernel PCA can be found in [101], page 15.

The discussion in this section supports the intuition that the general purpose of spectral embedding is to put similar data points near each other. It further explains why the stereotypical block diagonal spectral clustering example produces clusters that are separated by 90 degrees (corresponding to dot products of 0) [82], and why one embedding method normalizes all embedding vectors to length one (to bring points with low angular separation onto a hypersphere and nearer each other for easier clustering) [82]. In most methods a subset of the largest eigenvectors is used for embedding, giving an approximation of the matrix \mathbf{A} , or equivalently an estimate of the requested dot products⁵. In fact, it has been shown that this “truncation of

³This is one possible way to define the embedding vectors. Other functions of \mathbf{e} , generally limited to scaling of rows or columns of the \mathbf{U} matrix, may be used and some are described later in the chapter. Also, some embedding methods employ only the vectors \mathbf{e} and others also use $\mathbf{\Lambda}$.

⁴Kernel methods are algorithms that can be formulated in terms of dot products of the input data. For example kernel PCA requires a definition of the dot product between input data points, and it is equivalent to performing standard principal component analysis in the higher-dimensional space where that dot product holds. The actual mapping into that space does not need to be computed, however, because the method only depends on the dot products.

⁵However this does not directly apply to methods discussed later in this chapter, which are based on the Laplacian of a graph. There the problem is posed as a minimization and the smallest eigen-

the eigenbasis” moves high-affinity points nearer each other, and low-affinity points further apart, enhancing the cluster structure of the data [17].

As we discuss later in this chapter, the goal of the various embedding methods (PCA, MDS, spectral clustering, etc.) determines both the process for construction of \mathbf{A} and the equation used to calculate the embedding vectors.

4.3.2 Embedding Vectors and Eigenvectors of the Input Matrix

The relationship between embedding using rows of the eigenvector matrix (done in MDS and spectral clustering) and projecting the rows of \mathbf{A} onto the eigenvectors (done in PCA) can be seen from the following:

$$\mathbf{A} = \mathbf{U}\mathbf{\Lambda}\mathbf{U}^T \tag{4.3}$$

$$\mathbf{A}(\mathbf{\Lambda}\mathbf{U}^T)^{-1} = \mathbf{U} \tag{4.4}$$

$$\mathbf{A}\mathbf{U}\mathbf{\Lambda}^{-1} = \mathbf{U}. \tag{4.5}$$

The left hand side of (4.5) is a projection of the rows of \mathbf{A} onto a scaled eigenvector basis (where each eigenvector is divided by the corresponding eigenvalue), and this gives the right hand side, the rows of \mathbf{U} . This correspondence between rows of \mathbf{A} and rows of \mathbf{U} explains why in spectral embedding methods, row i of \mathbf{U} is used to produce the embedding vector \mathbf{e}_i for data point i , corresponding to row i of \mathbf{A} .

4.4 Matrices Used for Spectral Embedding

Here we define three types of matrices used for spectral embedding: Gram matrices, Laplacians of graphs, and Markov matrices. (Standard PCA uses a covariance matrix but that is not be discussed here.) In this section we give background information

vectors are used, those which contribute least to the dot products. Normalization of the Laplacian, however, produces an affinity matrix whose largest eigenvectors are used. Also, the optimization problem associated with the Laplacian embeds similar points near each other.

on each type of matrix. Section 4.5 then describes the use of Gram matrices and Laplacian matrices in spectral embedding.

4.4.1 The Gram Matrix

A Gram matrix is a symmetric, positive definite matrix containing dot products between all pairs of some set of vectors. It is positive definite if the kernel used to compute the dot products is a Mercer kernel [20]. This type of matrix is used in multidimensional scaling (Section 4.5.1). The affinity matrices used in many spectral clustering methods can also be seen as Gram matrices (unless they are not positive definite).

4.4.2 The Laplacian Matrix of a Graph

The Laplacian, \mathbf{L} , is a second-difference⁶ matrix that operates on functions defined on the nodes of a graph. In our case, the graph is the representation of the input data points to clustering, where each data point is a node of the graph. A function defined on the graph gives a function value at each node. (Note that the eigenvectors of the various matrices used in spectral embedding are functions defined on the graph.)

The Laplacian matrix is not only used in graph theory: it also describes the behavior of systems of springs and masses and electrical circuits [107]. Its eigenvectors have one physical interpretation as modes of vibration of a system of springs and masses. The Laplacian is derived from the incidence matrix \mathbf{I}_G , a first-difference⁷ matrix giving a geometric description of a graph. A graph $G = (V, E)$ consists of a set V of vertices and a set E of edges $\{i, j\}$. For an unweighted graph the incidence matrix is defined as

$$\mathbf{I}_{G_{ij}} = \begin{cases} \pm 1 & \text{if } \{i, j\} \in E \\ 0 & \text{otherwise} \end{cases} \quad (4.6)$$

where each column corresponds to a vertex and each row corresponds to an edge and

⁶A second difference is a discrete version of the second derivative.

⁷A first difference is a discrete version of the first derivative.

sums to 0 (the ± 1 is arbitrarily chosen, with $+1$ marking one vertex and -1 the other). Then \mathbf{L} is defined as

$$\mathbf{L} = \mathbf{I}_G^T \mathbf{I}_G. \quad (4.7)$$

The on-diagonal elements (\mathbf{L}_{ii}) equal the total number of edges for vertex i , and the off-diagonal elements (\mathbf{L}_{ij}) are -1 if there is an edge $\{i, j\}$, or 0 otherwise. The Laplacian can also be expressed as

$$\mathbf{L} = \mathbf{D} - \mathbf{A} \quad (4.8)$$

where \mathbf{D} is the diagonal degree matrix containing the number of edges for each node, and \mathbf{A} is the adjacency matrix with $\mathbf{A}_{ij} = 1$ if there is an edge $\{i, j\}$, and $\mathbf{A}_{ij} = 0$ otherwise. To refer to the i th entry on the diagonal of \mathbf{D} , \mathbf{D}_{ii} , we use d_i .

The Laplacian is also defined for weighted graphs, where the edge weights of 1 are replaced by E_{ij} for each edge. An example of a weighted graph was displayed at the beginning of this chapter in Figure 4-2. Component-wise, the Laplacian of a weighted graph is:

$$\mathbf{L}_{ij} = \begin{cases} \sum_k E_{ik}, & i = j \\ -E_{ij}, & i \neq j \text{ and } \{i, j\} \in E \\ 0 & \text{otherwise.} \end{cases} \quad (4.9)$$

For a weighted graph, \mathbf{A} , the adjacency matrix, becomes $\mathbf{A}_{ij} = E_{ij}$ if there is an edge $\{i, j\}$, and $\mathbf{A}_{ij} = 0$ otherwise. Generally \mathbf{A}_{ii} is 0 (no self-looping edges) though in current image segmentation applications that consider affinities rather than edges, that may not be the case (see discussion in Section 4.6). For weighted graphs, \mathbf{D} is the diagonal degree matrix containing the sum of edge weights for each node, which is equivalent to the row sum of \mathbf{A} . \mathbf{A} is often called the affinity matrix of a weighted graph because it contains edge weights or affinities, not adjacency indicators. An example affinity matrix was given in Figure 4-3.

In graph theory, \mathbf{L} is symmetric positive semidefinite. Its smallest eigenvalue is 0, corresponding to an eigenvector of all ones (because the row sums are 0). Other

eigenvectors are perpendicular to the vector of ones, thus must have positive and negative entries and sum to zero. Many objective functions related to the Laplacian have been defined for clustering and embedding, and in the ideal case the positive and negative entries of the eigenvectors are used to indicate membership in different clusters. Two methods that we cover in this chapter, the Normalized Cuts spectral clustering method and Hall’s quadratic placement algorithm, as well as others such as Laplacian eigenmaps [13], use matrices related to the Laplacian.

4.4.3 Markov Matrices

Each entry \mathbf{M}_{ij} in a Markov matrix describes the transition probability from state i to state j in a Markov chain. \mathbf{M} has all positive entries and its columns sum to 1 [108]. One normalization of the Laplacian

$$\mathbf{D}^{-1}\mathbf{L} \tag{4.10}$$

$$=\mathbf{D}^{-1}(\mathbf{D} - \mathbf{A}) \tag{4.11}$$

$$=\mathbf{I} - \mathbf{D}^{-1}\mathbf{A} \tag{4.12}$$

presented in [77] produces the normalized affinity matrix $\mathbf{D}^{-1}\mathbf{A}$ which is interpreted as a Markov matrix. Unlike Markov matrices presented by Strang in [108], the rows (not columns) of this matrix sum to 1. When using this matrix for embedding, the embedding is thought of as grouping input data points by the “similarity of their transition probabilities to subsets” of the graph [77]. So if part of the graph is minimally connected to another part, a random walk in that region has low probability of transitioning into the rest of the graph, and this is reflected in the eigenvectors of the matrix.

4.5 Spectral Embedding Methods

Here we discuss a few selected historical and recent embedding methods: multidimensional scaling [76], Hall’s quadratic placement algorithm [47], and the recent Nor-

malized Cuts spectral clustering method [104]. Each solves a different optimization problem by spectral embedding.

The optimizations are in the form of the Rayleigh quotient, $\frac{\mathbf{v}^T \mathbf{L} \mathbf{v}}{\mathbf{v}^T \mathbf{v}}$, which is minimized (maximized) by the eigenvector corresponding to the smallest (largest) eigenvalue when \mathbf{L} is a symmetric matrix [108]. Similarly, for a symmetric matrix \mathbf{M} the minimum of $\mathbf{v}^T \mathbf{M} \mathbf{v}$ subject to the constraint $\mathbf{v}^T \mathbf{v} = 1$ is given by the smallest eigenvalue of \mathbf{M} [107]. The Rayleigh quotient is equal to an eigenvalue when \mathbf{v} is an eigenvector:

$$\mathbf{M} \mathbf{v} = \lambda \mathbf{v} \tag{4.13}$$

$$\mathbf{v}^T \mathbf{M} \mathbf{v} = \mathbf{v}^T \lambda \mathbf{v} \tag{4.14}$$

$$\mathbf{v}^T \mathbf{M} \mathbf{v} = \mathbf{v}^T \mathbf{v} \lambda \tag{4.15}$$

$$\frac{\mathbf{v}^T \mathbf{M} \mathbf{v}}{\mathbf{v}^T \mathbf{v}} = \lambda. \tag{4.16}$$

4.5.1 Multidimensional Scaling

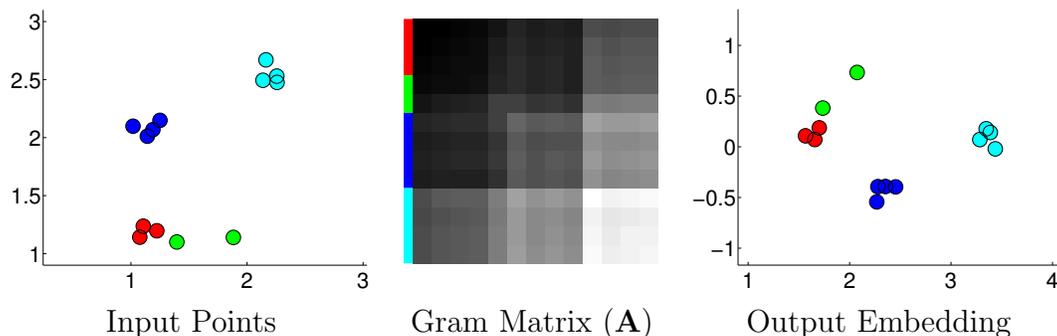


Figure 4-7: MDS example. Rows of the Gram matrix are labeled with the colors of the corresponding points. Embedding was performed using the second and third largest eigenvectors, scaled as in (4.19). This embedding preserves the known input dot products. The colorbar to the left of the matrix indicates the color of the input point and embedding vector corresponding to each matrix row.

Multidimensional scaling (MDS) gives intuition into the meaning of distances in the embedding space (see Figure 4-7 for a toy example). Proposed in the early 1950’s by Torgerson, MDS finds a low-dimensional representation of points whose pairwise

distances are known, with the specific aim of preserving these distances as well as possible [76]. The typical input example is a list of cities with pairwise distances, and one needs to estimate a possible 2D city map layout. The input to the algorithm is a matrix of pairwise distances between the input vectors and/or a Gram matrix of dot products between the input vectors. The conversion formula from dot products to distances is

$$d_{rs}^2 = (\mathbf{r} - \mathbf{s})^T(\mathbf{r} - \mathbf{s}) = \mathbf{r}^T \mathbf{r} + \mathbf{s}^T \mathbf{s} - 2\mathbf{r}^T \mathbf{s} \quad (4.17)$$

where \mathbf{r} and \mathbf{s} are two vectors and d_{rs} is the distance between them [76].

The output of the algorithm is an embedding of the input data points. To perform the embedding, the sum of squared differences between input distances (d_{ij}) and output distances (d_{ij}^*)

$$\sum_{i=2}^n \sum_{j=1}^{i-1} [d_{ij}^2 - (d_{ij}^*)^2] \quad (4.18)$$

is minimized by solving an eigenvalue problem [76]. MDS finds the eigenvectors \mathbf{U} of the matrix of dot products \mathbf{A} , and defines the embedding of the points as

$$\mathbf{E}_{mds} = \mathbf{\Lambda}^{1/2} \mathbf{U} \quad (4.19)$$

where the rows of \mathbf{E}_{mds} are the output embedding [76].

This scaling of \mathbf{U} finds embedding vectors whose dot products are exactly the input dot product values in \mathbf{A} , thus distances are preserved since they are sums of dot products according to (4.17). The relation of the original embedding vectors \mathbf{e} in (4.2) to these vectors \mathbf{e}_{mds} (rows of \mathbf{E}_{mds}) is $\mathbf{e} = \mathbf{\Lambda}^{-1/2} \mathbf{e}_{mds}$. To see that the dot products of the \mathbf{e}_{mds} vectors are those in \mathbf{A} we substitute this into (4.2):

$$\mathbf{A}_{ij} = \mathbf{e}_i^T \mathbf{\Lambda} \mathbf{e}_j = \mathbf{e}_{mds_i}^T \mathbf{\Lambda}^{-1/2} \mathbf{\Lambda} \mathbf{\Lambda}^{-1/2} \mathbf{e}_{mds_j} = \mathbf{e}_{mds_i}^T \mathbf{e}_{mds_j}. \quad (4.20)$$

4.5.2 Hall's r-Dimensional Quadratic Placement Algorithm

This method explains the optimization problem that lies behind spectral embedding using the Laplacian matrix (see Figure 4-8 for a toy embedding example). In 1970,

Hall proposed a solution to the quadratic placement problem, which is finding a layout of n input points which minimizes the weighted sum of squared distances

$$z = \frac{1}{2} \sum_{i=1}^n \sum_{j=1}^n (x_i - x_j)^2 a_{ij} \quad (4.21)$$

where the values a_{ij} are “connections” or affinities between input points and the x values are one-dimensional coordinates [47]. Note that to minimize this equation, points with high affinity will need to be separated by a small distance, i.e. placed near each other. When this minimization problem is put in matrix form, it contains the Laplacian of the graph. The steps are reproduced from [47], adding (4.26) for clarity:

$$z = \frac{1}{2} \sum_i \sum_j (x_i - x_j)^2 a_{ij} \quad (4.22)$$

$$= \frac{1}{2} \sum_i \sum_j (x_i^2 - 2x_i x_j + x_j^2) a_{ij} \quad (4.23)$$

$$= \frac{1}{2} \left(\sum_i x_i^2 a_{i.} - 2 \sum_i \sum_j x_i x_j a_{ij} + \sum_j x_j^2 a_{.j} \right) \quad (4.24)$$

$$= \sum_i x_i^2 d_i - \sum_j \sum_{i \neq j} x_i x_j a_{ij} \quad (4.25)$$

$$= \mathbf{X}^T \mathbf{D} \mathbf{X} - \mathbf{X}^T \mathbf{A} \mathbf{X} \quad (4.26)$$

$$= \mathbf{X}^T \mathbf{L} \mathbf{X}. \quad (4.27)$$

Here \mathbf{X} is a column vector of embedding coordinates (this is for the one-dimensional embedding case), the notation $a_{i.}$ is the sum of the i th row of affinity values while $a_{.j}$ is the sum of the j th column, and \mathbf{L} is the Laplacian of the graph. Equation 4.25 follows because \mathbf{A} is symmetric, so the row and column sums both equal the degree d_i of the node, and the $i \neq j$ is because a_{ii} is defined as zero (this last was in the original paper but is not necessary for the derivation).

Hall introduces the constraint $\mathbf{X}^T \mathbf{X} = 1$ and consequently z is minimized when \mathbf{X} is an eigenvector of \mathbf{L} , specifically the one corresponding to the second-smallest

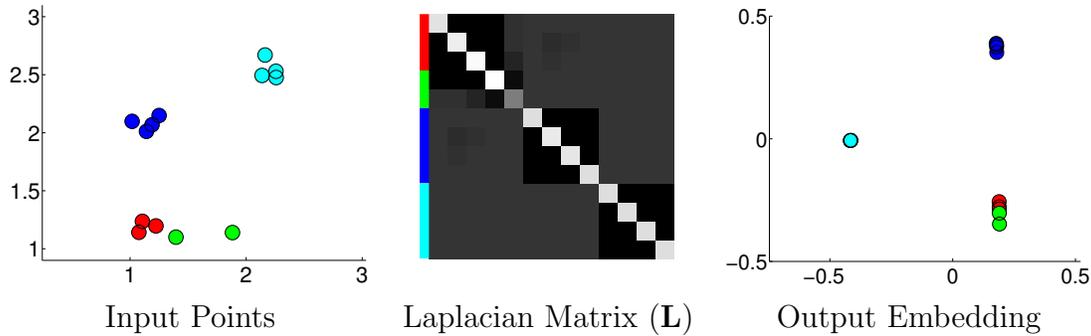


Figure 4-8: Laplacian embedding example. Rows of \mathbf{L} are labeled with the colors of the corresponding points. \mathbf{A} was calculated as $e^{-\frac{d^2}{2\sigma^2}}$ (where d was pairwise distance and σ was 0.3). Embedding was performed using the second and third smallest eigenvectors of \mathbf{L} . The colorbar to the left of the matrix indicates the color of the input point and embedding vector corresponding to each matrix row.

eigenvalue (as the smallest eigenvalue of 0 gives the trivial solution of a constant eigenvector where all points map to the same location). Note that this means that the weighted sum of distances, z , is equal to the eigenvalue, so in this type of spectral embedding the eigenvalues have a clear meaning. Hall uses the rows of the Laplacian's eigenvector matrix directly to perform embedding, without using the eigenvalues:

$$\mathbf{E}_{hall} = \mathbf{U}_L. \quad (4.28)$$

Hall also shows that mappings to higher dimensions can consecutively use the sequence of smallest available eigenvectors, and this satisfies the minimization.

4.5.3 Normalized Cuts

Hall's method was for layout of points; the early application of those ideas to graph partitioning was due to Fiedler who analyzed the second eigenvector of the Laplacian (now called the Fiedler vector) in 1973. Some current spectral clustering methods use the second (and higher) eigenvectors of a normalized Laplacian, where the normalization corresponds to an objective function related to a graph cut. A toy example of Normalized Cuts embedding is shown in Figure 4-9.

The Normalized Cuts objective function for image segmentation using spectral

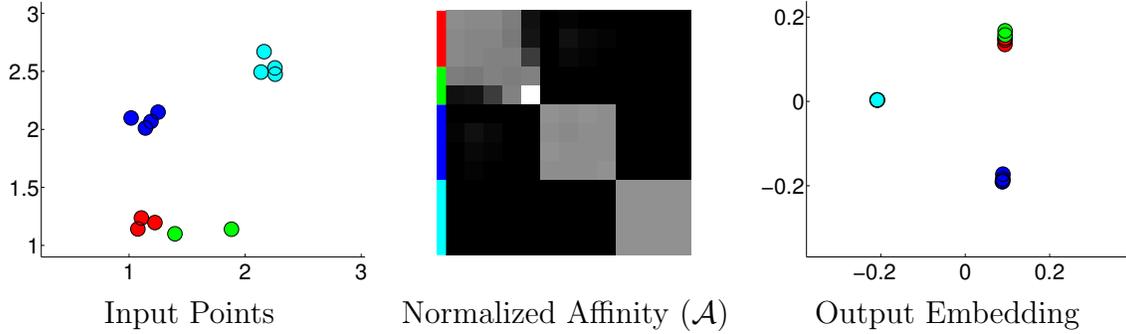


Figure 4-9: Normalized cuts embedding example. Rows of \mathcal{A} are labeled with the colors of the corresponding points. \mathbf{A} was calculated as $e^{-\frac{d^2}{2\sigma^2}}$ (where d was pairwise distance and σ was 0.3), then \mathbf{A} was normalized as in (4.37) to form \mathcal{A} . Embedding was performed using the second and third largest eigenvectors of \mathcal{A} . The colorbar to the left of the matrix indicates the color of the input point and embedding vector corresponding to each matrix row.

clustering was introduced by Shi and Malik in 1997. In Normalized Cuts, instead of posing the problem as an exact minimization, it is expressed as an NP-hard minimization of the normalized cut value:

$$Ncut(A, B) = \frac{cut(A, B)}{assoc(A, V)} + \frac{cut(A, B)}{assoc(B, V)} \quad (4.29)$$

where A and B represent the two sets of nodes in a bipartition of the graph (where the set of all nodes is V), the *cut* is defined as the sum of edges which must be cut to separate A and B , and *assoc*(A, V) is defined as the total connection (edge sum) of nodes in A to all nodes V in the graph [104]. The rationale behind this formula is to improve on straightforward minimization of the cut value (“min cut”), by avoiding its bias for separating small sets of less-connected points [104].

After much algebra [104], minimization of $Ncut$ is found to be equivalent to minimization of

$$\frac{\mathbf{y}^T(\mathbf{D} - \mathbf{A})\mathbf{y}}{\mathbf{y}^T\mathbf{D}\mathbf{y}} \quad (4.30)$$

where \mathbf{y} is an indicator vector for partition membership, \mathbf{D} and \mathbf{A} are the degree and affinity matrices defined previously, and $(\mathbf{D} - \mathbf{A}) = \mathbf{L}$, the Laplacian. When the original constraint of \mathbf{y} being an indicator vector is relaxed, allowing \mathbf{y} to take

on real values, since 4.30 is in the form of the Rayleigh quotient it is minimized by eigenvectors \mathbf{y} of

$$(\mathbf{D} - \mathbf{A})\mathbf{y} = \lambda\mathbf{D}\mathbf{y}. \quad (4.31)$$

So the discrete clustering problem is approximated by a continuous optimization problem. The actual function which is exactly minimized by the Normalized Cut (from [104]) is then quite similar to the function minimized by Hall (4.21), with a new normalization in the denominator. It is

$$\frac{\sum_{i=1}^n \sum_{j=1}^n (y_i - y_j)^2 a_{ij}}{\sum_{i=1}^n y_i^2 d_i} \quad (4.32)$$

subject to $\sum_{i=1}^n y_i d_i = 0$, where y_i is a one-dimensional embedding coordinate, a_{ij} is the affinity value, and d_i is the degree of node i (sum of row i in \mathbf{A}). The normalized Laplacian is calculated in [104] by converting (4.31) to a standard eigensystem using the substitution $\mathbf{z} = \mathbf{D}^{1/2}\mathbf{y}$:

$$(\mathbf{D} - \mathbf{A})\mathbf{D}^{-1/2}\mathbf{z} = \lambda\mathbf{D}\mathbf{D}^{-1/2}\mathbf{z} \quad (4.33)$$

$$(\mathbf{D} - \mathbf{A})\mathbf{D}^{-1/2}\mathbf{z} = \lambda\mathbf{D}^{1/2}\mathbf{z} \quad (4.34)$$

$$\mathbf{D}^{-1/2}(\mathbf{D} - \mathbf{A})\mathbf{D}^{-1/2}\mathbf{z} = \lambda\mathbf{z} \quad (4.35)$$

$$(\mathbf{I} - \mathbf{D}^{-1/2}\mathbf{A}\mathbf{D}^{-1/2})\mathbf{z} = \lambda\mathbf{z}. \quad (4.36)$$

The normalized affinity matrix

$$\mathcal{A} = \mathbf{D}^{-1/2}\mathbf{A}\mathbf{D}^{-1/2} \quad (4.37)$$

has the same eigenvectors as the normalized Laplacian,

$$\mathcal{L} = \mathbf{D}^{-1/2}(\mathbf{D} - \mathbf{A})\mathbf{D}^{-1/2}, \quad (4.38)$$

because the matrices differ by a change of sign and addition of the identity matrix,

operations which do not change eigenvectors. The eigenvalues, however, differ: for the normalized affinity matrix they are one minus the eigenvalues of the normalized Laplacian.

The eigenvectors \mathbf{z} of \mathcal{A} are then used for embedding. In the recursive spectral bipartitioning method, the second eigenvector is thresholded to calculate a bipartition, then this is recursively done for each partition, giving a dendrogram [104]. In this thesis we employ multiple eigenvectors to cluster based on the Normalized Cut as in [40], similar to the multiple eigenvector embedding for spectral clustering discussed in [82] and to the use of multiple eigenvectors proposed by Hall [47]. The embedding formula from [40] is (with modified notation):

$$\mathbf{E}_{ncut} = \mathbf{D}^{-1/2}\mathbf{U}_{\mathcal{A}} \quad (4.39)$$

where the rows of \mathbf{E}_{ncut} are embedding vectors, $\mathbf{U}_{\mathcal{A}}$ is the matrix of eigenvectors of \mathcal{A} , and $\mathbf{D}^{-1/2}$ divides each row of \mathbf{U} by the square root of the row sum of the corresponding row of \mathbf{A} . This computes the same embedding as given by the generalized eigensystem in (4.31) because each eigenvector \mathbf{z} in (4.36) is converted back to \mathbf{y} in (4.31). The uninformative leading eigenvector is not used for embedding (this eigenvector has eigenvalue 1 and is all constant: the minimum cut is no cut at all). In addition, only a subset of the leading eigenvectors is used for embedding: in the case of an ideal block diagonal affinity matrix, the number of informative eigenvectors equals the number of clusters [82]. This corresponds to using columns 2 through k of the matrix \mathbf{U} , where k is the number of clusters/eigenvectors. However in practice when the clusters are not disconnected in the graph, the number of useful (non-noisy) eigenvectors can be significantly less than the desired number of clusters.

After embedding to put similar data points near each other and to enhance cluster structure [17], calculating the actual clusters is usually done using the k-means algorithm [40, 82]. The k-means algorithm is a simple iterative method for finding a local minimum of the sum of squared distances to cluster centroids. This objective

function is

$$\Phi(\text{clusters}, \text{data}) = \sum_{i \in \text{clusters}} \left[\sum_{j \in i\text{thcluster}} (\mathbf{x}_j - \mathbf{c}_i)^T (\mathbf{x}_j - \mathbf{c}_i) \right] \quad (4.40)$$

where the \mathbf{x} values are coordinates of the points and \mathbf{c}_i is the centroid of the i th cluster [39].

4.6 A Note on the Diagonal of \mathbf{A}

In the traditional graph-theoretic description and in one popular spectral clustering paper [82] the definition of \mathbf{A} has specifically stated that $\mathbf{A}_{ii} = 0$. However, in spectral clustering for image segmentation including the Normalized Cuts method [104, 40, 119], the affinity matrix has $\mathbf{A}_{ii} = 1$.

The relationship between $\mathbf{A}_{ii} = 0$ (interpreted as no edge connects a vertex to itself) and $\bar{\mathbf{A}}_{ii} = 1$ (interpreted as self-similarity is maximal) is the following. $\bar{\mathbf{A}} = \mathbf{A} + \mathbf{I}$, and consequently the row sum increments by one: $\bar{\mathbf{D}} = \mathbf{D} + \mathbf{I}$. When embedding is performed using the Laplacian itself, $\mathbf{L} = \bar{\mathbf{D}} - \bar{\mathbf{A}} = (\mathbf{D} + \mathbf{I}) - (\mathbf{A} + \mathbf{I}) = \mathbf{D} - \mathbf{A} = \mathbf{L}$, so there is no change. However when the Laplacian is normalized, $\bar{\mathbf{D}}^{-1/2} \mathbf{L} \bar{\mathbf{D}}^{-1/2} \neq \mathbf{D}^{-1/2} \mathbf{L} \mathbf{D}^{-1/2}$. If the diagonal of the affinity matrix \mathbf{A} is constant and zero, it will be unchanged after normalization, giving constant-length embedding vectors (if the metric \mathbf{A} is used in length measurement). On the other hand, if the diagonal is constant and one, then it will no longer be constant after normalization, corresponding to varying length embedding vectors. We are unaware of any analysis of the effect of the diagonal of \mathbf{A} on clustering results, but we note that if the row sums in \mathbf{D} are constant, the eigenvectors calculated after normalizing \mathbf{A} and $(\mathbf{A} + \mathbf{I})$ will be equal.

4.7 The Nystrom Method for Eigenvector Estimation

The idea behind this method is that an approximation to the eigensystem of a large matrix can be calculated by using a randomly sampled submatrix. For clustering of large datasets this is useful because one doesn't want to (or can't) store a large matrix in memory or find its eigenvectors.

When applied to symmetric matrices, the Nystrom method is equivalent (up to scale factors) to projection onto the eigenvector basis as described at the beginning of this chapter in (4.5) and as performed in kernel PCA [101]. The connection to kernel PCA has been mentioned in two papers that applied the Nystrom method to image segmentation [40] and kernel machine speedup [127]. For nonsymmetric matrices, the SVD can be approximated by a similar projection [37].

In some sense, "Nystrom method" is just a complicated name for a weighted sum. The Nystrom method is so named due to a connection to the Nystrom technique for approximating integrals:

$$\int_a^b y(s)ds = \sum_{j=1}^N w_j y(s_j) \quad (4.41)$$

where the weights w_j come from some quadrature rule (such as the trapezoidal rule or Gaussian quadrature) [96]. The connection to eigenvector approximation is demonstrated by writing the continuous form of equation (4.5) as an integral: the projection of a row of \mathbf{A} onto eigenvectors scaled by eigenvalues becomes a projection onto eigenfunctions [40, 127].

4.7.1 Eigenvector Calculation

Concretely, when applying the Nystrom method, the interpolation weights w_j are the affinity values for one data point (which are weights on the nodes of the graph of \mathbf{A}), and the function $y(s_j)$ being interpolated is the eigenvector (which is a function defined on nodes of the graph of \mathbf{A}). The goal is to estimate the eigenvector at a node which was not present in the graph of \mathbf{A} , based on how similar that node is to nodes

in \mathbf{A} . Following the notation of [40], the entire affinity matrix can be seen as divided into 3 parts, \mathbf{A} , \mathbf{B} , and \mathbf{C} (see Figure 4-10). Affinities are calculated from a random sample of data points to create the square, symmetric \mathbf{A} submatrix. The additional data points are then compared to those that were sampled in order to generate the \mathbf{B} and \mathbf{B}^T submatrices. Finally, the \mathbf{C} submatrix is the largest of all and does not have to be calculated.

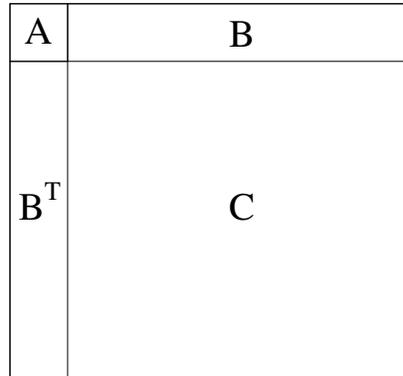


Figure 4-10: Diagram of the parts of the affinity matrix when using the Nystrom method.

Once \mathbf{A} and \mathbf{B} have been obtained, the eigenvectors \mathbf{U} of \mathbf{A} are calculated and used to estimate the eigenvectors $\bar{\mathbf{U}}$ of the full matrix using the following formula from [40]

$$\bar{\mathbf{U}} = \begin{bmatrix} \mathbf{U} \\ \mathbf{B}^T \mathbf{U} \mathbf{\Lambda}^{-1} \end{bmatrix}. \quad (4.42)$$

In equation 4.42 the ordering of the rows of $\bar{\mathbf{U}}$ is such that those corresponding to rows of \mathbf{A} are first, followed by those corresponding to rows from \mathbf{B}^T . (Note $\mathbf{U} = \mathbf{A} \mathbf{U} \mathbf{\Lambda}^{-1} = \mathbf{A}^T \mathbf{U} \mathbf{\Lambda}^{-1}$ has the same form as $\mathbf{B}^T \mathbf{U} \mathbf{\Lambda}^{-1}$ in equation 4.42, showing that the rows of both matrices, \mathbf{A} and \mathbf{B}^T , are projected onto the orthogonal basis $\mathbf{U} \mathbf{\Lambda}^{-1}$. This basis defines the embedding space.)

4.7.2 Matrix Normalization

The row sums of the entire affinity matrix are needed to perform matrix normalization (to create the normalized Laplacian and/or the normalized affinity matrix used in

Normalized Cuts). From [40], the formula for this estimation is:

$$\hat{\mathbf{d}} = \begin{bmatrix} \mathbf{a}_r + \mathbf{b}_r \\ \mathbf{b}_c + \mathbf{B}^T \mathbf{A}^{-1} \mathbf{b}_r \end{bmatrix} \quad (4.43)$$

where \mathbf{a}_r and \mathbf{b}_r are column vectors containing the row sums of \mathbf{A} and \mathbf{B} , and b_c is the column sum of \mathbf{B} . (Of course here \mathbf{A} and \mathbf{B} refer to those matrices before the normalization operation.) In (4.43) the row sums of the uncalculated matrix, \mathbf{C} , are each estimated as a weighted sum of row sums from the known rows of \mathbf{B} .

4.7.3 Performance

The Nystrom method performs well if the entire affinity matrix can be approximated well using the Nystrom eigenvectors (and eigenvalues). The quality of the approximation depends on how well the randomly sampled data points represent the entire problem. Fowlkes and Malik performed empirical performance testing by measuring the repeatability of the Nystrom eigenvectors versus the number of samples. They found that when clustering image pixels, repeatable embeddings were obtained using only one percent of the input pixels to form the matrix \mathbf{A} [40]. In Chapter 8 we perform similar testing for tractography data. In addition, Fowlkes and Malik compared the Nystrom method to the Lanczos sparse eigensolver, which was given a sparsified \mathcal{A} matrix. They found the Nystrom method to be significantly faster.

4.8 Clustering Example, Revisited

We now review the clustering example that was presented at the beginning of the chapter, placing it in the context of Normalized Cuts spectral clustering. Figure 4-11 gives an overview of the entire process.

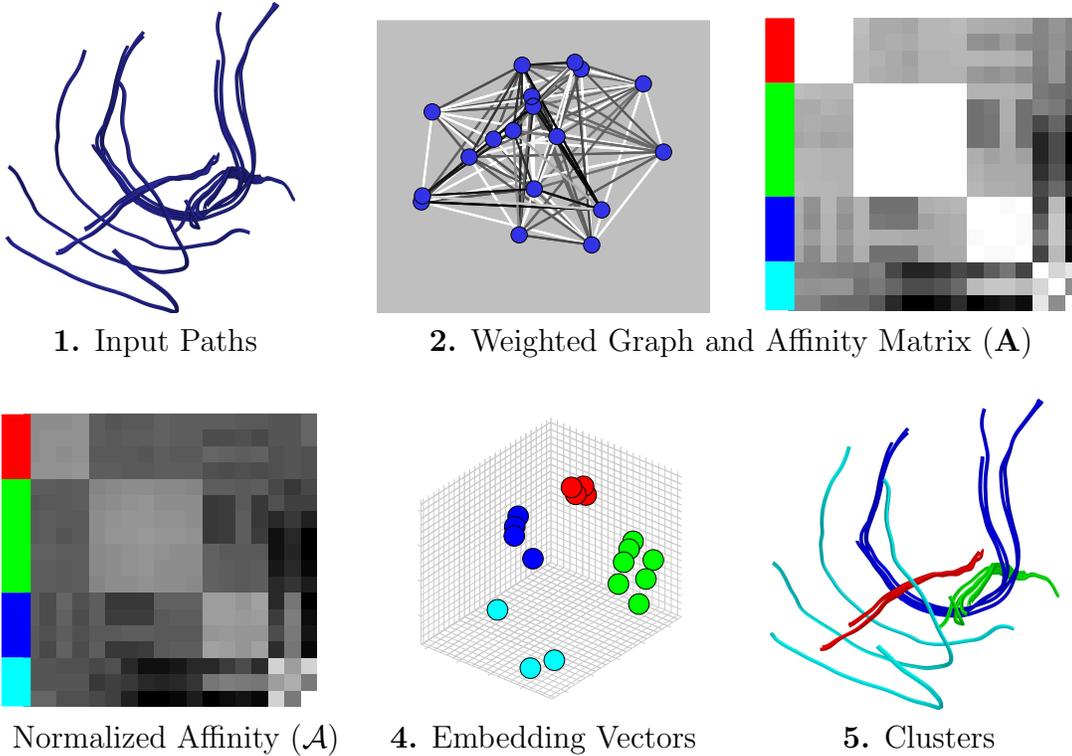


Figure 4-11: Clustering example, revisited. Each step is demonstrated for Normalized Cuts spectral clustering of the tractography example in the upper left. Beginning with input paths (1) from tractography, affinities (2) are computed which can also be thought of as graph edge weights (2). Next the affinity matrix is normalized (3) and its eigenvectors are used to perform spectral embedding (4). Finally, k-means is run in the embedding space, giving output clusters (5). In Chapter 5 we describe how the affinity matrix is computed in our tractography clustering method. This adds another intermediate step (pairwise distance computation between tractographic paths) between steps (1) and (2).

4.9 Summary

We began this chapter with a simple tractography clustering example to give a high-level overview of spectral clustering and to explain how it relates to the problem addressed in this thesis. Then we gave more general background information about spectral embedding and matrices that may be employed for spectral embedding. Next we described the Normalized Cuts and Nystrom techniques which we employ in thesis research. Finally, we returned to the simple tract clustering example to summarize the clustering background material and again relate it to the thesis topic. This leads

us directly to the next chapter which describes our tractography clustering method.

Chapter 5

White Matter Fiber Tract Clustering

The raw output of whole-brain tractography contains many thousands of trajectories which we refer to as tractographic paths. For the purpose of visualization, it is useful to group, or cluster, similar paths in an automatic way so they may be viewed separately from the rest of the data; it is even more useful to augment the clusters with anatomical information to allow display of specific white matter structures. For the purpose of DTI analysis, it is useful to define regions of interest using tractography, in order to measure diffusion properties (for example FA) in subregions of the white matter. However, standard clustering solutions provide no correspondence across subjects, a property which is needed for group analysis is desired for visualization in order to give a similar segmentation and color scheme across all subjects. One proposed method does postprocessing on clusters, attempting to match them across subjects [134]. In this chapter we describe a method for automatic clustering of tractography, for tractography matching across subjects, and for tractography matching across hemispheres¹.

In order to produce cluster correspondences across subjects, rather than a two-step process, we propose integration of the matching and the clustering such that the

¹Some portions of this chapter have been previously published in [87] and a general explanation of the clustering method was published in [86].

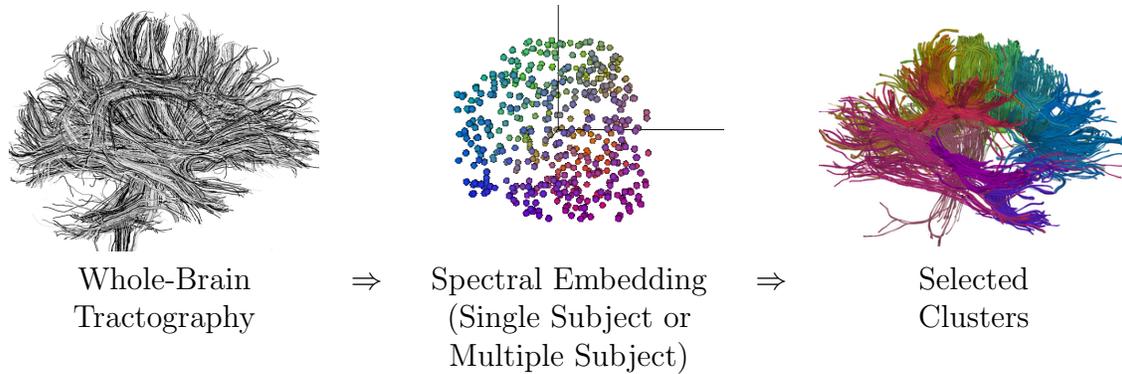


Figure 5-1: The input to our method (left) is whole-brain tractography. Then spectral embedding (center) is performed using the Normalized Cuts method. This produces a representation of each path from tractography as a point in the embedding space. We can embed a single subject or multiple subjects simultaneously (to find matching clusters in all subjects.) Finally, a clustering solution is found in the embedding space, and used to give a segmentation of the input tractography (right). The data shown in this figure is from the population clustering experiment presented in the Results section. The embedding vectors are an actual sample of 500 embedding vectors from the whole population, with the first 3 dimensions displayed. The colors of the clusters (right) are assigned according to the embedding coordinates, as described later in this chapter. In many cases, an anatomical structure will be subdivided into many clusters (for example see the corpus callosum in the image at right). By associating anatomical labels with the clusters, all of the clusters in one structure may be grouped, either for visualization or analysis.

clustering takes advantage of tractography information from all subjects simultaneously, as in Figure 5-1. Our approach to single- and multiple-subject tractography segmentation uses spectral clustering and the Nystrom method. In order to perform spectral clustering in all subjects together, we must get their embedding spaces² to match. Matching embedding spaces is non-trivial: if two brains are clustered separately, there will be no particular relationship between the two embedding spaces, which are defined by the eigenvectors of each brain’s affinity matrix. Furthermore, when using the Nystrom method, the embedding spaces are defined by the eigenvectors of an affinity matrix from a random sample of paths from each brain: thus it is even less likely that a correspondence could be found. However, by clustering multiple brains simultaneously (essentially treating them as one brain and putting all

²The embedding space was explained in Chapter 4; in the embedding space, each path from tractography is represented as a point, and similar paths are near each other.

of the paths from tractography together before performing embedding) they will all use the same embedding space. Clustering in this space can then take advantage of information from multiple subjects, and the clusters that are found in this space will correspond across subjects. By “correspond” we mean that when viewing the clusters in separate subjects, cluster number i in each subject comes from cluster number i in the multiple-brain embedding space, thus all paths in that cluster are be similar to each other, both within and across subjects. Similar to how we “cluster brains together,” we are also able to “cluster hemispheres together” in order to automatically find corresponding clusters across hemispheres.

We introduce three novel ideas in this chapter: population clustering to find a matching tractography segmentation across subjects, automatic matching of tractography across hemispheres, and the introduction of anatomical knowledge in the form of a per-cluster label. The labels allow visualization and analysis to take advantage of the low-level grouping according to cluster, or a higher-level grouping according to anatomical structure.

In this chapter we first present related work on tractography clustering in section 5.1, followed by a description of our method in section 5.2. Section 5.2 is organized according to the six main steps of the method, which are: (1) tractography, (2) generation of pairwise tract distances with optional cross-hemisphere matching, (3) conversion of distances to (normalized) affinity values, (4) spectral embedding and clustering, (5) coloring of paths according to cluster location in embedding space, and (6) postprocessing (optional) to introduce expert knowledge to anatomically label clusters. The description of steps (2) and (3) of the method specifies how we create affinity matrices to perform spectral embedding of tractography as depicted in Figure 5-1. After explaining the method, we demonstrate single-subject and population clustering in section 5.3, showing our ability to automatically discover and match anatomical structures across hemispheres and across subjects. We demonstrate that our method is able to define regions consistent with expected white matter fiber tract anatomy.

5.1 Related Work

Most related work in segmentation of tractographic paths has employed clustering approaches that divide the problem into two parts: the choice of clustering method and the decision of how to quantify distance (or similarity) between paths. Two general types of clustering methods have been employed in the literature, hierarchical clustering [33, 43, 134], and the spectral clustering approach [18, 52, 87, 86].

Several distance measures have been proposed in the literature. In one of the earliest approaches, Ding et al. calculated the mean distance separating paths using pointwise correspondences between path segments [35]. Their method was specific to paths which had been seeded in one image slice, but in later approaches more general distances have been defined. Early work by Brun et al. performed spectral embedding based on distance between path endpoints, then colored paths using their embedding coordinates to give a soft visual perception of connectivity [19]. Later Brun et al. introduced a 9-D tract shape descriptor, which they defined as the mean (3 numbers) and lower triangular part (6 numbers) of the covariance matrix of the points on a path (giving 9 numbers in total) [18]. They computed the distance between paths as the Euclidean distance between these 9-D shape descriptor vectors. Gerig et al. and Corouge et al. proposed distances that do pointwise comparison of tract shapes: they defined three measures related to the Hausdorff distance including the mean closest point distance we employ [43, 33]. Zhang et al. also employed a distance based on average distance between pairs of closest points, where they average all distances that pass a threshold for significant separation of paths [134]. A gray matter atlas was used to initialize clustering in [131], after which the mean closest point and Hausdorff distances were used to group trajectories. In work by Jonasson et al. (who used paths through high angular resolution diffusion data) a path similarity measure was calculated based on the number of times two paths shared the same voxel [52]. There is some convergence in the literature with respect to the choice of distance measure, as multiple authors employ some type of mean closest point distance [33, 35, 43, 86, 134, 131] and it was found to be the most effective in a small

study where the ground truth clusters were known [78].

There are several recent (2005) approaches to matching tractographic paths across subjects. In work by Zhang et al., a two-step process first performed clustering on each subject, then described the clusters with 9-D vectors (the average start point, end point, and “middle point” of all paths in a cluster). These feature vectors were then used to match clusters across patients [134]. Another two-step method was proposed in [61]. In a different approach by Maddah et al., a fiber tract atlas was manually created and used to label multiple subjects, where each new path was labeled as belonging to the atlas structure possessing the most similar path [73].

5.2 Methods

The steps in the method are listed in Table 5.1. In the rest of this section we explain each step. Investigation into reproducibility of the method and its sensitivity to parameter settings is performed in Chapter 8.

Step	Description	Section
1	Whole-brain tractography	5.2.1
2	Pairwise tract distances (optional cross-hemisphere matching)	5.2.2
3	Conversion of distances to affinity values; affinity normalization	5.2.3
4	Spectral embedding and clustering	5.2.4
5	Coloring of paths according to cluster centroid embedding vector	5.2.5
6	Introduction of expert knowledge: anatomical cluster labels	5.2.6

Table 5.1: Steps in our clustering method.

5.2.1 Step 1: Tractography for Clustering

The first step in our method is to generate the input: whole-brain tractography. We use Runge-Kutta order two tractography [95, 3] with a fixed step size of 0.5mm. We have three threshold parameters for tractography: T_{seed} , T_{stop} , and T_{length} . The first

two are anisotropy thresholds based on C_L [123, 124]

$$C_L = \frac{\lambda_1 - \lambda_2}{\lambda_1} \quad (5.1)$$

The goal of the anisotropy thresholds is to limit tractography to the white matter. We seed (initiate) tractography in every voxel in the brain with C_L higher than the threshold T_{seed} . Then tractography stops when C_L on the path falls below T_{stop} , indicating gray matter or planar areas. Example images of C_L thresholds are presented in Figure 5-2. Finally, we employ a length threshold, T_{length} , to remove very short

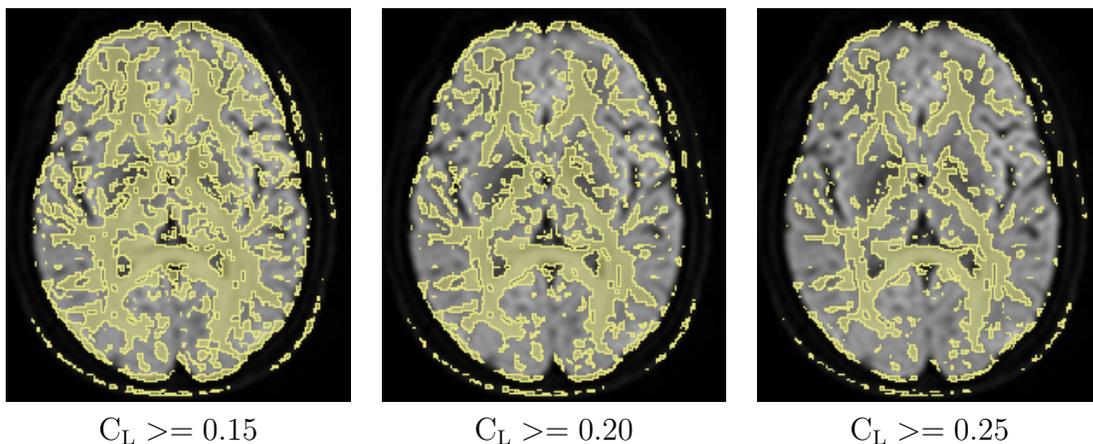


Figure 5-2: C_L thresholds can be used to identify white matter regions for seeding tractography. In each image, the region with C_L greater than the threshold value is shown in yellow. The background image is the average of the diffusion-weighted images used to create the tensor.

paths from the clustering. The thresholds vary by application, generally higher (less inclusive) for group clustering where the major tracts are desired, and lower (more inclusive) for neurosurgical visualization where one would like to see “everything.” Useful ranges for the thresholds are listed in Table 5.2. The settings used for experimental results are specified in section 5.3.

Threshold	Description	Standard Value
T_{seed}	start tractography where $C_L > T_{\text{seed}}$	C_L of 0.15 to 0.3
T_{stop}	stop tractography where $C_L < T_{\text{stop}}$	C_L of 0.15 to 0.2
T_{length}	cull paths whose length $< T_{\text{length}}$	10 to 50mm

Table 5.2: Thresholds used in tractography.

Anisotropy Thresholds for Tractography

In this section we explain why we have chosen to use C_L thresholds for tractography for clustering. In almost all published studies [55, 116, 50], a lower limit (threshold) on fractional anisotropy

$$FA = \frac{1}{\sqrt{2}} \frac{\sqrt{(\lambda_1 - \hat{\lambda})^2 + (\lambda_2 - \hat{\lambda})^2 + (\lambda_3 - \hat{\lambda})^2}}{\sqrt{\lambda_1^2 + \lambda_2^2 + \lambda_3^2}} \quad (5.2)$$

is used to decide when tractography should halt. This threshold for FA is generally in the neighborhood of 0.2, meaning that FA must be greater than or equal to 0.2 for tractography to continue tracing a fiber trajectory. However, the FA is not a particularly good indicator of the presence of a single fiber tract, as it can be relatively high in regions of planar anisotropy, which can indicate tract crossings or branchings [38]. Tractography that employs an FA threshold also requires a curvature threshold in order to avoid such ambiguous areas which may cause sharp turns (see Figures 5-3 and 5-4). The curvature threshold generally uses the radius of curvature, calculated as one over the curvature, where the curvature is the rate of change of the unit tangent vectors on the path. The radius of curvature is the radius of the osculating circle at a point on the curve (this is the circle that best approximates the curve; it has the same tangent and curvature as the curve does locally [121]). Paths with a low radius of curvature are turning quickly, thus a threshold on radius of curvature is a lower limit.

For clustering we want to focus on the more reliable regions that may be delineated using tractography, with the goal of avoiding areas of planar anisotropy where tractography may erroneously cross from one structure to another (though avoiding

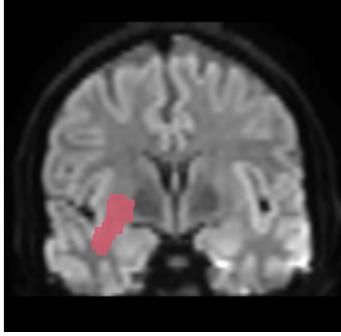


Figure 5-3: One slice through the ROI used to seed tractography in Figure 5-4.

such areas entirely may not be possible). The linear measure (C_L) is a better indicator than FA of the presence of a single fiber tract. If the linear measure is high, it is unlikely that the region is susceptible to eigenvalue ordering problems, and it is more likely that the tensor model is reasonable in that region. A plot of FA vs. C_L is shown in Figure 5-5, as well as images of the areas included with an FA threshold and not a C_L threshold, and vice versa. The C_L threshold, without a curvature constraint, gives comparable (and better for our application) performance to an FA threshold with curvature constraint, with the added benefit of needing only one parameter setting (see Figure 5-4).

Random Sampling of Tractographic Paths

Depending on the voxel resolution of the scan and the various thresholds, whole-brain tractography produces between 10,000 and 100,000 paths per subject. In practice we randomly sample from these paths to obtain a practical number, where approximately 10,000 paths per subject is a reasonable upper limit for visualization (however the total number of paths needed for sampling of tensors and measurement of scalar invariants is under investigation). All random sampling of paths in this thesis was performed without replacement.

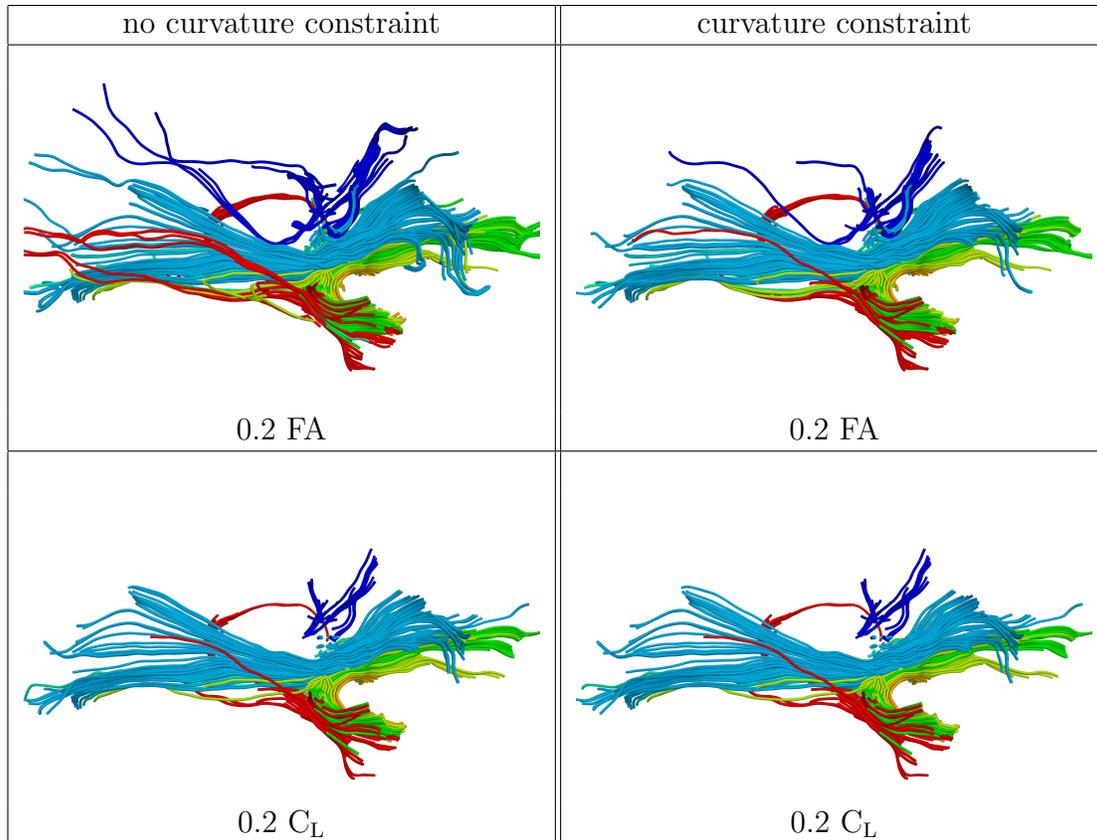


Figure 5-4: Stopping tractography with C_L (bottom) is comparable to using FA (top) and a curvature constraint, and specifically avoids areas with planar anisotropy because C_L is specific to linear anisotropy (cigar-shaped tensors). The two images with a C_L cutoff are almost identical, demonstrating that when using C_L to stop tractography, a curvature threshold is less important and often not necessary.

5.2.2 Step 2: Fiber Trajectory Distance and Matching Across Hemispheres

The second step in the method is the generation of symmetric pairwise tract distances. In this section we describe the distance measure we employ as well as our method for obtaining cluster correspondences across hemispheres. The assumption behind all existing tractography clustering methods is that paths which begin near each other, which follow similar trajectories, and which terminate near each other should belong to the same anatomical structure. This is the best assumption possible without additional information, though since tractography (and axons as well) enter

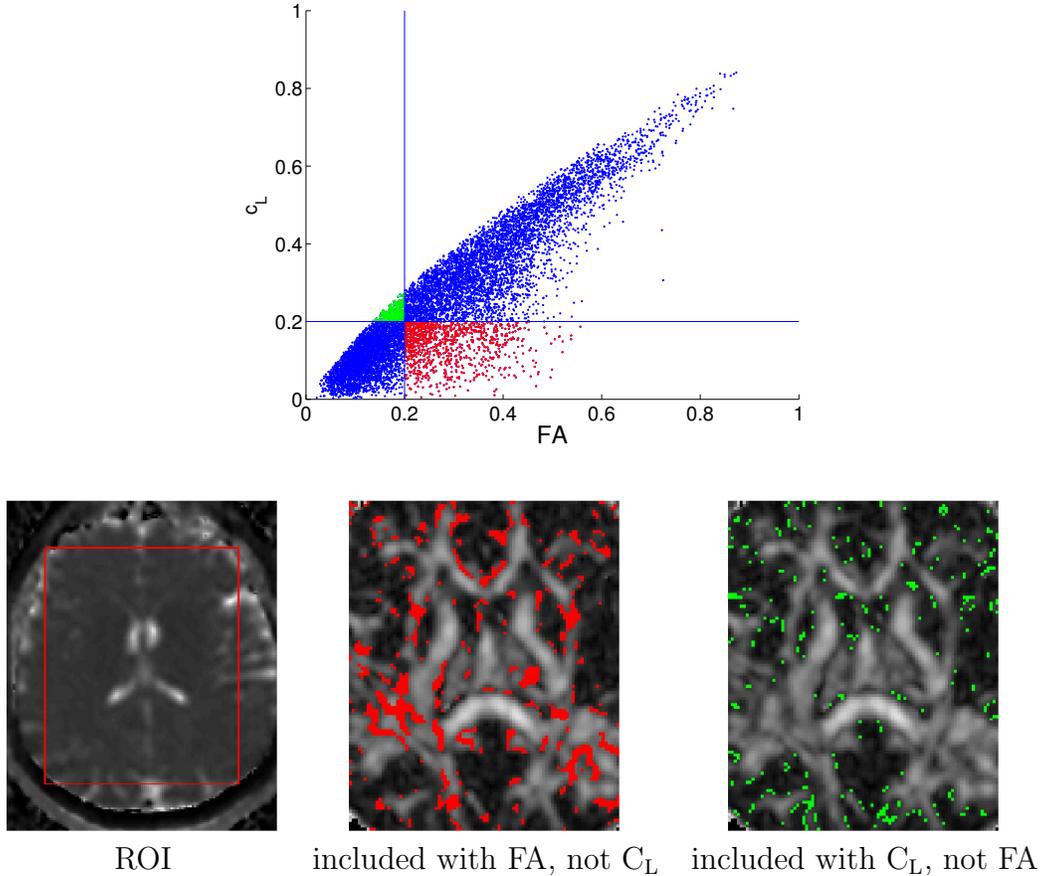


Figure 5-5: Comparison of FA and C_L thresholds of 0.2. Plot of C_L vs. FA in a region of interest (top). The region of interest (lower left), and voxels included by one anisotropy threshold but not the other (lower middle and right, in color over the FA image). Red and green voxels correspond to plot colors (top). Many borders between tracts (“bad” regions for tractography) are in red, i.e. only included with FA.

and leave some structures (the cingulum for instance, as illustrated in Figure 2-6) this assumption is not perfect.

Distance Measure

As summarized in section 5.1, various distance measures have been proposed in the literature to quantify this assumption [19, 18, 43, 33, 52, 134]. We employ the mean closest point distance,

$$d_{\text{mcp}_{ij}} = \frac{1}{n} \sum_{k=1}^n d_k, \quad (5.3)$$

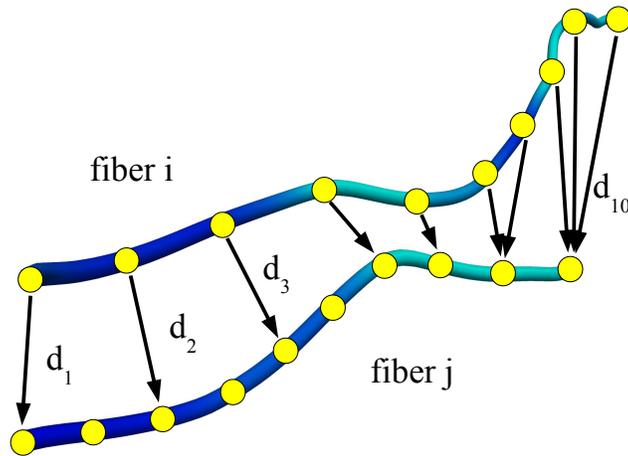


Figure 5-6: Illustration of computation of mean closest point distance measure between two fibers. In this example each path is represented by 10 yellow points. The directed closest point distances (from fiber i to j) are represented with black arrows.

where n is the number of points on path i , and d_k is the distance between point k on path i and the closest point on path j . This distance is just the average distance between pairs of closest points on the two paths. Note that this is a directed distance so $d_{mcp_{ij}} \neq d_{mcp_{ji}}$. Figure 5-6 illustrates the computation of this distance measure.

This distance measure is presently the most successful in the tractography clustering literature, as it is currently in use by three separate groups [43, 33, 86, 134] and was found to be better than others in an evaluation of tractography clustering methods [78]. We have implemented this distance as well as the distances suggested by Brun et al. and while Brun’s measures are more rapid to compute, the mean closest point performs better on whole-brain clustering problems, in that it produces clusters which are qualitatively most similar to the desired anatomical regions. The mean closest point measure has also been applied to trajectory clustering in outdoor video [117]. Practically, it is not necessary to use all points on each path to compute the measure: we currently estimate it using 15 equidistant points along the path, including both endpoints. The path lengths range from around 10mm (20 points) to 150mm (300 points) so using only 15 points per path significantly decreases computation time.

We symmetrize the mean closest point distance by taking the minimum of the two

possible distances $d_{\text{mcp}_{ij}}$ and $d_{\text{mcp}_{ji}}$

$$d_{ij} = \min(d_{\text{mcp}_{ij}}, d_{\text{mcp}_{ji}}). \quad (5.4)$$

We have also employed the mean of the two distances, and we find that the minimum performs qualitatively better when clustering using bilateral matching (Section 5.2.2), while the mean is generally better when doing standard clustering. The minimum encourages the grouping of shorter paths with longer paths, if they run parallel for some distance; this is beneficial for clustering similar anatomy in both hemispheres.

Similarity Across Hemispheres (Anatomical Matching)

A modification to the distance computation allows us to consider symmetry across hemispheres. When reading brain MRI images, doctors rely on the presence or absence of symmetry across hemispheres in order to evaluate whether the scan is normal; consequently providing a symmetric clustering or cluster coloring is an important goal³. Cluster correspondence across hemispheres facilitates both visual and quantitative comparison of anatomical structures that are present bilaterally. A soft correspondence is obtained if cluster colors (based on spectral embedding coordinates [19]) can be approximately matched, while a hard correspondence is obtained where clusters actually contain paths from both hemispheres. Examples of soft and hard correspondences are given in Figures 5-7 and 5-8.

To achieve these goals, before computing the similarity metric we can reflect one side of the brain across the midsagittal plane⁴, such that paths with similar shapes and locations in either hemisphere can cluster together, automatically giving anatomical correspondences. This is a new approach in clustering of tractography. We find that this method produces better separation of some anatomical structures, for example the inferior parts of the cingulum from the inferior parts of the fornix. We believe the

³Thanks to Dr. Alexandra Golby for explaining the importance of bilateral symmetry and motivating research into automatic cluster correspondences. Chapter 7 further addresses this concept.

⁴The midsagittal plane can be defined by AC-PC alignment of the brain images, and translation of the AC-PC plane to the origin, so that taking the absolute value of the R (right-left axis) coordinate performs reflection across the midsagittal plane.

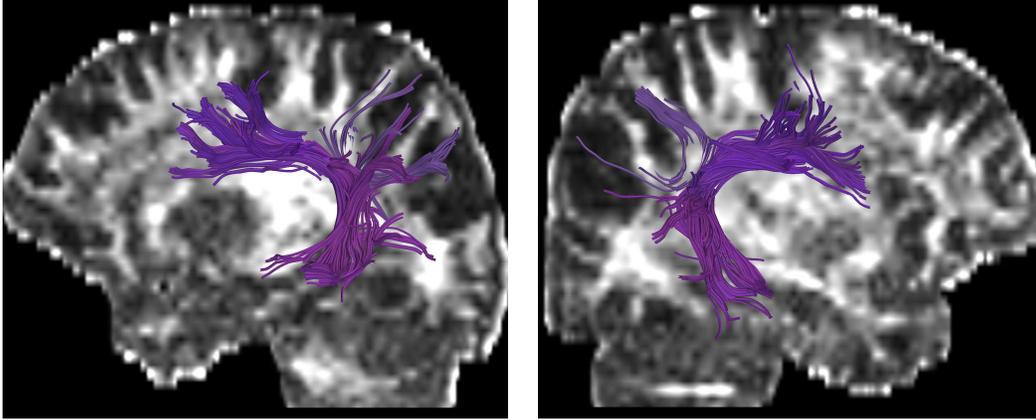


Figure 5-7: Soft visual color correspondences: many clusters in the arcuate fasciculus region are colored in shades of purple. An FA image medial to the clusters is shown in the background. (Dataset from population II.)

improvement in clustering is because reflecting across the midsagittal plane effectively doubles the number of prototype brain examples input to the clustering process.

Note that our clustering method (and atlas creation method described in Chapter 6) are not dependent upon this reflection approach, however the bilateral matching is a useful additional property that we can obtain. Also note that the success of the approach would decrease in subjects with midline shift or other asymmetries, though a more flexible surface-based representation of the midsagittal plane could be utilized. In the case of pathology, the soft visual correspondence based on spectral embedding colors may still be obtained without hard cluster correspondences.

5.2.3 Step 3: Conversion to Affinity

The third step in our method is the conversion of the distance measures to affinity measures suitable for spectral clustering. After computation of pairwise distances, each distance (d_{ij}) is converted to an affinity measure (a_{ij}) via a Gaussian kernel

$$a_{ij} = e^{-d_{ij}^2/\sigma^2}, \quad (5.5)$$

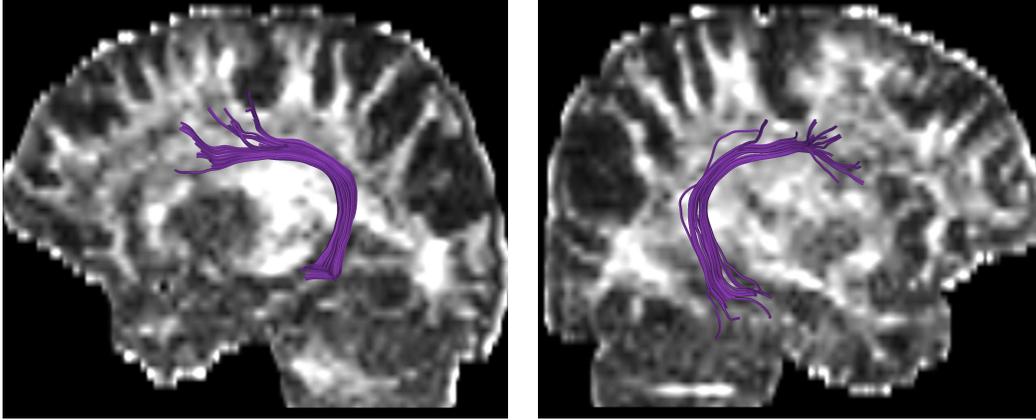


Figure 5-8: Hard cluster correspondence: one c-shaped cluster in the arcuate fasciculus region. An FA image medial to the cluster is shown in the background.

a method that is frequently employed in the clustering literature [82, 104, 119]. Since the distances are symmetrized, this conversion produces a symmetric affinity matrix for clustering. The role of σ in (5.5) is to define the scale of the problem by setting the distance over which paths can be considered similar. We standardly choose σ based on our clustering experience to be 30mm with bilateral matching (and minimum symmetrization), and 60mm in standard clustering (with mean symmetrization). For smaller clustering problems (not the whole brain), values as low as 10mm can give better clusterings. Note that a larger value of σ incorporates more information from anatomically neighboring structures, which is beneficial for whole-brain clustering. A benefit of this affinity measure is that the clustering is insensitive to small registration errors and to small anatomical differences (when clustering multiple subjects) due to the capture range of the mean closest point distance and the Gaussian kernel. Experiments demonstrating that the method is robust to the choice of σ are presented in Chapter 8. The affinity matrix is normalized by dividing each element by the square root of the row and column sum at that location [40].

5.2.4 Step 4: Clustering of Tractography: Single and Multiple Subjects

The fourth step in the method is spectral embedding and clustering based on the pairwise tract affinity values. After calculating (a subset of all) fiber affinities, we employ the method of Normalized Cuts with the Nystrom approximation, as described in [40] and explained in Chapter 4. In the Nystrom method, instead of computing the entire affinity matrix directly, a random sample of paths is selected, and the affinity measure is calculated between each path and that sample. Two matrices are created: the matrix \mathbf{A} contains affinities between paths in the sample, and the matrix \mathbf{B} contains affinities between other paths and the random sample. For a single brain, we find that 1,500 paths are sufficient for the \mathbf{A} matrix (see Chapter 8 for this result).

The matrices are then normalized and spectral embedding of all paths is performed using the eigenvectors of the normalized \mathbf{A} matrix (using the Nystrom method as described in Chapter 4). A user-selected number of clusters is found by running k-means in embedding space. Interesting anatomical clusters exist at many size scales so choosing the number of clusters is not trivial, however the number must be large enough to avoid combining dissimilar paths. Empirically we have found that 100-200 clusters is a reasonable number for whole brain tractography clustering, and experimental results supporting this are presented in Chapter 8. In general more clusters are needed for higher-quality data or if more of the u-fibers (short cortico-cortical connections) have been included due to a lower C_L and length threshold. In addition, we find that approximately 20 eigenvectors are useful for spectral embedding (see Chapter 8 for embedding reproducibility experiments motivating this choice). The parameter settings are specified along with the results presented in this thesis.

Simultaneous Clustering of Multiple Subjects

By simultaneously performing clustering in a group of registered subjects, we find population clusters which represent common structures present in tractography, and actually correspond to expected locations of well-known white matter anatomy. The

layout of the affinity matrix for group clustering, for an example five-subject clustering problem, is illustrated in Figure 5-9. Note that the **A** and **B** matrices contain information from all subjects together. We randomly sample an equal number of paths from each subject to create **A**. In the current implementation, the size of the **A** matrix is limited to the eigenvalue problems that can be solved by matlab, so our present size limit is approximately $4,000 \times 4,000$.

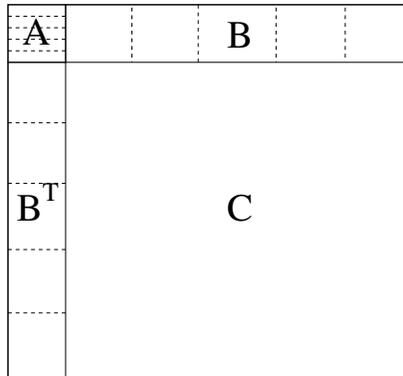


Figure 5-9: Diagram of the parts of the affinity (tract similarity) matrix. This is an example for a five-subject clustering problem. The regions within **A** and **B** represent individual subjects. **C** is the part of the matrix that does not need to be calculated when using the Nystrom method.

5.2.5 Step 5: Cluster Color Assignment

The fifth step in the method is the assignment of cluster colors based on cluster centroid locations in embedding space. As proposed by Brun et al. [19], we scale the chosen embedding coordinates to fit into the interval $[0, 1]$ and use them to determine RGB (red, green, and blue values). The first three embedding coordinates are chosen, corresponding to the second, third, and fourth largest eigenvalues/eigenvectors⁵. We color each cluster according to the location of its cluster centroid in embedding space. An example illustrating the first three dimensions of the embedding, along with the colors and example colored clusters, was shown in Figure 5-1.

⁵As explained in Chapter 4, the first eigenvector is constant since the row sums of the affinity matrix are 1, so this eigenvector is not used for embedding. When embedding using the Laplacian instead of a normalized affinity matrix, the eigenvectors corresponding to the smallest eigenvalues are used. With either matrix, the desired eigenvectors are the “smoothest” ones.

5.2.6 Step 6: Expert Labeling

The sixth and optional step is use of expert information to create a higher-level grouping and anatomical labeling. Since the actual anatomical white matter tracts come in many sizes and shapes (the corpus callosum is much larger than the uncinate fasciculus), we err on the side of oversegmentation of the tractographic paths, with the goal of achieving a correct clustering, i.e. separating all paths which are truly different. Consequently some of the anatomical structures will be subdivided, which in general is a desired property of the method (allowing for example the automatic segmentation of subregions of the corpus callosum as described in Chapter 6).

To group small clusters into larger bundles according to higher-level expert information, each cluster is given an anatomical label. The introduction of a per-cluster label allows data interaction and analysis at multiple scales: visualization and measurement of scalar invariants can be performed at the individual cluster level or the whole anatomical structure level. We discuss in Chapter 6 how this anatomical label information can be applied to a novel subject.

Cluster labeling is performed as follows. After cluster generation, k clusters have been defined in each subject. As explained at the beginning of this chapter, when performing group clustering the clusters correspond across subjects. For example this means that cluster number 10 represents approximately the same region for every subject (because cluster 10 was created as a single cluster in the multiple-subject embedding space). Thus providing higher-level anatomical information is reduced to the problem of defining labels for each of the k clusters. In the labeling process, the tractography clusters are visualized using three dimensional graphics. Single or multiple clusters are selected with the mouse and given an anatomical label. New labels can be created by the user so the naming is very flexible. Theoretically the labeling could be done by inspecting the clustering result for only one subject, but due to anatomical or tractography differences, a small number of clusters are generally empty in each subject. We find the simplest approach is to label one subject, transfer the labels to the next, and work through each subject in this manner, ensuring that

at the end all clusters have a high-level anatomical description. Due to the fact that tractography may cross from one anatomical structure to another, these anatomical labels represent the best approximate description of the regions discovered in group clustering.

5.3 Experimental Results

5.3.1 Anatomical Matching Across Hemispheres

Data and Parameter Settings

The dataset was one subject from population II (for a description of the populations, see Appendix A). The tractography settings were as follows: voxel seeding (initiation of tractography) where $C_L > 0.25$, termination of tractography where $C_L < 0.15$, and minimum path length of 25mm. 10,000 paths were randomly sampled from this tractography dataset and used for clustering. 1,500 paths were sampled to form **A**. Midsagittal reflection was used for distance computation via the mean closest point distance measure. Distances were symmetrized with the minimum operation and converted to affinities using σ of 30mm. 20 eigenvectors were used and $k = 120$ clusters were found.

Results

Figure 5-10 shows the bilateral matching in an example subject (single-brain clustering). To generate this visualization, the paths in each cluster were colored based on the centroid of that cluster in the embedding space. The figure contains the raw whole-brain output with red, green, and blue colors assigned according to the fourth, fifth, and sixth embedding coordinates, respectively. (Normally we employ the second through fourth coordinates but in this case with the midsagittal reflection, these colors had lower contrast so we explored coloring with other coordinates.) Figure 5-11 illustrates a subset of the clusters that are more medial and not visible in Figure 5-10. The selected clusters are from the regions of the cingulum bundles, uncinate

fasciculi, middle cerebellar peduncle, and part of the internal capsule/corona radiata. They demonstrate automatic anatomical correspondence across hemispheres.

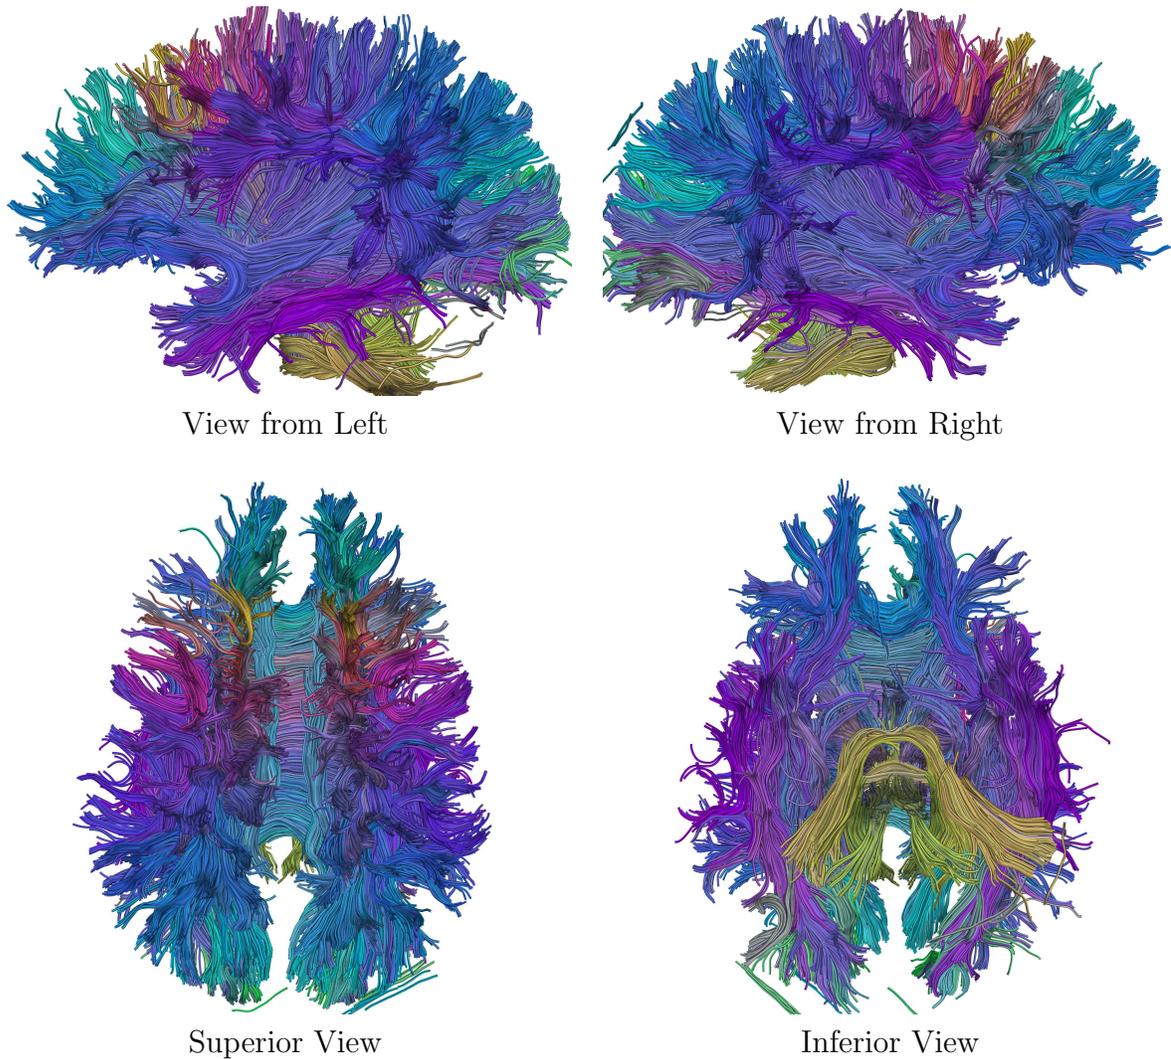


Figure 5-10: Result of clustering with bilateral affinity measure (the colors demonstrate correspondence). This is the whole brain result; see Figure 5-11 for selected clusters.

Discussion of Anatomical Matching Results

The bilateral matching affinity measure introduced in this chapter is dependent on the existence of symmetry in the DTI dataset, and this assumption may break down in EPI (echoplanar imaging) data which often has susceptibility artifacts in the temporal lobe. These artifacts are local magnetic field distortions found especially at

the interface of air and tissue, such as near the sinuses. Since these artifacts may affect the DTI data differently on each side of the brain, we find that with EPI data, we are more likely to compute unilateral clusters in the temporal region than in other regions. However when direct cluster correspondences are not found, the colors from spectral embedding are still similar across hemispheres in the temporal lobe (for example in the occipitotemporal fasciculus). In addition, any clusters that aren't automatically found bilaterally may still be labeled as the same structure when anatomical knowledge is introduced, so the method can recover from this error.

Some outlier paths may be included in our clusters because currently we label everything (all paths which are input to the clustering). For example in the bottom image in Figure 5-11, it is apparent that a single path is present in the cerebellum (seen in the bottom of the image with gold color). This path is an outlier because it has clustered with the other gold paths in the middle cerebellar peduncle due to the fact that very few paths were generated in the cerebellum, so there was no cerebellar cluster. For the purposes of this visualization, removing this path would not be a significant improvement. However for some applications such as the measurement of diffusion values, removal of outliers such as this path could be helpful. Related clustering approaches for vehicle/pedestrian trajectories remove outlier paths a priori (before clustering or embedding) based on their (high) distance to nearest neighbors [117]; however the a priori removal of some of the data is not appropriate for all applications of our method. In addition, we do not compute the pairwise distances of all paths, but this type of nearest-neighbor outlier removal approach generally requires all pairwise distances. We have instead explored removing outliers as a postprocessing step, based on their (high) distance to cluster centroids, and in the future, outlier paths could potentially be marked as such using this approach.

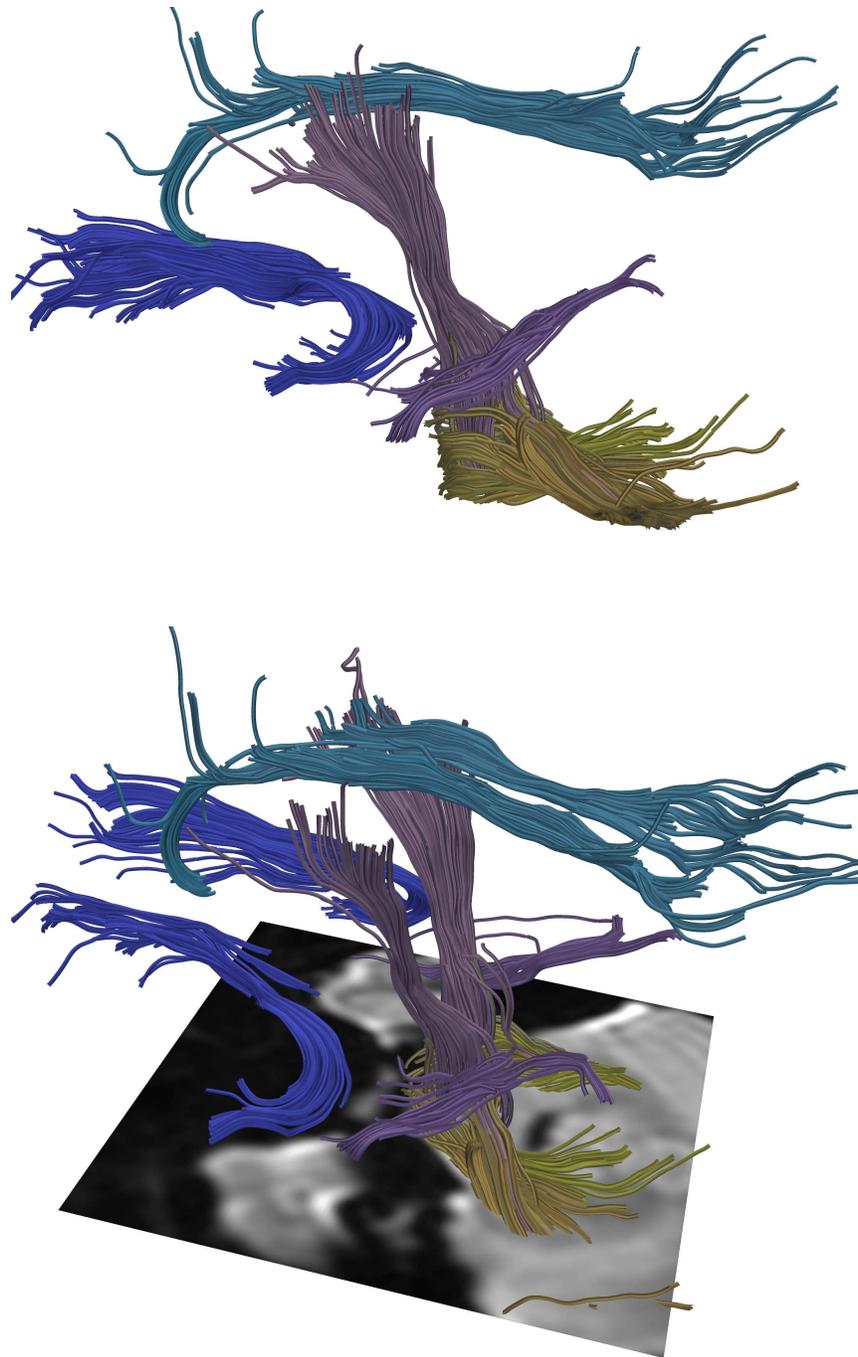


Figure 5-11: Selected clusters demonstrate anatomical matching across hemispheres using bilateral affinity measure (same subject as shown in Figure 5-10). Clusters are located in the following regions: the cingulum bundles (superior and inferior portions) in green and purple, part of the cerebral peduncles/internal capsule/corona radiata in purplish brown, part of the uncinata fasciculi in blue, and the middle cerebellar peduncle in shades of yellow. These are individual clusters found by the algorithm except for the middle cerebellar peduncle, where several clusters are displayed.

5.3.2 Population Clustering

Data and Parameter Settings

We employed the population clustering method to cluster tractographic paths in a population of 5 brains (a subset of population I, described in Appendix A). The tractography settings were: voxel seeding where $FA > 0.3$, termination of tractography where $FA < 0.1$ or curvature > 1.1 , and minimum path length of 30mm. The mean was used for symmetrization of distance values. The total number of tractographic paths from all brains was 36,003. We randomly sampled 1000 tract paths from the population to form **A**. We then performed spectral clustering to generate 100 clusters. These results are from our original population clustering experiment (for more details see [87]) before we refined parameter settings, which is why FA and curvature were used. Since these scans were acquired with AC-PC alignment (see Appendix A) we explored performing clustering without rigid registration.

Results

The output population clusters are shown in each subject in Figure 5-12. The correspondence of the colors directly demonstrates the correspondence in embedding space: the first three components of the cluster centroid embedding vectors were used to define red, green, and blue. In Figure 5-13 we provide more detailed anatomical correspondences in the corpus callosum, cingulum, and uncinate regions. These regions were selected after anatomical labeling of clusters. The cluster indices aren't in themselves significant, but clusters 3, 8, 12, 14, 19, 33, 54, 66, 77, and 81 were found to correspond to the corpus callosum, cluster 31 contained both cingulum bundles, and clusters 29 and 42 represented the uncinate fasciculi. These anatomical regions are displayed in all brains to demonstrate correspondence.

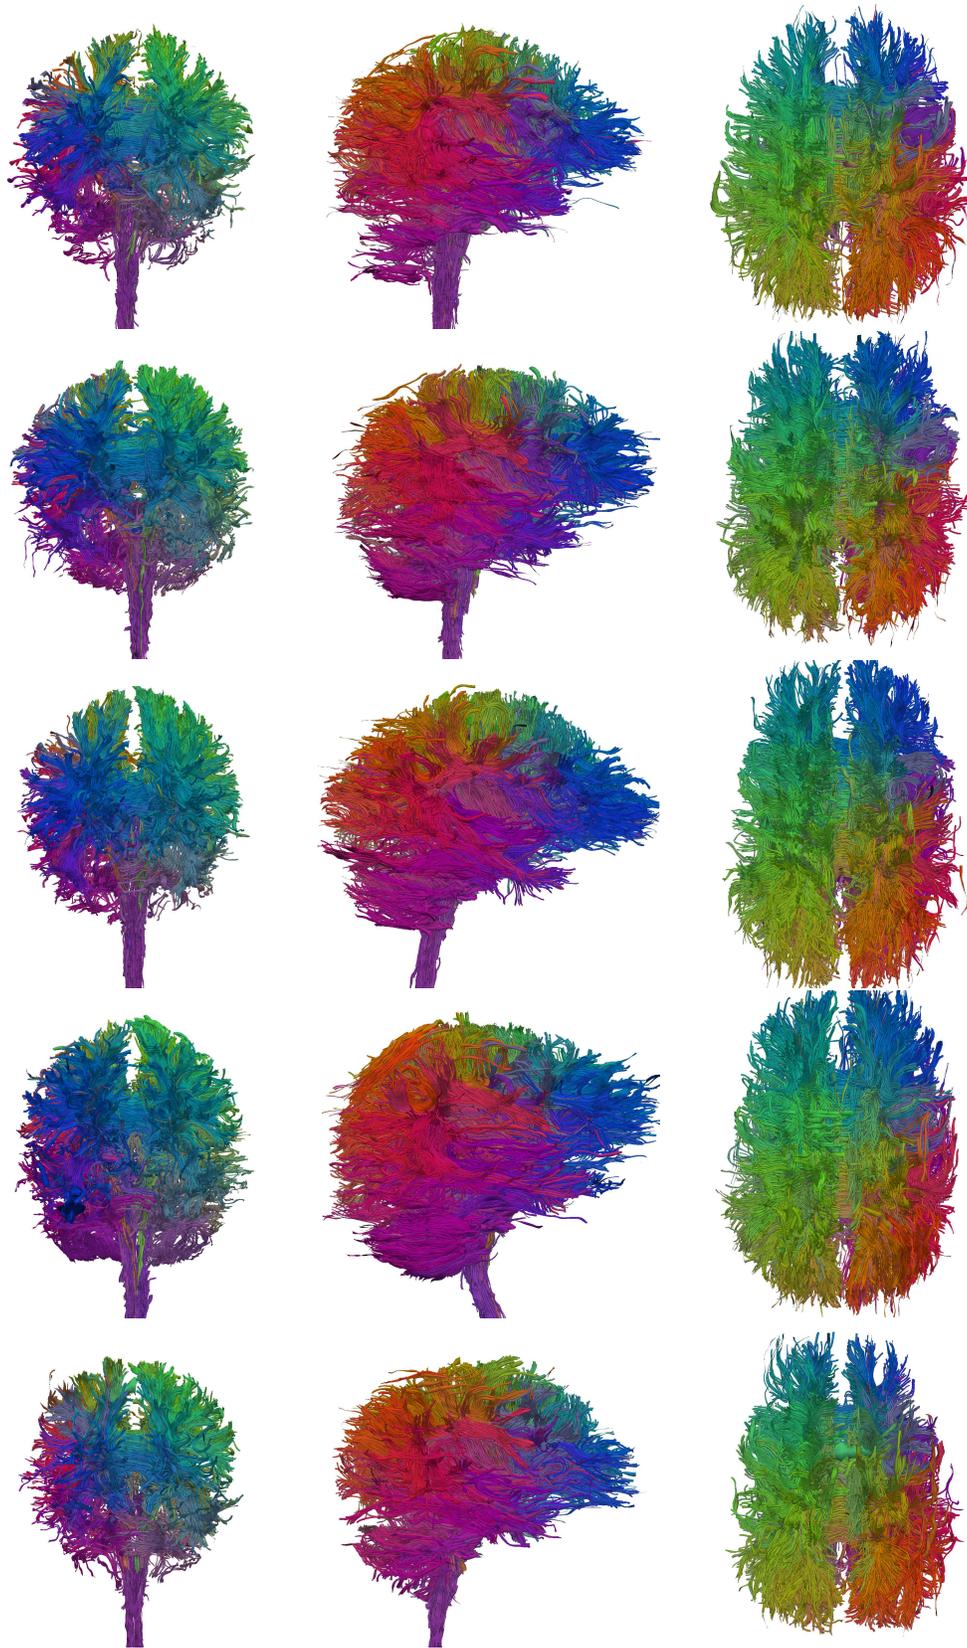


Figure 5-12: Anterior, right, and superior views of the five brains in the population. The colors demonstrate correspondence via spectral embedding. From [87].

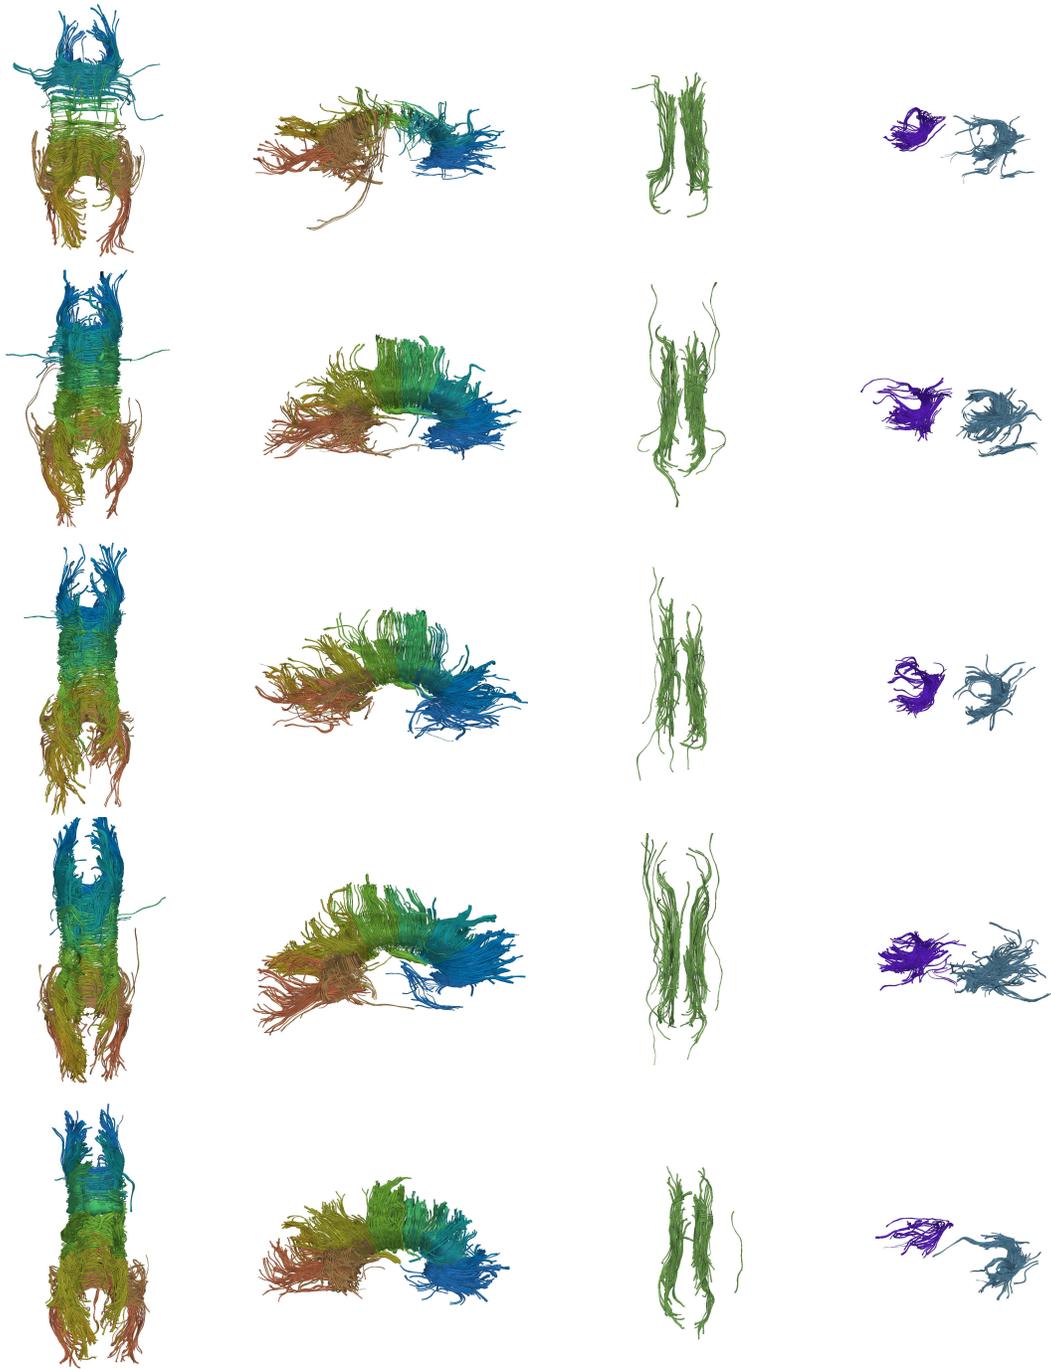


Figure 5-13: Anatomical correspondences after expert labeling: selected clusters, displayed in all 5 brains. The two leftmost images illustrate the corpus callosum clusters, viewed superiorly and from the right. Of the 100 clusters found, 10 were manually labeled as belonging to the corpus callosum. The third images from the left demonstrate a single cluster containing the cingulum bundles, viewed superiorly. Finally, the rightmost images illustrate the two clusters that contain the left and right uncinate fasciculi, viewed anteriorly and from the right. From [87].

Discussion of Population Clustering Results

In Figure 5-12, note that in the fourth brain (from the top), the brain stem is angled very differently from the rest of the brains. This brain was misregistered to the rest due to the rotation about the left-right axis (head tilt). Since these scans were acquired with AC-PC alignment (see Appendix A) we explored performing clustering without rigid registration. The robustness of the method is demonstrated by the overall success of the correspondence. This demonstrates the capture range of the similarity measure in combination with the spectral embedding. Because the affinity measure falls off gradually with distance, when paths from several brains have similar trajectories which are near each other but not perfectly aligned, the paths are able to map near each other in the embedding space. A path in one brain does not have to be similar to corresponding paths in all other brains; if it is similar to paths in some other brains, and those are similar to paths in the rest of the brains, the long-range similarities are automatically discovered by the spectral embedding. This advantage of spectral clustering has been demonstrated in a clustering example in [40], where it was shown that all input data points on the perimeter of a circle could map to the same cluster.

As seen in the body of the corpus callosum of the first subject in Figure 5-13, the output of tractography may vary across subjects when there are differences in the DTI data. Some of the DTI data used for this population clustering experiment was from schizophrenic subjects, and the first subject had visible differences in major eigenvector orientation in several coronal slices in the area where the corpus callosum generally begins to curve superiorly. It can be seen that tractography halted here, and there are fewer paths in the body of the first subject's corpus callosum than in the other subjects (likely because the paths did not pass the length threshold). However, the clustering did manage to group the paths in this region with the corpus callosum clusters. When applying our method to a population of mixed controls and patients, it is important to be aware of possible differences in the DTI data across subjects.

5.4 Discussion of Fiber Tract Clustering Method

5.4.1 Choice of Clustering Algorithm

Our choice of spectral clustering for the application of whole-brain tractography segmentation was originally motivated by the fact that very large datasets may be clustered by using the Nystrom approximation, because pairwise affinities do not need to be computed for all possible pairs of paths from tractography. The Nystrom approximation allows us to cluster a larger number of paths per brain than other published methods [18, 43] and enables population clustering: we are able to cluster many subjects at once to learn a representation of their common anatomy. To give an idea of the numbers of paths involved, there are about $256 \times 256 \times 50$, or 3 million, voxels in a dataset, and seeding (initiation of tractography) in a subset of white matter voxels produces upwards of 10,000 paths per subject. With the Nystrom method, we don't need to store affinity matrices of over 10,000 by 10,000 in memory nor find their eigenvectors directly.

The other clustering method that has been applied to the tractography clustering problem is hierarchical clustering. Agglomerative hierarchical clustering methods iteratively join similar items until some threshold is reached. Such methods don't require precomputation of large affinity matrices because the affinity values can be computed during the clustering operation. However the hierarchical clustering methods in a sense make more local decisions than spectral clustering, because in spectral clustering many affinity values are used simultaneously to compute an embedding. The hierarchical clustering methods lend themselves less well to matching of clusters across subjects because they lack the useful embedding space representation.

An important advantage of our approach is the representation of each path via spectral embedding, which allows us (see Chapter 6) to learn the locations of population tractography clusters in the embedding space, and to apply this information to embed and label novel datasets. The fact that clustering is based only on pairwise affinities is also an advantage for the application, because our choice of pairwise tract affinity measure is relatively robust to slight misregistrations or anatomical differences

when clustering. Other attractive properties of spectral clustering are the underlying minimization of the Normalized Cut values (giving clusters with high within-cluster similarity and low between-cluster similarity) and the inherent noise reduction by using only the top eigenvectors. The downside of spectral clustering (and indeed of many computer vision methods) is the need to investigate multiple parameter settings. However for the application of whole-brain clustering the same clustering parameters work reasonably for all subjects, as demonstrated by the similarity of the tractography segmentations we calculate across subjects (shown in the Results sections of this and later chapters). Chapter 8 of this thesis addresses parameter settings and reproducibility.

5.4.2 Tractography as Input

The main issue with the presented method is that it depends on paths from tractography, a process prone to errors due to noise and partial volume effects, with limited validity in regions of low tensor anisotropy. The alternative approach of working directly in voxel space to cluster or segment tensors by defining a suitable voxel-to-voxel similarity measure is possible, however capturing long range connectivity may be much more difficult. In addition, when specifically clustering voxels it would be difficult to allow more than one tract to share a voxel, unlike our method which allows any paths to be clustered together, regardless of their relationship to voxels. Our method is a reasonable approach to take, despite its reliance on tractography. Although each individual path from tractography may have errors, when taken globally as an entire dataset and segmented in the same way across subjects, the results are strikingly similar across subjects (as seen in this chapter and the next).

Our method is also general: our method can be applied without any modifications to paths produced by another tractography algorithm or paths through data with another diffusion model (such as the ODF model described in chapter 3). With current DTI data, some white matter tracts such as the anterior commissure are not defined well by tractography. Consequently as tractography and diffusion MRI data improve in the future, our method can be applied to improved input tractographic

paths.

5.4.3 Tractography for ROI Definition

The opinion that tractography is useful in defining regions of interest is prevalent in the DTI analysis field, as much recent research has focused on such methods. The current state of the art uses manual interactive methods (as discussed in Chapter 3) to define regions of interest, either by converting trajectories to a voxel representation or by sampling tensors along the trajectories. For example, interactive definition of a (voxel) region of interest using tractography has been shown to be more robust than manual outlining of the same region [93], and several studies including [28, 45, 54, 56] have employed interactive methods for tractography segmentation and used the trajectories for measurement of scalar invariants. Our global population clusters are a useful and relatively robust method for defining regions of interest in the white matter. Our method allows identification of regions in tractography (see the results in this and the next chapters) which are common in an entire population of subjects, as opposed to the manual interactive approach which necessitates detailed three-dimensional neuroanatomical knowledge and analyzes a single subject at a time.

5.4.4 Motivation for Atlas Generation

It makes sense to ask how the population clustering method would scale to a large population of brains. With limited random sampling of paths from each subject, a 64-bit computer with enough memory to hold the affinity matrices, and the willingness to wait for clustering to finish (currently it takes from 45–75 minutes to cluster approximately 30,000 paths and for a large population that would increase⁶), large-scale population clustering could (and should) be attempted. However, in the standard case where one performs a study and may want to add additional subjects as they are

⁶Each iteration of k-means computes the squared distance from each of the n points to all k centroids, where the distance computation uses a number of operations proportional to the number of dimensions, d , of the embedding space. The maximum number of iterations is a constant (a parameter of the method in the implementation we are using), so the running time to calculate a k-means clustering is $O(dnk)$.

scanned, or for the case in which one would like to use population cluster information to label a novel brain, the ability to cluster a large number of subjects would not solve the problem at hand. For these reasons, in the next chapter we propose a method that allows us to take advantage of what we have learned from group clustering and apply that information for novel subjects. The group clustering is used to define an atlas embedding space, and we map tractography from novel subjects into that space in order to automatically segment tractography. This extension of our method makes it suitable for practical application to any DTI dataset, without needing a powerful computer or even performing the slow clustering step. However the true power of the method comes from the ability to associate additional anatomical information with each cluster, and transfer it automatically to a new subject. This allows consistent grouping and visualization across subjects without time-consuming manual interaction, and it addresses the real problem at hand for users of tractography, namely how to find the corpus callosum/uncinate fasciculus/etc. in every subject.

5.5 Summary

In this chapter we presented our framework for single- and multiple-subject tractography clustering, including a method for finding cluster correspondences across hemispheres. We demonstrated multi-subject cluster correspondence results from the application of the method to a small population of five subjects, and bilateral anatomical matching results from two subjects, illustrating matching clusters in the arcuate fasciculus region and in the cingulum, uncinate, and corticospinal regions. In the next chapter these clustering ideas are extended to perform high-dimensional atlas creation, enabling automatic segmentation and anatomical labeling of tractography from novel subjects.

Chapter 6

A White Matter Fiber Tract Atlas

Our automatic atlas-based white matter tractography segmentation method¹ has two parts: high-dimensional atlas generation and automatic tractography segmentation. The method is illustrated in Figure 6-1.

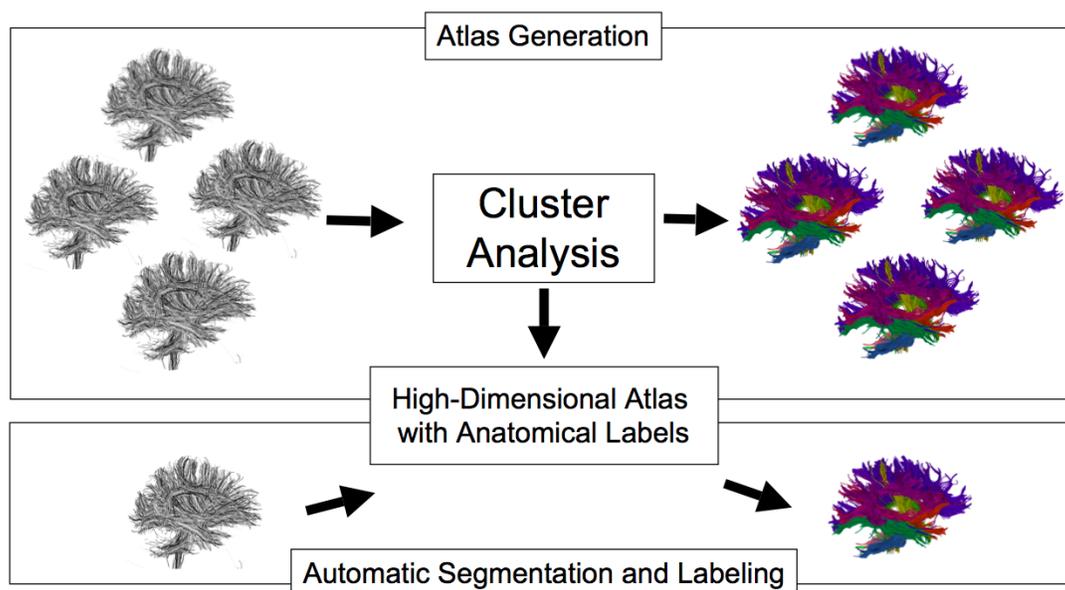


Figure 6-1: Overview of our method. In atlas generation, clustering of tractography from multiple subjects discovers common white matter structures that are anatomically labeled and stored in a high-dimensional atlas. The atlas is then used to automatically segment and label tractography from a new subject.

In this chapter we define a framework both for building white matter atlases and

¹Some portions of this chapter have been previously published in [88] and [89].

for labeling novel brains using such atlases. The atlas generation process automatically finds white matter structures present in a group of tractography datasets and it has a natural extension to allow automatic segmentation of those structures in new datasets, as illustrated in Figure 6-2. The atlas is created by learning a pop-

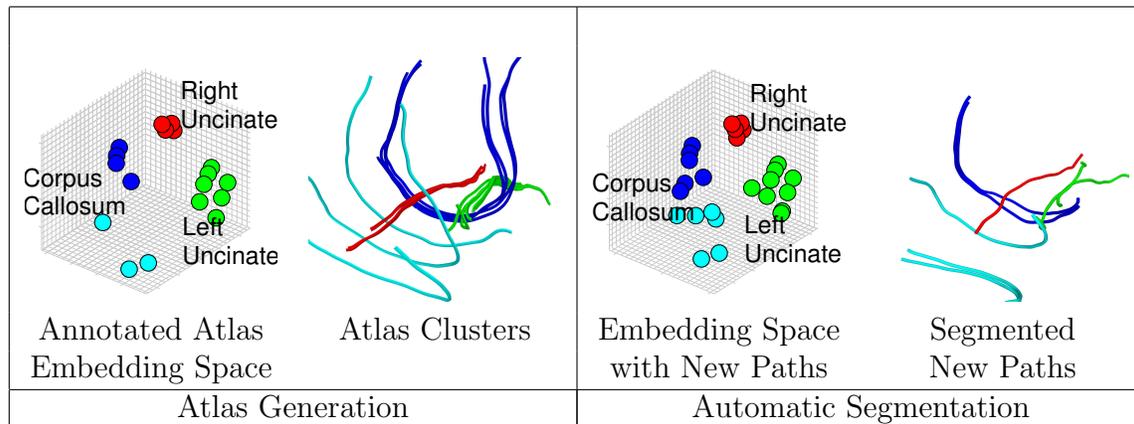


Figure 6-2: Example of atlas generation process (left) and automatic segmentation (right). Normally we create an atlas by clustering multiple subjects, but to illustrate the idea we use the clustering example presented in Chapter 4. Spectral embedding of paths from tractography (left) is performed as described in Chapter 5. (These are the real embeddings of this data, not cartoons.) The embedding space (the first three dimensions are shown here) is then annotated with anatomical information using expert knowledge. Then embedding can be performed for new paths (right). Each new path is represented as a point in the atlas embedding space (we explain how in this chapter). The new paths are automatically segmented and labeled according to the original clusters.

ulation clustering in spectral embedding space, then associating expert anatomical knowledge with each cluster. The stored cluster information is represented in the high-dimensional spectral embedding space described in Chapters 4 and 5. We call our high-dimensional representation of white matter structure an atlas, however it differs from traditional digital (voxel-based) atlases because it represents long-range connections from tractography rather than local voxel-scale information. To automatically segment a novel brain, we embed its paths in the high-dimensional atlas embedding space using the Nystrom method and label the paths by assigning them to the nearest cluster centroid. Our work can be applied both to visualization of fiber paths (performing automatic organization of the thousands of input trajectories to

display) and to population studies of white matter.

In this chapter we first present related work, then we explain the method, showing how to embed novel data to take advantage of the existing Normalized Cuts spectral clustering solution. Finally, we demonstrate the application of the method to two different datasets. In our initial atlas-building experiment we perform population clustering on a five subject dataset, and show that the atlas which is produced can generalize to label a sixth subject. The second experiment uses a larger dataset: we create an atlas using ten brains (where a random sample of paths is used from each) and we use the atlas to segment 10,000 paths each from the ten training subjects and from five additional subjects. In the results section we show anatomical regions that are learned by the method, its application to corpus callosum subdivision², and FA measurement results from the fifteen subjects.

6.1 Related Work

This is the first work to use unsupervised tractography clustering to enable atlas creation. We are aware of three related tractography segmentation approaches which use atlas information, all demonstrated in the past year. Maddah et al. manually created a fiber tract atlas from a single subject and used it to label tractography from multiple subjects [73]. Their further work created an atlas using labeled tractography from several subjects (results were shown in the corpus callosum) and used an EM framework to classify paths from novel subjects [72]. Their atlas contained statistical models of fiber tracts created from expert labeled tractography, represented as the average and standard deviation of paths in fiber bundles, where the paths were represented using spline coefficients [72]. Another approach to clustering of tractography by Xia et al. incorporated information from a gray matter atlas [131].

Related work exists for the general white matter segmentation problem: a rule-based parcellation method has been proposed to label regions of standard anatomical MRI data [74]. Specifically for parcellation of the midsagittal plane of the corpus

²Thanks to Marek Kubicki for motivating research into corpus callosum segmentation.

callosum (one goal which can be achieved with our method), there is a large body of related work including the seven Witelson subdivisions of the midsagittal slice [129] which have inspired a rule-based parcellation into seven segments [74]. Subdivision of the corpus callosum is of interest because many studies have localized inter-subject differences to specific regions of the corpus callosum, for example in callosal atrophy with age [97] and in FA changes in Alzheimer’s disease [99]. DTI-based approaches have mostly aimed to parcellate according to whether corpus callosum tractography goes to temporal, occipital, frontal, or parietal regions [49, 109, 31] though one method employed a statistical model of tractography based on the Witelson subdivisions [72]. Another method parcellated the occipital connections of the corpus callosum based on cortical segmentation from fMRI visual field maps [36]. Unlike our method which finds clusters in an unsupervised way using only DTI tractography, all of these methods employ some type of additional knowledge (rule-based subdivision, cortical segmentation or expert tractography).

In the machine learning literature, using spectral clustering for embedding novel data was proposed by Bengio et al. in 2004 [15]. Bengio et al. describe how to embed out-of-sample (non-training) data using the Nystrom method, for various methods of spectral embedding including Normalized Cuts and MDS [15], but they do not discuss use of clustering results in the embedding space for labeling or segmenting the out-of-sample points. In contrast, our approach segments novel data using spectral embedding and the results of an existing clustering solution, and we augment the learned cluster model with additional anatomical information specific to our segmentation problem. Bengio et al. see the pairwise affinity and matrix normalization steps together as a data-dependent kernel that is learned from the training data and applied to the novel data [15]. Other related work is the Nystrom extension to Normalized Cuts by Fowlkes et al. [40] on which our approach is based. However Fowlkes et al. discuss only eigenvector approximation of a matrix where all data is at hand, not the use of the method for embedding new data. It has also been pointed out that the embedding of new data via the Nystrom method is equivalent up to scale factors to the method of projection onto eigenvectors of kernel PCA [127]. In this chapter we

show specifically how to embed novel data using the framework for Normalized Cuts spectral clustering and matrix normalization of Fowlkes et al. [40], and we show that our tractography data can be segmented according to its distance to existing atlas cluster centroids.

6.2 Methods

Our method includes both high-dimensional atlas generation and segmentation of novel subjects using the atlas. In this section we explain these two parts of the method and describe the contents of the atlas. In addition we show how the segmentation results can be used to measure scalar invariants in corresponding anatomical regions in multiple subjects. We also describe how the tractography clusters may be converted to voxel ROIs (regions of interest).

6.2.1 Atlas Generation

The steps in atlas generation are the same as those explained for population clustering in Chapter 5, including production of whole-brain tractographic paths for all subjects, group clustering, and expert anatomical labeling of clusters. Our high-dimensional white matter atlas contains the following: cluster centroids in the spectral embedding space, information for embedding a novel subject’s paths in this space using the Nystrom method, and anatomical labels for each cluster centroid. The full contents of the atlas are listed in Table 6.1 and will be further explained in this section and the next.

To motivate why this information is needed in the atlas and to provide background information for the next section on labeling novel subjects, we give a brief summary of the Nystrom method here. For more details see Chapter 4 of this thesis and [40, 127]. The Nystrom method approximates the value of each eigenvector for a new data point:

$$\bar{\mathbf{U}} = \begin{bmatrix} \mathbf{U} \\ \mathbf{B}^T \mathbf{U} \mathbf{\Lambda}^{-1} \end{bmatrix} \quad (6.1)$$

Atlas Data	Description
paths	Random sample of paths used for affinity computation (those whose pairwise affinities are in the matrix \mathbf{A}).
D_{param}	Parameters of distance computation (for mean closest point distance this is the method for symmetrization: mean or minimum).
σ	Parameter in conversion of distances to affinities.
$\mathbf{A}^{-1}\mathbf{b}_r$ $\mathbf{a}_r + \mathbf{b}_r$	For estimating row sums of affinities which are not computed. Row sum of \mathbf{A} and \mathbf{B} in atlas affinity matrix (equivalent to column sum of $\begin{bmatrix} \mathbf{A} \\ \mathbf{B}^T \end{bmatrix}$).
$\mathbf{U}\mathbf{\Lambda}^{-1}$	Basis vectors for embedding space.
centroids	Cluster centroids in atlas embedding space.
labels	Anatomical information associated with each centroid.
colors	Colors to use for display of each cluster and anatomical structure.

Table 6.1: Contents of the high-dimensional atlas. Information used for affinity computation is listed first, followed by information used for matrix normalization and spectral embedding, and then by information used for classification and labeling.

where $\bar{\mathbf{U}}$ is the matrix of output eigenvectors, \mathbf{U} and $\mathbf{\Lambda}$ contain the eigenvectors and eigenvalues of \mathbf{A} (the $n \times n$ matrix of pairwise affinities for a randomly sampled subset of n paths out of the total N paths), and the $n \times (N - n)$ matrix \mathbf{B} holds the affinities for the paths to which we want to extend the eigenvector solution. The affinities in \mathbf{B} are calculated by comparing the n sampled paths to all other $(N - n)$ paths, giving an $n \times (N - n)$ matrix. The matrices \mathbf{A} and \mathbf{B} are parts of the whole affinity matrix between all pairs of paths as shown in Figure 6-3.

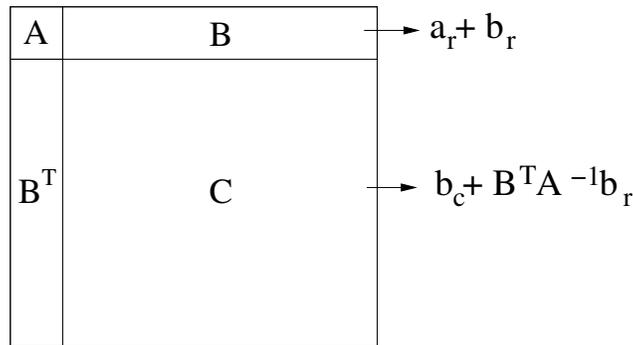


Figure 6-3: Diagram of the parts of the affinity matrix and row sums for atlas construction. \mathbf{A} and \mathbf{B} contain affinities for all subjects and \mathbf{C} is never computed. The row sums of each part of the matrix (those for \mathbf{C} are estimated) are shown at right.

For matrix normalization (needed for Normalized Cuts spectral clustering) the row sums of the entire affinity matrix are estimated from the calculated affinities as

$$\hat{\mathbf{d}} = \begin{bmatrix} \mathbf{a}_r + \mathbf{b}_r \\ \mathbf{b}_c + \mathbf{B}^T \mathbf{A}^{-1} \mathbf{b}_r \end{bmatrix} \quad (6.2)$$

where \mathbf{a}_r and \mathbf{b}_r are vectors containing the row sums of \mathbf{A} and \mathbf{B} , and \mathbf{b}_c is the column sum of \mathbf{B} [40]. $\hat{\mathbf{d}}$ is a column vector containing the sum for each row. (Note that here \mathbf{A} and \mathbf{B} refer to those matrices before the normalization operation.)

To map new paths into the atlas embedding space, the necessary information is the random sample of input paths used in the affinity calculation, the parameter settings for computing distances and affinities, the basis vectors in $\mathbf{U}\mathbf{\Lambda}^{-1}$, the matrix $\mathbf{A}^{-1}\mathbf{b}_r$ for estimating the row sums, and the row (also column) sum $\mathbf{a}_r + \mathbf{b}_r$. Table 6.1 lists the contents of the atlas.

6.2.2 Atlas-Based Automatic Segmentation: Extending Spectral Clustering to New Data

Our framework provides a natural way to transfer the atlas information to a new subject. To apply the atlas to segment a novel subject, each new tractographic path is embedded in the same space in which clustering was performed originally. Then cluster labels and anatomical information are assigned according to the nearest cluster centroid. In this section we describe how the contents of the atlas are used for automatic tractography segmentation in four steps: (1) affinity computation, (2) matrix normalization, (3) spectral embedding, and (4) classification and labeling.

Step 1: Affinity Computation

Figure 6-4 shows the parts of the affinity matrix with a novel subject included. \mathbf{S} holds affinity values for the new subject. To classify a new path, its distance is first measured to each path in the random subset of paths that was chosen during atlas creation. The distances are symmetrized using the same method (mean or minimum)

that was employed in atlas creation. Next, the distances are converted to affinities using the same σ value as in atlas creation. This produces a new row in the affinity matrix for the novel subject (a row of \mathbf{S}^T in Figure 6-4). See Chapter 5 for details on distance and affinity computation.

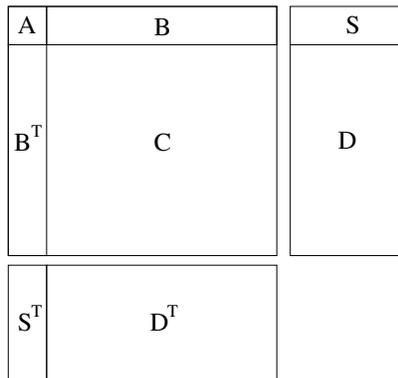


Figure 6-4: Diagram of the parts of the affinity matrix for extension of spectral embedding to new data. \mathbf{S} contains affinities for embedding a new subject and \mathbf{D} is not computed. \mathbf{S} is shown as having a large size to emphasize that we can label a large number of paths per subject with this approach, without the overhead of actually clustering that many paths using k-means in the embedding space. The clustering takes around an hour for a moderate sized atlas but the labeling of a new subject does not need that time-consuming step.

Step 2: Matrix Normalization

To perform embedding using the Normalized Cuts framework, the row must be normalized by division of each element by the square root of the row and column sums at that location [40]. The required row sums are estimates of the row sums of $[\mathbf{S}^T \mathbf{D}^T]$ in Figure 6-4, which can be calculated as

$$\hat{\mathbf{d}}_{row} = \left[\mathbf{s}_c + \mathbf{S}^T \mathbf{A}^{-1} \mathbf{b}_r \right] \quad (6.3)$$

where \mathbf{s}_c is the column sum of \mathbf{S} . For normalization of each row by column sum we employ the column sum from the original matrix,

$$\hat{\mathbf{d}}_{col} = \left[\mathbf{a}_r + \mathbf{b}_r \right]. \quad (6.4)$$

Performing the scaling in this way makes sense for two reasons. First, if we re-embed a path that we have already seen (whose information was in \mathbf{A} or \mathbf{B}) it will be mapped to the same location in the embedding space, because the matrix normalization will be the same. This is why the column sum should not be updated with new information from \mathbf{S} . Secondly, we would expect that each individual new path (row of \mathbf{S}^T) would not significantly change the column sum (6.4) of the entire original affinity matrix, due to the fact that thousands of paths are used in creation of the original atlas affinity matrices. Thus the scaling applied to a novel path (row of \mathbf{S}^T) is basically the same as that which would have been applied if it were part of the original clustering problem.

Step 3: Spectral Embedding

After normalization of the \mathbf{S} matrix, the eigenvectors are estimated using the Nystrom method:

$$\bar{\mathbf{U}}_S = \mathbf{S}^T \mathbf{U} \mathbf{A}^{-1} \quad (6.5)$$

and the Normalized Cuts embedding vectors (4.39) are given as the rows of $\bar{\mathbf{U}}_S$, each divided by the square root of the corresponding row sum from $\hat{\mathbf{d}}_{row}$. This is the same as the embedding process employed for the \mathbf{B} matrix in atlas creation.

Step 4: Classification and Labeling

Finally, automatic segmentation is performed by applying the cluster information from the atlas to the new embedding vectors. The new subject's embedding vectors are labeled according to the nearest cluster centroid, giving a cluster label for each path. In addition, the per-cluster anatomical labels and any additional information (for example the color to use for display of the anatomical structure) are transferred to the novel subject.

6.2.3 Quantitative Measurement of Scalar Invariants

We use the clusters obtained via automatic segmentation as regions of interest for quantitation of scalar invariants of the diffusion tensor field (such as FA and trace). We sample tensors at locations along the paths, then we calculate the invariants. Currently our paths are represented as points with 0.5 mm spacing, so a tensor is sampled every 0.5mm. Any scalar invariant (C_L , FA, mode, trace, eigenvalues, etc.) or the tensors themselves can be measured. Since we have automatic cluster correspondence, statistics may be done to compare the invariants in each cluster across populations. In this chapter we present results showing mean values of invariants in clusters and anatomical regions. In the future, obtaining point correspondences based on alignment of paths from tractography as proposed in [32] will allow statistical analysis of features along white matter tracts.

Related work on tractography-based quantitation of scalar invariants includes recent work on fiber-tract oriented statistics which uses a Riemannian framework for tensor sampling/averaging [32]. This and related methods attempt to average tensors such that the average of two tensors with the same determinant will have that determinant, rather than a possibly higher determinant as with linear componentwise averaging. However, the determinant of the diffusion tensor has not, to our knowledge, been shown to have any biological utility. The trace of the tensor, on the other hand, is the only attribute of the tensor which is in clinical use (for early stroke detection as described in Chapter 3). The trace is known to be nearly constant in all of the white and gray matter [94], thus focusing on the trace when interpolating would seem to be a logical approach. Linear componentwise interpolation of two tensors with the same trace will produce a tensor with that trace. We do not advocate that trilinear componentwise interpolation of tensors is the best interpolation method for sampling from a tensor field (there clearly are more advanced kernels for interpolation, and it is likely that sampling from the original diffusion-weighted images then computing the tensor is the best approach), however the current focus of our work is not on tensor interpolation methods. The trilinear interpolation method is available to us and we

have used it for sampling tensors to produce the measurement results given later in this chapter.

6.2.4 Creation of Voxel ROIs

Similar to the approach used by Mori et al. to color image slices according to tractography [80], we can use our clusters to label voxels. Our current approach is to label each voxel with the ID number of the cluster which has the maximum number of paths that pass through the voxel. Another possible method is fractional labeling, where each voxel would preserve information about the percentage of paths crossing it from each cluster. As an alternative to sampling tensors along paths, the voxels of the original DTI dataset could be labeled in this way according to cluster or anatomical structure, allowing direct use of the original tensors for measurement. In the future we would like to investigate both approaches for calculation of scalar invariants, sampling and voxel measurement. In the Results section we demonstrate use of clusters to create voxel ROIs in the region of the midsagittal plane of the corpus callosum.

6.3 Experimental Results: Atlas 1

6.3.1 Methods

In the first atlas creation experiment, we analyzed DTI data from 6 subjects from population I (see Appendix A). Since the datasets were acquired with AC-PC alignment we explored performing atlas creation without an additional group registration step. We used tractography from subjects 1-5 for atlas creation then we segmented tractography from subject 6 using the automatic atlas-based segmentation process described in this chapter. Tractography was performed in each subject using Runge-Kutta order two integration with the following parameters: seeding threshold of C_L 0.35, stopping threshold of C_L 0.2, step size 0.5mm, and minimum total length 50mm. 5,000 paths from each of subjects 1-5 were used in clustering and 5,000 paths were labeled from the subject 6. For affinity computation we used the mean closest point

distance, σ of 60mm, and mean symmetrization. 1,500 paths were used to create **A**, 20 eigenvectors were used for embedding, and 100 clusters were generated. (Chapter 8 gives results relevant to these parameter choices.) Images were made of the whole brain clustering/labeling and of selected clusters in all brains.

6.3.2 Segmentation Results

Anatomical regions which were found in atlas clustering are shown in Figure 6-5. All 200 clusters are shown in the top rows of images, followed by specific regions in the lower images. These images show automatic cluster correspondence across subjects as well as generalization of the atlas to label the 6th (novel) subject.

6.3.3 Discussion

We have chosen to label novel paths using the nearest centroid, not the nearest embedded point. This allows our atlas to be compact (we only need to store the centroids) though it would make sense to investigate the behavior of the labeling at the borders between clusters. In Chapter 8 we give results regarding clustering stability vs. number of clusters for a single brain.

Unlike Atlas 2, Atlas 1 was created without an additional registration step. The effect can be seen for example in the cingulum region, where four clusters were found. Note the different-colored cingulum bundles in the two rightmost brains in Figure 6-5, which were translated slightly above the rest.

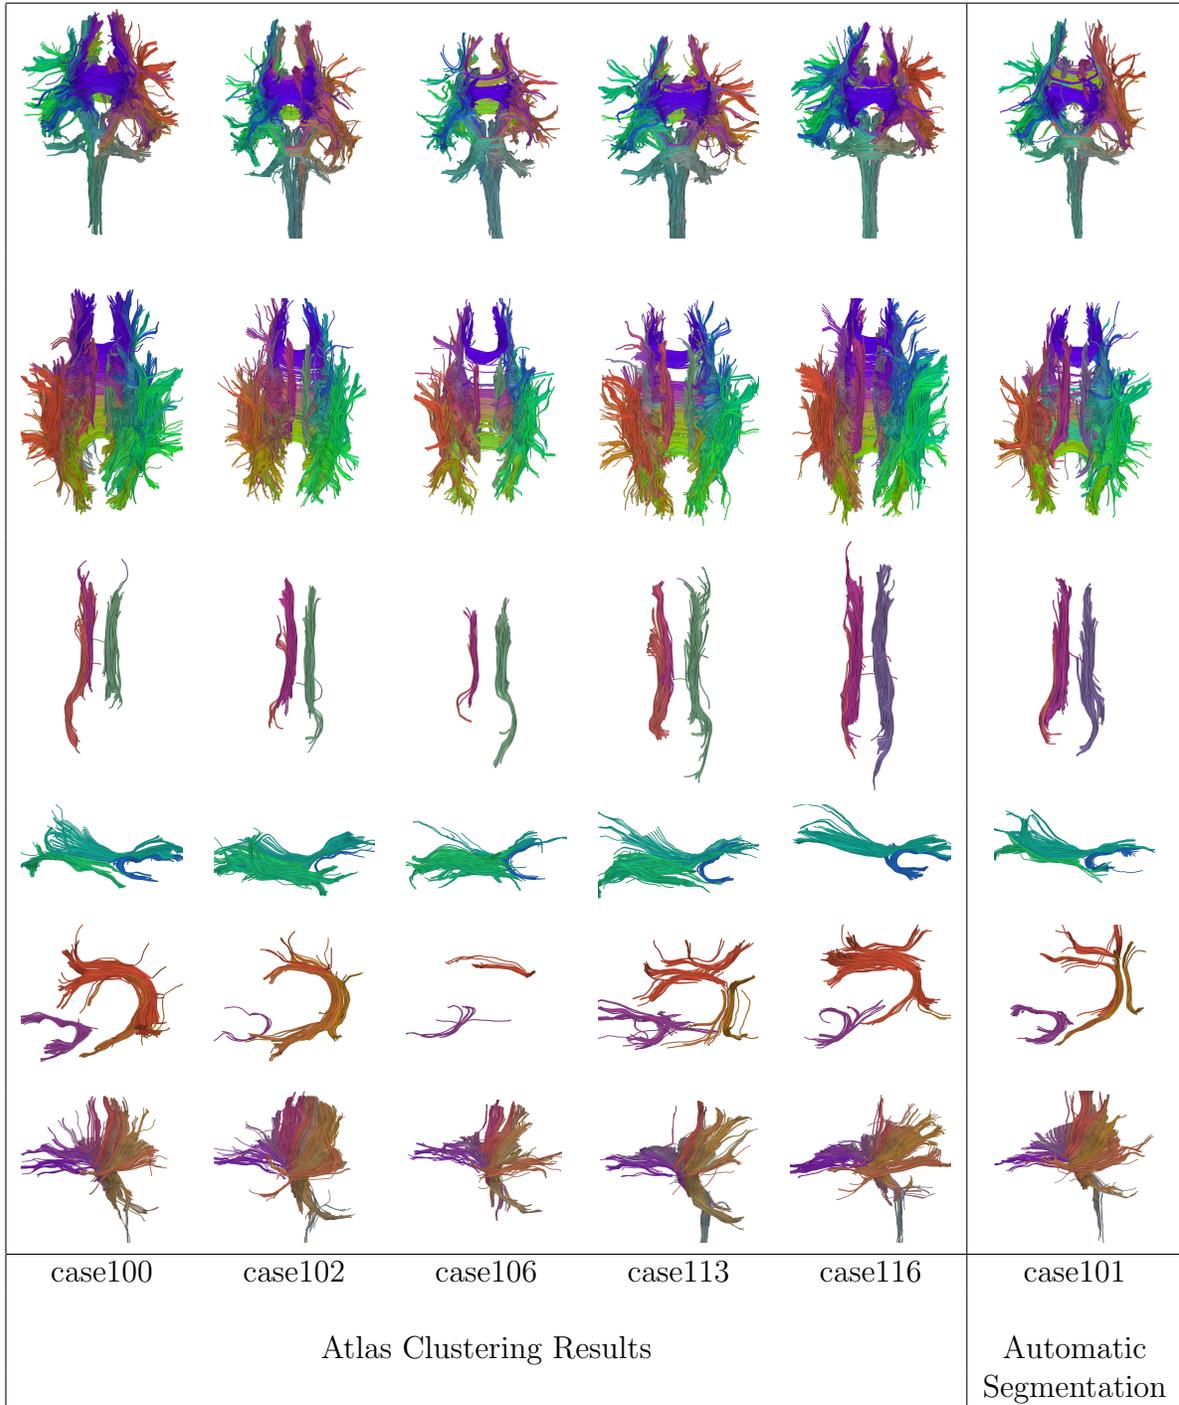


Figure 6-5: Atlas 1: atlas clusters (left) and a novel subject labeled with the atlas (right). The coloring is by the first 3 embedding coordinates. The regions by row, starting from the top, are: whole brain (anterior); whole brain (superior); cingulum bundles, occipitotemporal, occipitofrontal, and uncinate region (from right); arcuate fasciculus/SLF and uncinate (from left); and cerebral peduncle/internal capsule/corona radiata region (from left). Despite the fact that the subjects were slightly misregistered (which affected the clustering, producing for instance several colors in the cingulum region), the atlas can generalize to label the sixth subject.

6.4 Experimental Results: Atlas 2

6.4.1 Methods

In the second atlas creation experiment, we analyzed DTI data from 15 subjects from Population II (see Appendix A). Group registration of subject fractional anisotropy images was performed using the congealing algorithm [135] to calculate rotation, translation, and scaling (no shear terms). This registration was then applied to the paths generated via tractography. To form the atlas, we used tractography from subjects 1 through 10 for clustering and anatomical labeling. Tractography was performed in each subject using Runge-Kutta order two integration with the following parameters: seeding threshold of C_L 0.25, stopping threshold of C_L 0.15, step size 0.5mm, and minimum total length 25mm. We used the midsagittal reflection method for distance computation, where the midsagittal plane was defined using the average group registered FA image. For affinity computation we used the mean closest point distance, σ of 30mm, and minimum symmetrization.

To create the atlas, 3,000 paths were randomly selected from each of the 10 subjects as input to the clustering, giving 30,000 total paths to cluster. The size of the random sample used to create the \mathbf{A} matrix was 2,500, 20 eigenvectors were used for embedding, and $k = 200$ clusters were generated. Finally, the 200 atlas clusters from each subject were visually inspected and labeled with anatomical names. Due to the fact that tractography may cross from one anatomical structure to another, these anatomical labels represented the best approximate description of the regions discovered in group clustering.

We then performed automatic segmentation of tractography from all 15 subjects using the atlas. We labeled 10,000 paths from each of the 15 (10 used in atlas creation plus 5 novel) subjects using the atlas. This 10,000 was a random sample from the total tractography for each subject, which contained between 80,000 and 100,000 paths. (Note that in the 10 training subjects, this random sample was unlikely to contain many of the 3,000 paths which were used for clustering during atlas creation, so performing segmentation of the 10,000 paths was not the same as re-classifying the

original data.)

The 10,000 labeled paths per subject were then used to produce the results shown in this chapter. First, images were created to show selected clusters/anatomical regions of interest for all subjects. Voxel regions of interest (voxel segmentations) were also created for each subject's midsagittal plane, using the clusters in the region of the corpus callosum to perform corpus callosum parcellation or subdivision. Measurements of FA and other scalar invariants were then made within clusters and anatomical regions for each subject.

6.4.2 Segmentation Results

Ten training subjects and five novel subjects were labeled using the atlas, and Figures 6-6 through 6-8 display several views of the whole-brain segmentation result from all subjects. In each figure, the top two rows contain the training data and the bottom row contains the test data (the novel subjects). Note that the labeling is consistent across subjects, and that major structures such as the corpus callosum, arcuate fasciculus, uncinate fasciculus, etc. are found and labeled by the automatic segmentation method. Some clusters are not shown in these images, in particular the u-fibers (small tracts near the cortex).

Figures 6-9 through 6-11 focus on results in specific white matter regions. Figure 6-9 demonstrates the success of the segmentation in the bilateral uncinate fasciculus and arcuate fasciculus regions. Figure 6-10 shows the performance of the method in the limbic system. There are some challenges to the method in the region of the fornix (green) where tractography is more variable (as seen in the figure, trajectories may cross to the corpus callosum, anterior thalamic radiation, or traverse part of the arcuate and terminate near the fornix). Figure 6-11 displays the individual clusters which were found in the cingulum region. This illustrates one of the strengths of the atlas representation: in the atlas we are able to store the expert information that these clusters represent the same anatomical structure.

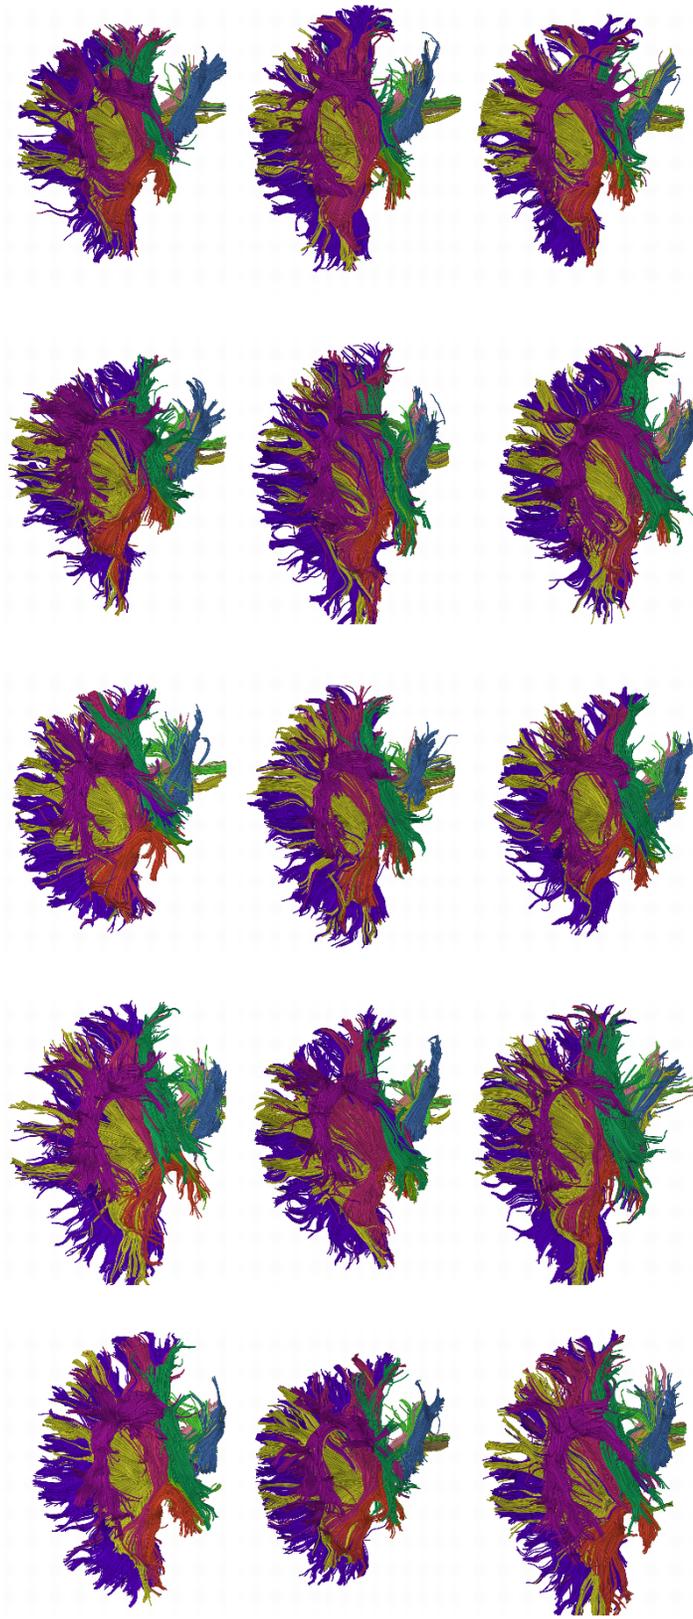


Figure 6-6: Result of automatic segmentation of all subjects (view from left). The arcuate fasciculus region is shown in purple, the corpus callosum in blue, the internal capsule in yellow, the uncinate fasciculus in orange, the middle cerebellar peduncle in light blue, and the inferior occipitofrontal fasciculus in pink.

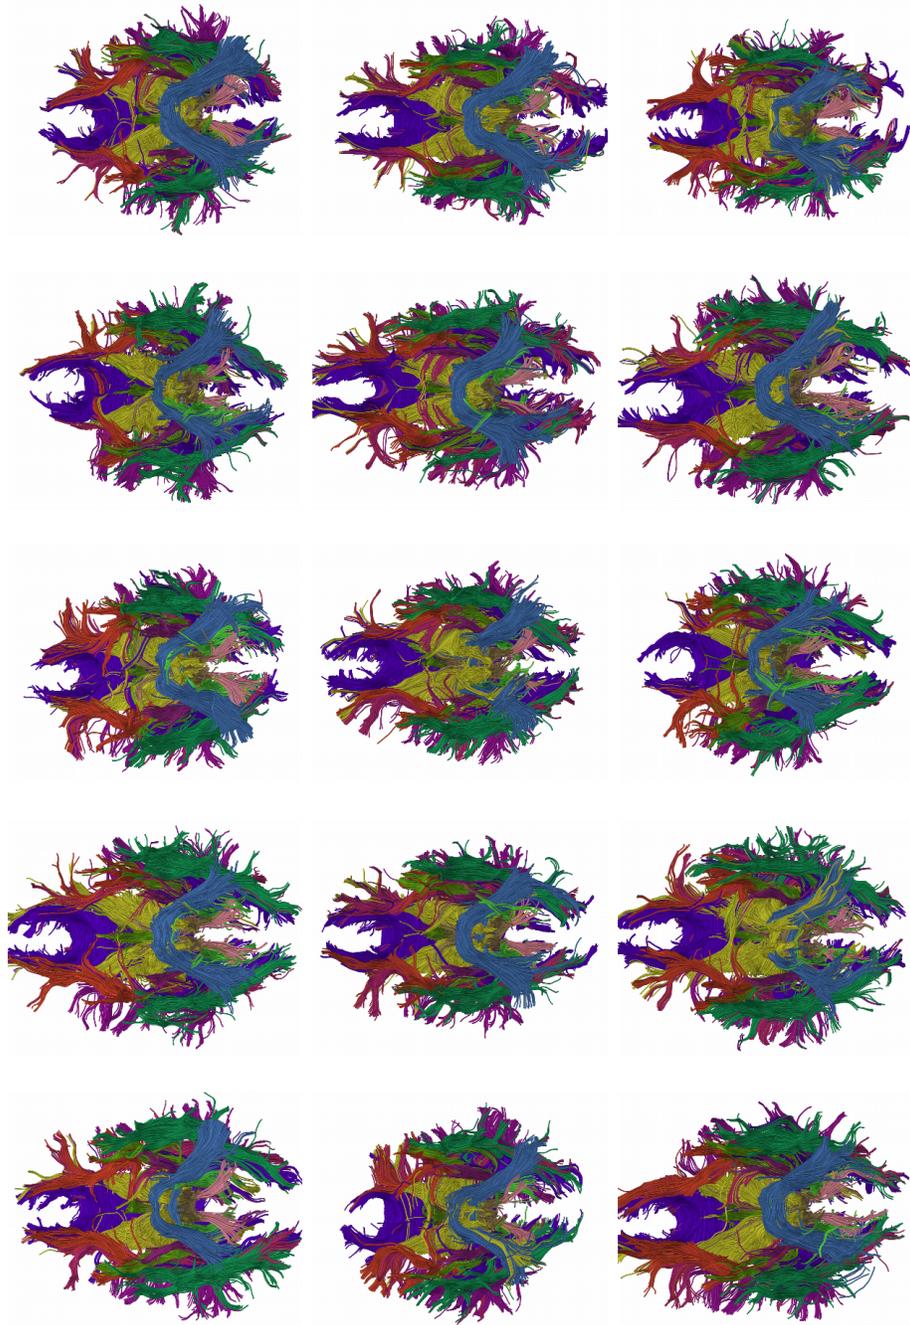


Figure 6-7: Result of automatic segmentation of all subjects (inferior view). The middle cerebellar peduncle is visible in light blue, the inferior longitudinal fasciculus region in green, the corpus callosum in blue, and the superior cerebellar peduncles in light pink.

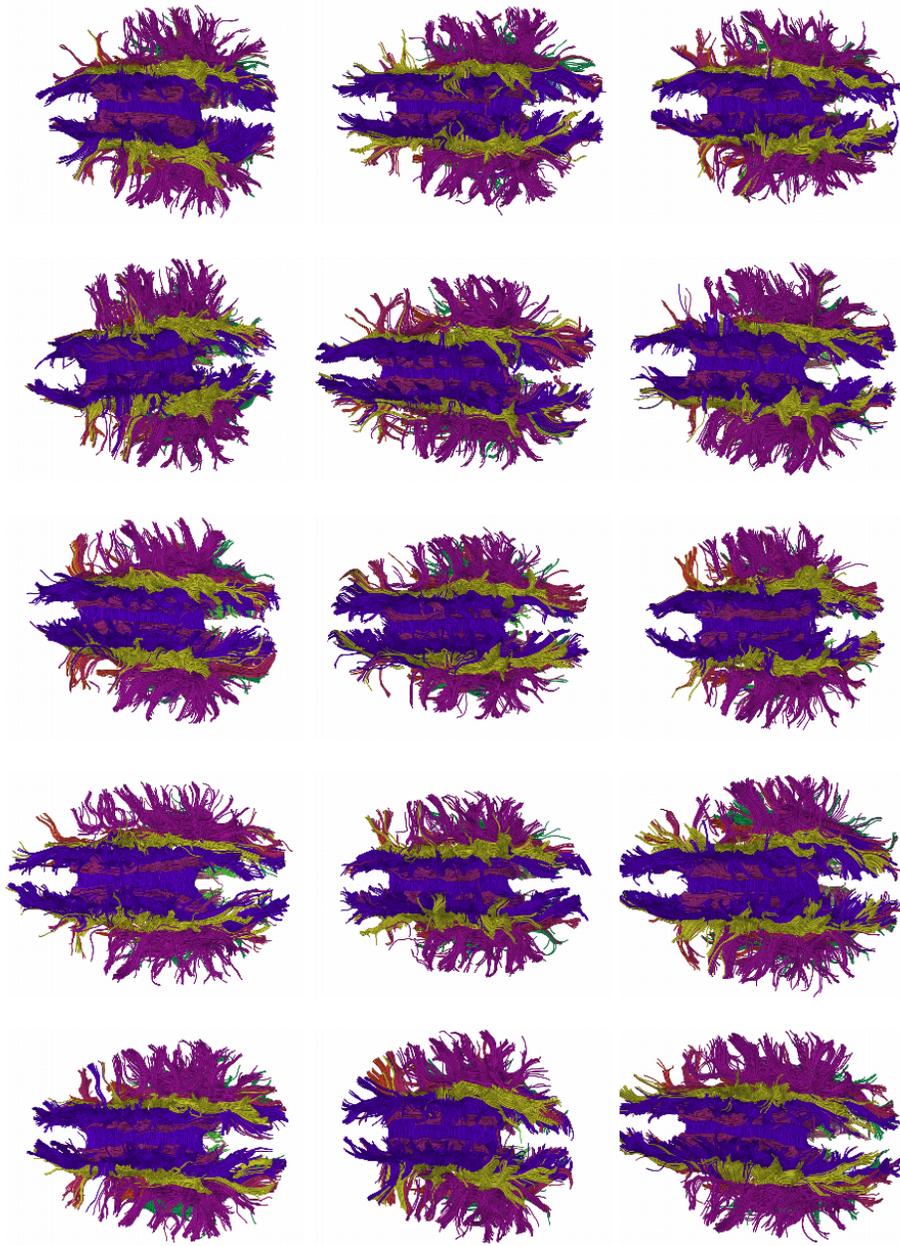


Figure 6-8: Result of automatic segmentation of all subjects (superior view). The corpus callosum is shown in blue, the corona radiata in yellow, and the cingulum bundles in purple (located medially). Clusters in the region of the arcuate fasciculus are also in purple (located laterally, on each side of the images).



Figure 6-9: Result of automatic segmentation of all subjects (view from left). The bilateral arcuate fasciculus cluster is shown (purple) as well as clusters in the region of the uncinate fasciculus (orange) and cerebral peduncle/internal capsule/corona radiata region (yellow).



Figure 6-10: Result of automatic segmentation of all subjects (view from left). The cingulum is shown in purple and the fornix in green.

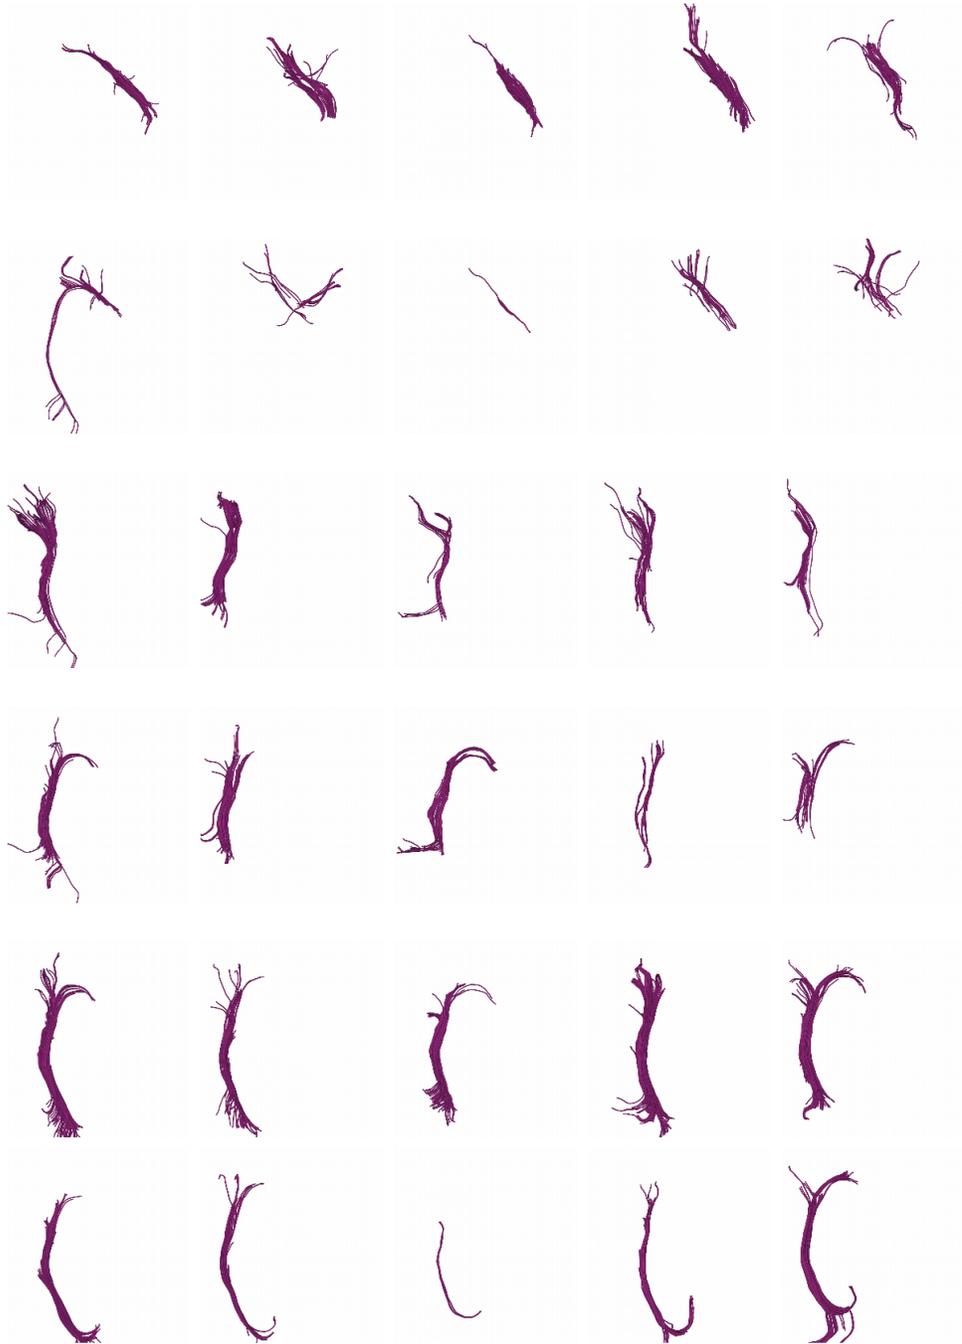


Figure 6-11: Result of automatic segmentation of cingulum clusters in test subjects (view from left). Each row corresponds to one subject, and each column to a cluster. The most anterior clusters are to the left and the posterior and inferior ones are to the right.

6.4.3 Automatic Corpus Callosum Subdivision

The corpus callosum region contained 32 clusters. The clusters were manually grouped into 12 regions (to simplify the display of the results) and assigned unique colors. The results for all subjects are shown in Figure 6-12, and example midsagittal plane voxel segmentations are demonstrated in Figure 6-13.

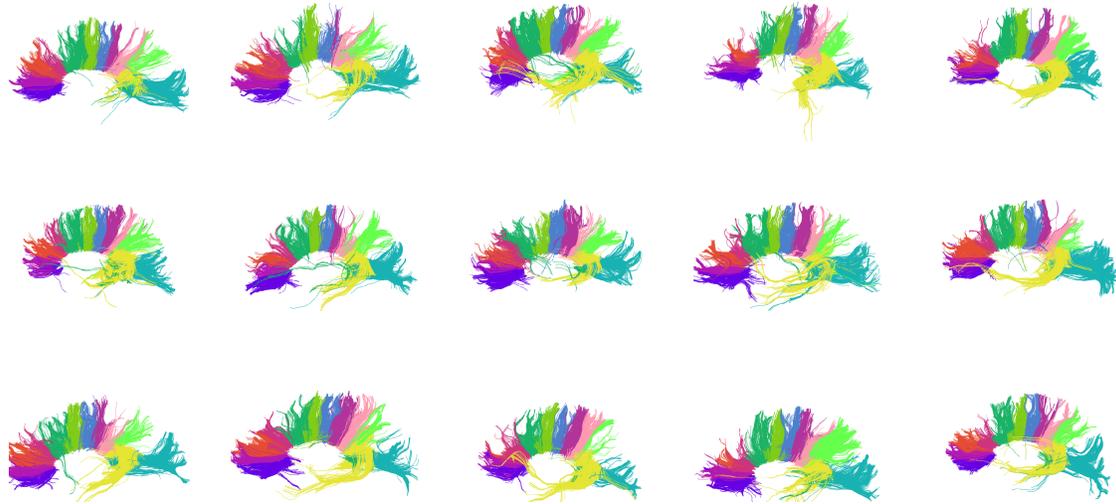


Figure 6-12: Automatic subdivision of corpus callosum in all subjects. Training subjects are in the top two rows, followed by test subjects in the bottom row.



Figure 6-13: Automatic subdivision of corpus callosum, converted to $2 \times 2 \times 2$ mm voxels, in three example subjects.

6.4.4 Measurement of Scalar Invariants in the Population

Figure 6-14 shows the mean FA measured in selected anatomical regions in all subjects. The FA was measured bilaterally (both hemispheres together). The mean FA

is relatively similar across subjects (in each structure) yet differs across structures, as would be expected from anatomically consistent group measurements. We have used bar graphs in order to show the pattern across subjects where some were found to have consistently high or low FA relative to other subjects, as seen clearest in the corpus callosum and “corona radiatae” graphs. Note that for lack of a better term, we have used “corona radiatae” to refer to the large structure including part of the cerebral peduncles, internal capsule, and corona radiata proper.

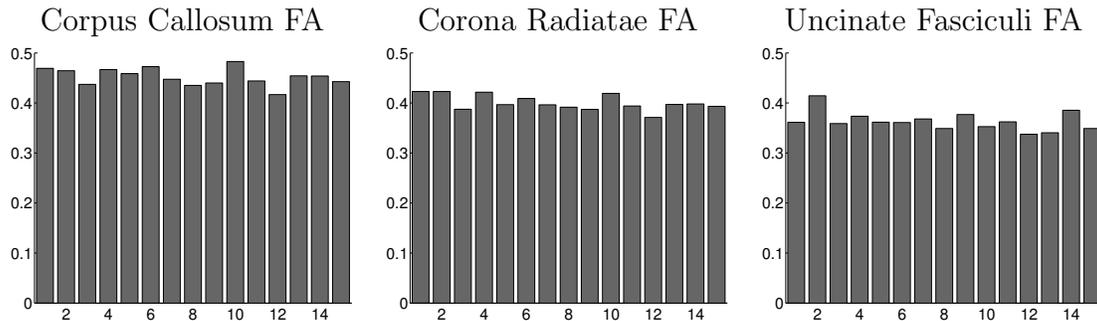


Figure 6-14: Mean FA in three regions for 15 subjects. The bar graph format shows the pattern that some subjects have generally higher or lower FA than others.

Figure 6-15 shows scalar invariants (FA, C_L mode, and trace) measured in each structure for all 15 subjects. The measurements are displayed as dots and each subject’s values are connected with a colored line to show the pattern across subjects.

6.4.5 Discussion

Overall the correspondence of the structures across subjects is very good. This is shown by the images, by the consistency of the FA measurements across subjects and within structure, and by further experiments performed in Chapter 8. However in some regions the tractography is more variable and may cross between structures, as demonstrated in the fornix region in Figure 6-10. Another example is seen in Figure 6-11 in the cingulum cluster that includes trajectories that take the posterior anatomical curve (this is anatomically correct but rarely seen in the tractography, so this cluster result is less compact or similar within cluster). These regions present

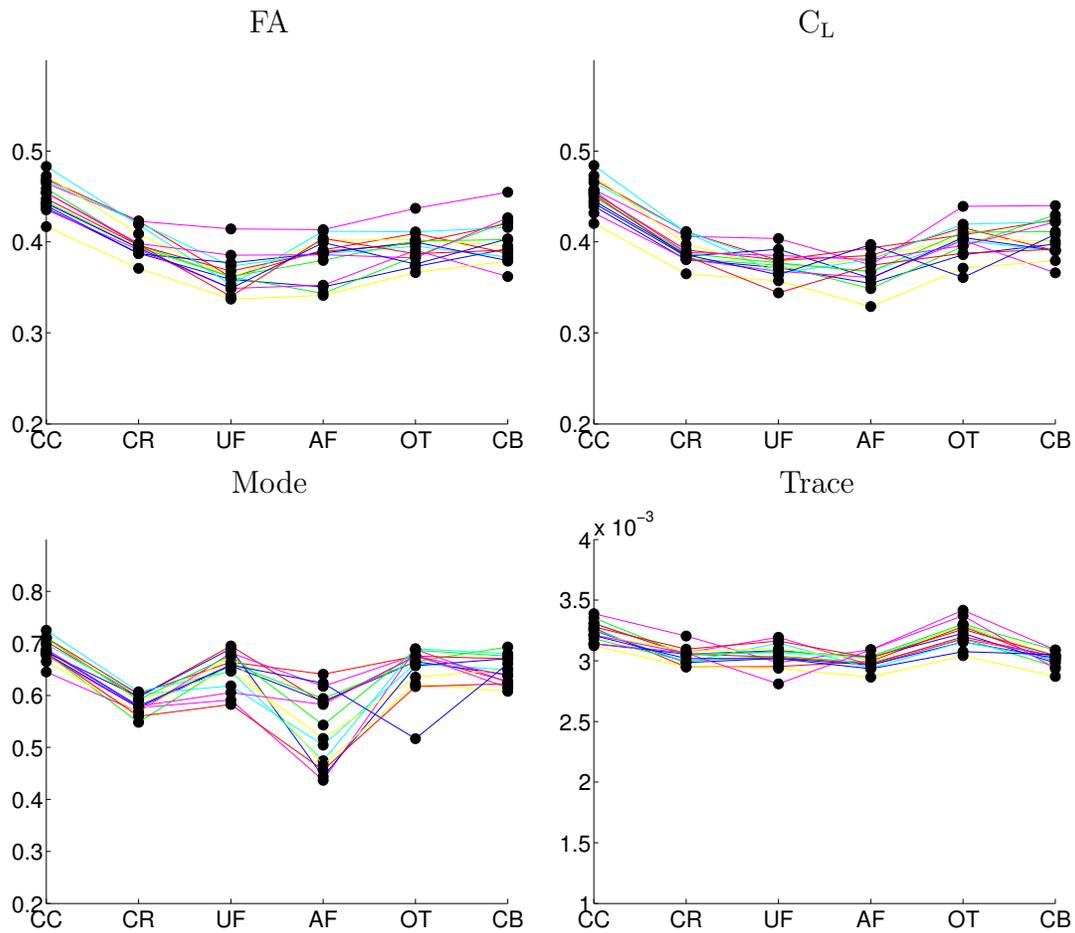


Figure 6-15: Scalar invariants in all regions for 15 subjects. Colors are per-subject (i.e. blue connects the dots from all measurements from one subject, cyan another, etc. to show the pattern across subjects). Abbreviations are: CC, corpus callosum; CR, corona radiatae; UF, uncinate fasciculi; AF, arcuate fasciculi; OT, occipitotemporal fasciculi; and CB, cingulum bundles. The arcuate measurements were made in one bilateral cluster containing the traditional C-shaped fibers.

challenges to the method, and it will be of interest in the future to try to detect unusual or crossing paths.

As regards the group registration step for Atlas 2, we did not use any shearing or nonrigid warping because those operations, when applied to the tensor fields, could change the trajectories of the paths from tractography. Rather we used translation, rotation, and scaling, which keep each trajectory the same shape, and produce the same result whether applied to the tensors or the paths (we applied the transformations to the paths). Because our method is relatively robust to slight misregistrations (as mentioned in the Discussion section in Chapter 5) due to the capture range of the similarity measure, it may not be necessary to use nonrigid registration for tractography atlas creation.

Our method for corpus callosum subdivision is unlike other methods which are rule-based [74] or depend on information from a gray matter segmentation [49, 109, 31], because our method learns a parcellation in an unsupervised way based on DTI tractography data alone. It is not clear at this point whether this is an advantage over using more information, however we point out that it allows us to find finer subdivisions than methods which have subdivided into 4 or so regions based on cortical segmentations. It is also promising that the parcellation can be learned from a group of subjects rather than imposed by predefined region or cluster definitions.

6.5 Summary

We have shown how to extend a Normalized Cuts spectral clustering solution to label new examples using the Nystrom method. Using this technique, we have extended our population clustering method to create a high-dimensional white matter fiber atlas which we have used to automatically segment tractography from novel subjects. We have presented segmentation results both of whole-brain tractography and of voxels in the midsagittal plane of the corpus callosum (a new method for corpus callosum subdivision). We have also demonstrated that this method can be used to measure scalar invariants of the tensor field in automatically defined regions of

interest that correspond across subjects. Our ability to find corresponding white matter regions enables quantitative neuroscientific studies and investigations into white matter anatomy.

Chapter 7

Clinical Application: Neurosurgery

In this chapter we present the results of a small investigation applying tract clustering to neurosurgical visualization. We believe that this is the first use of whole-brain tractography and of tract clustering for neurosurgical visualization. We begin the chapter with background information on neurosurgery and DTI in tumors, then we discuss neurosurgical planning questions that can potentially be addressed with DTI and with tractography clustering methods. Finally we present visualization results from two surgical cases. In one case, language function was interrupted during surgery, and the tract clusters we obtain support the hypothesis that the language interruption was due to the proximity of the arcuate fasciculus to the tumor. Whole-brain clustering was performed on the two subjects retrospectively (after surgery) in collaboration with Dr. Alexandra Golby (Assistant Professor, Department of Neurosurgery, Harvard Medical School).

7.1 Introduction

7.1.1 Tumors, Fibers, and Eloquent Regions

The potential utility of DTI in neurosurgery becomes clearer when one realizes that the white matter fibers can't be seen during neurosurgery. Surprisingly, the unfixed brain has a consistency not dissimilar to Jello and portions of the brain can actually

be resected (removed) using suction. Though the fibers can't be seen, a tumor may be more apparent: the consistency of a tumor is different and can be more fibrous.

During surgery, electrical stimulation of the cortex of the brain is used to interrupt neural function (speaking, counting, movement) in order to assess proximity to eloquent regions. The eloquent regions have measurable functions, as opposed to "silent" regions of the brain. The electrical stimulation is done on the cortex, not in the white matter, so it does not directly provide information on which white matter regions are functional.

7.1.2 DTI In Tumors

The main potential use of DTI in neurosurgery is to evaluate whether functional white matter is present in a tumor. However, DTI doesn't measure function, it measures structure, and it is not a simple problem to correlate DTI anisotropy measures and principal directions with white matter functionality. However, by assuming that function is preserved in white matter that looks "normal" on DTI, methods have been suggested for assessing the integrity of fiber tracts.

Using DTI color images¹ the following categories have been introduced to describe white matter tracts in the vicinity of a tumor: displaced, edematous, infiltrated, and disrupted [130, 50]. Their categorization was performed in the following way: Displaced tracts were said to be those with normal anisotropy but abnormal location or orientation. Edematous tracts looked normal (orientation and anisotropy) in the DTI color images but had high T2 intensity. Infiltrated tracts were those which had reduced anisotropy but could be identified with DTI color images. Finally, disrupted tracts were classified as such if they could not be identified on the DTI color images.

In addition to the clinical application of DTI colormaps, various groups have employed tractography to visualize the relationship between tumors and white matter [83, 111, 60]. The author of this thesis was involved in one early (2003) investigation of

¹These "DTI color" or "DTI colormap" images display the orientation of the principal eigenvector using red, green, and blue channels. The brightness of the image is controlled by tensor anisotropy (usually FA). Examples are seen in Figures 3-4, 7-1, and 7-5.

Tumor Type	%	Description
Glioma	50%	From glial cells, supporting cells for neurons This is a broad category with a wide range of malignancy.
Meningioma	15%	From meninges that cover the brain and spinal cord
Pituitary	10%	Located in pituitary gland (may secrete hormones)
Pineal	1%	Located in pineal gland
Metastatic		Originates elsewhere in the body
Lymphoma		Primary tumor, or rarely secondary (metastatic)

Table 7.1: Common Tumor Types in Adults (percentages from [1])

tractography for surgical planning [111] which used the author’s DTMRI module for DTI visualization in the 3D Slicer program [44]. In one recent paper, intraoperative tractography was used to visualize and quantify the shifting of white matter fibers (primarily the pyramidal tract but in one case the corpus callosum) that occurs during surgery [83]. This study found average shifts of 2.7 mm, with a range from -8 to +15mm, where negative and positive numbers indicated movements toward and away from the craniotomy opening, respectively. To our knowledge, other than in this investigation, whole-brain tractography has not been employed to date for neurosurgical visualization or planning.

Common types of tumors and their characteristics are listed in Table 7.1. Malignancy in brain tumors is defined not by whether they metastasize, since this does not happen with primary brain tumors, but rather by how quickly they grow and by their cellular differentiation as measured by pathology (whether the cells look normal or have many mitoses, etc.). However tumors in other parts of the body can metastasize to the brain.

7.1.3 The Neurosurgical DTI Visualization Problem

In this section we aim to define the problem of DTI visualization for surgery, focusing especially on how to create useful tractography images. (Currently, Dr. Golby’s lab primarily uses slice-based glyph visualization for presurgical DTI.) We have created a list of questions relevant to neurosurgery that can potentially be answered with DTI or DTI-derived images, shown in Table 7.2. In Table 7.3 possible methods are

proposed to answer each question. Later in this chapter we apply some of the methods to neurosurgical DTI visualization.

1	Are the tracts symmetric across hemispheres?
2	Are there tracts in the tumor? Behind the tumor?
3	Exactly where is the 3D tractography in relation to the image slices?
4	In which direction should the initial incision be made?
5	How “true” is the tractography?

Table 7.2: Neurosurgical planning questions potentially answerable by DTI and tractography.

1	Color code matching tractography on right and left.
2	Color the part of the tractography that is inside the tumor, if tumor boundaries are known. Otherwise color the tractography according to an MRI image which clearly shows the tumor.
3	Show the tractography on the 2D images.
4	Display the fiber orientation near the cortex for incision planning.
5	Provide an estimate of how certain it is that the tract exists using anisotropy values or a bootstrap method.

Table 7.3: Proposed methods to address the questions.

Question 1 is “Are the tracts symmetric across hemispheres?” This relates to the general concept that doctors analyze brain images by comparing the left and right sides to look for abnormalities. It is not clear how to do that with three-dimensional information from tractography. One potential method, listed in Table 7.3, is to aid visual comparison of tractography in both hemispheres by color-coding corresponding tracts on the right and left sides. This is not trivial as it depends on having either anatomical labels for the tracts or an algorithm for matching them across hemispheres which is robust to pathology. (We have discussed an automatic solution to this problem that is successful in normal brains, and we have presented results in chapters 5 and 6).

Question 2 is “Are there tracts in the tumor? Behind the tumor?” This is especially important when deciding where to stop resection if the tumor borders a tract (meaning the tract is behind the tumor with respect to the resection path). The question “What is in front/behind/inside the tumor?” cannot be easily answered visually,

both due to computer graphics issues and to the fact that the true tumor border may not be known. In the 3D Slicer visualization system [44], it is possible to make a surface model of the tumor semi-transparent to see what is inside, but then it is not possible to differentiate whether tractography is inside or behind the tumor. Answer number 2 in Table 7.3 lists potential methods to address this ambiguity by painting information from MRI images onto the tractography trajectories, augmenting the 3D view of the trajectories by coloring them where they pass through the tumor. We have tested these ideas in a surgical case, and the results are not conclusive but do provide more visual information than before (see Figure 7-4).

Question 3 is “Exactly where is the 3D tractography in relation to the image slices?” This question, like the previous one, has to do with the 3D information being spatially unclear. When viewing tractography in 3D along with 2D image slices, the 3D extension of a tractographic path often visually occludes its intersection with the 2D slice and therefore with the tumor. One approach is to show the paths in the 2D slices instead or in addition to the 3D view. Painting voxels according to tractography is a method which has been employed by Mori et al. in creation of a voxel atlas [80] and by Conturo et al. in one of the earliest DTI tractography papers [30].

The fourth question in Table 7.3 is “In which direction should the initial incision be made?” When neurosurgeons make the initial incision, they try to make it parallel to white matter fibers in order to cut as few as possible. They do not “scoop” out a cavity, but rather cut a line in the surface and then follow that direction inward. The orientation of this line on the surface is determined by the known anatomical orientation of the fibers, however this is not patient-specific and may not be correct in all cases. This orientation may vary in neuroanatomical illustrations, as in for example the fibers in Meyer’s loop around the inferior horn of the lateral ventricles, where an incision is made to approach tumors of the ventricles. Images that illustrate which way the fibers are going below the surface, for example of the cortex of the brain, would aid in planning the orientation of the incision. We have not implemented any method to address this question, and it is not straightforward because the DTI fiber orientation information is less reliable near the cortex where white matter transitions to gray

matter and anisotropy decreases. However, an appropriate experiment is to paint the cortex or another surface such as the lateral ventricle using the standard RGB color scheme, based on the orientation of white matter that is somehow “nearest” to each location on the surface.

Question 5 is “How true is the tractography?” This relates to the fact that the fibers displayed using tractography are most assuredly not representations of individual axons (with a scale of microns), but rather are paths which follow the most probable fiber orientations measured on a scale of millimeters using diffusion MRI. Due to noise, partial voluming, or fiber crossing, tractography makes anatomical errors, and it is important to a) display only paths that pass quality thresholds and b) provide a visual quality estimate. Point a) is addressed by the anisotropy thresholds which are used for seeding and stopping tractography, though for the case of tumor visualization one may not wish to apply stringent thresholds which could discard or shorten many paths. Rather it may be preferable to address the quality visually without making hard decisions to remove many paths: one possibility for b) is to color paths based on anisotropy values. A better approach, if more than the minimum of 6 diffusion-weighted images have been acquired, is to use a bootstrap method similar to the method proposed for display of uncertainty of tensor fields [53].

7.2 Experimental Results

7.2.1 Methods

Retrospective visualizations were created for two surgical cases in which tumors had been resected by Dr. Golby. Each case had presurgical data including T1- and T2-weighted MRI, DTI, and fMRI (functional MRI).

Creation of the visualizations had several steps. We summarize the steps here; chapter 5 provides more detail about the method. The first step was whole-brain tractography, which was initiated in each voxel (of a region of interest including the whole brain) where C_L was above .3. Tractography stopped when C_L fell below

0.15. This relatively low stopping value was chosen with the goal of “not leaving out anything” in the tractography. The minimum path length was 40mm. It is not possible (on a current PC) to rapidly interactively visualize all of the paths output from whole-brain tractography, so a random subset of the paths produced was selected for further processing in each subject. We selected 10,000 for case 1 and 8,000 for case 2 (because after analyzing case 1, we desired faster interaction speed during the final manual cluster labeling and coloring step). Note these numbers of paths are still larger than those clustered with other methods in the literature [18, 43].

The second step in production of the visualizations was an automatic spectral clustering of the paths from tractography to produce 120 bundles. We chose 120 because (as demonstrated in Chapter 8) there are more than 100 clusters inherent in single-subject whole-brain tractography data. The fiber affinity measure was the mean closest point distance, with symmetrization using the mean value (see Chapter 5 for more details).

The clustering step was followed by an interactive manual anatomical labeling of clusters, performed with repeated reference to the white matter anatomy described in [84, 50, 116]. Finally, in response to feedback from the neurosurgeon, the clusters were colored to match bilaterally. The coloring was done manually (according to the anatomical labels) as an experiment for these two cases, and this motivated the research into automatic cluster correspondence across hemispheres that was presented in Chapter 5.

7.2.2 Case 1

Patient 1, a 54 year old male, presented with a seizure and word finding difficulty, and on MRI had a large left temporal mass. The tumor was resected and the pathologic diagnosis was anaplastic oligoastrocytoma, WHO Grade III. During surgical resection in the posterior region, the patient had difficulty with speech; superiorly, electrical stimulation interrupted language function. About half of the tumor was resected while the rest was left to preserve function. Dr. Golby interpreted the speech arrest during surgery as due to the proximity of functional white matter to the resection

region. The goal of the retrospective visualization, consequently, was to investigate whether there was a fiber tract in the tumor. Specifically, the hypothesis was that the arcuate fasciculus, the tract connecting speech production areas, was involved in the tumor.

Figure 7-1 shows a standard 2D slice visualization at one axial location for patient 1. Due to the lack of anisotropy, according to the categorization of [130, 50] the DTI colormaps seem to indicate that tracts are infiltrated or disrupted in this level of the tumor, or displaced elsewhere. However there are some bright spots in the tumor region seen in the FA image, and from this single slice image one can't tell about tract presence/absence in the rest of the tumor.

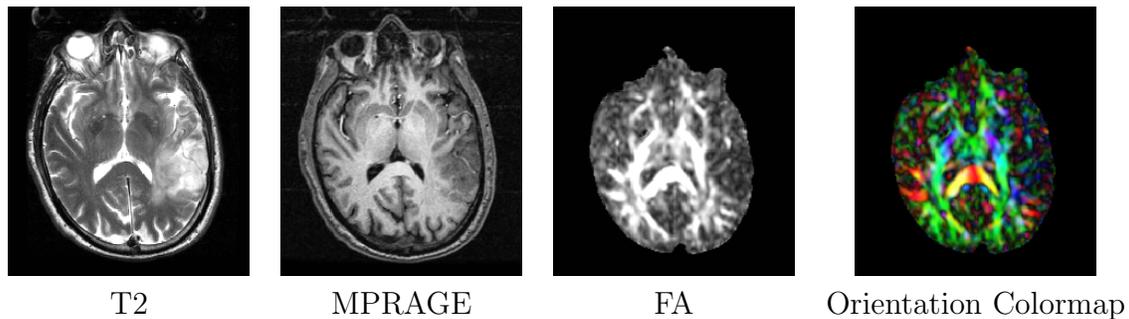


Figure 7-1: Slice-based DTI visualization of patient 1. The left two images are MRI structural images. The tumor (and surrounding edema) are bright in the T2 image, and the tumor is dark in the MPRAGE image. The FA image shows the diffusion anisotropy. The orientation colormap displays the direction of the major eigenvector of the tensor: red represents left-right orientation, green is anterior-posterior, and blue is superior-inferior. Brightness of the colormap is controlled by anisotropy. Decreased brightness is seen in the tumor area, indicating abnormal white matter.

Figure 7-2 displays the raw output of the tractography clustering, where colors were determined by the clustering algorithm (similar colors appear when tractographic paths map near each other in the high-dimensional space in which clustering is performed). The anatomical interpretation of this initial clustering is not clear when all paths are viewed at the same time as in Figure 7-2.

To improve the visual anatomical information, the tractography in the left and right hemisphere was manually labeled using colors matched by structure, in order to facilitate right/left comparisons. This labeling was done manually as an experiment

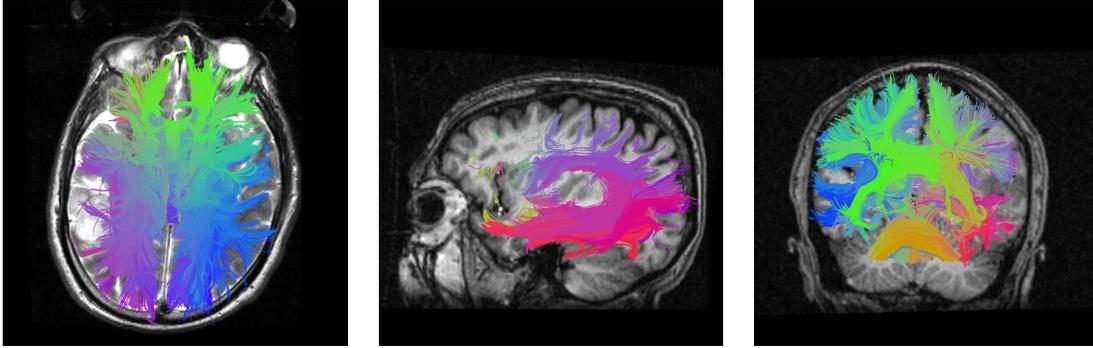


Figure 7-2: Raw output of whole-brain clustering for patient 1, shown with axial, sagittal, and coronal images. Each unique color represents a cluster.

to see if the colors are useful for surgical planning. Figure 7-3 illustrates a subset of the clustered and labeled tracts, with consistent anatomical coloring bilaterally. fMRI activations from a language task are rendered in bright yellow. The fMRI task was categorizing nouns as living or nonliving, where patient would respond with either “living” or “nonliving” for each noun. The larger size of the activations on the right side could indicate migration of function to that side.

Due to the proximity of the tumor to the arcuate fasciculus region, and the speech arrest which took place during resection, an attempt was made to clarify if this tract was passing through the tumor. Figure 7-4 demonstrates one possible method for clarifying whether a tract is in a tumor. The tractographic paths were “painted” with the T2 image volume (keeping the same window/level as the T2 images) in order to transfer its information about tumor boundaries to the tract. In general this method can be used with any available images.

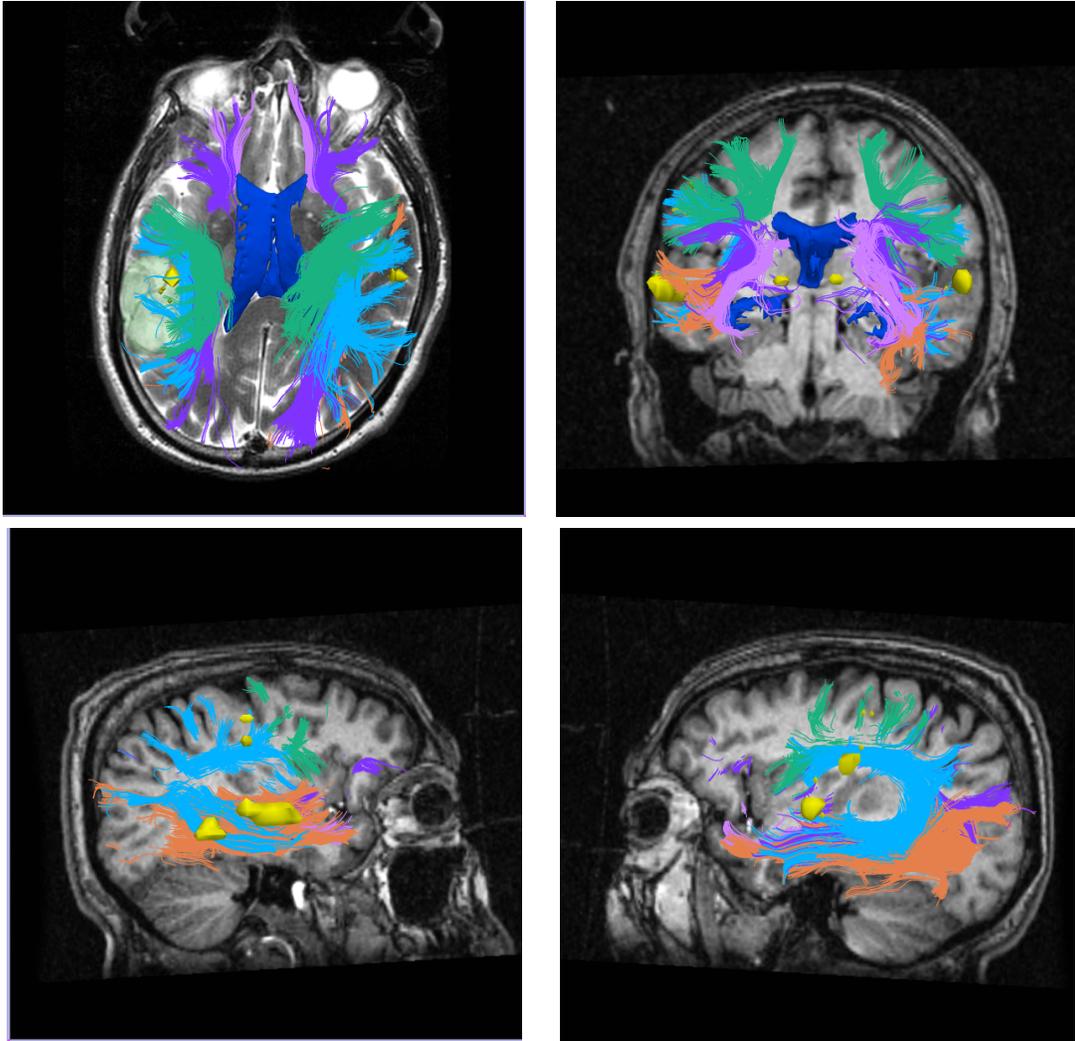


Figure 7-3: Three-dimensional DTI visualization of patient 1, with consistent coloring on both sides to facilitate left/right hemisphere comparison. Note that the tumor has displaced tracts inferiorly in the coronal view. The colors correspond to anatomy as follows: dark blue, ventricles; light blue, arcuate fasciculus; orange, inferior longitudinal fasciculus; dark purple, inferior fronto-occipital fasciculus; green, u-fibers (may include some SLF); light purple, uncinate fasciculus; and yellow, fMRI language activations. All labeled tracts can't be seen at once, so a subset is displayed.

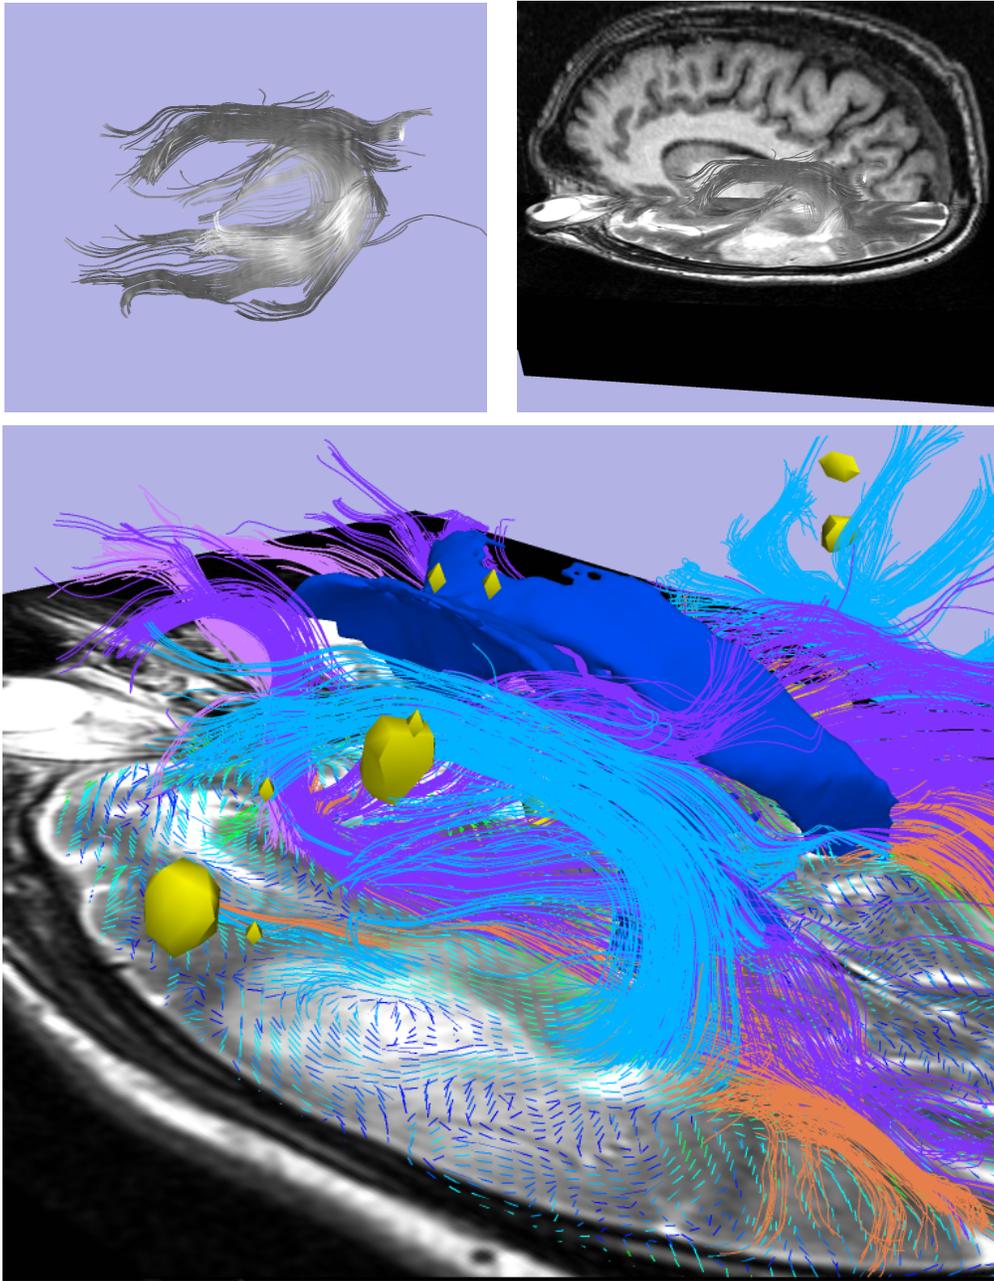


Figure 7-4: Spatial relationship of arcuate fasciculus cluster and tumor. The top two images show the left arcuate fasciculus, as defined by fiber clustering and labeling, colored using the T2 weighted images (the tumor is bright). In the top right image the axial T2 is displayed to allow comparison of intensities in the tumor and on the tract, while the sagittal T1 is shown for background contrast to make the tract more visible. The bottom image illustrates the arcuate fasciculus cluster in bright blue with fMRI language activations in yellow. The eigenvector glyphs in the bottom image are colored by anisotropy, with bright blue being highest.

Discussion

The original hypothesis that the arcuate fasciculus was involved in the tumor region is supported by our tract clustering and visualization method. The problem addressed in Figure 7-4, whether the arcuate fasciculus is literally inside the tumor, is difficult to answer because the T2 volume could be bright due to edema rather than tumor, so the exact tumor boundaries are not known. However we have proposed a method that can aid the surgeon's judgement. Once the surgeon makes a decision about tumor boundaries in the T2 (or other MRI) image, the T2 intensity information can be transferred to the tractography, providing more information about the tract-tumor spatial relationship. The bright white areas on the tract in Figure 7-4 indicate that it is very near the tumor, and the judgement of the surgeon can now be employed to decide about tumor boundaries just as when using the T2 MRI images directly. This painting method was considered to be a potentially useful tool by the neurosurgeon.

In general it is hard to determine where a tract ends in a tumor. It is possible that edematous tissue contains functioning fibers, but that due to edema the anisotropy in the region is low, causing the tractography algorithm to stop tracing². Note that in Figure 7-4, in the right image with eigenvector glyphs, there appears to be a region of coherent orientation that was not traced by tractography. The termination of tractography is possibly due to lower anisotropy caused by edema or tumor infiltration. However the lack of trajectories in this region could be due to our length cutoff to remove short paths before clustering. In the future when shorter u-fibers (short cortico-cortical connections) are of interest due to their proximity to a tumor, lowering the path length threshold will be tested.

The visualizations that we have created were found by the surgeon to be more useful when colored in a corresponding way bilaterally. For example in Figure 7-3 in the coronal view, the displacement of tracts in an inferior direction can clearly be seen in the hemisphere affected by the tumor.

²In fact, anisotropy has been found to be reduced in the peritumoral area of both gliomas and metastatic lesions, but the mean diffusion (ADC) was higher at the border of metastatic lesions, apparently due to a higher amount of edema [106]. The anisotropy decrease bordering gliomas was partially attributed to infiltration and there was less edema [106].

7.2.3 Case 2

Patient 2 is 28 years old and presented with seizures. On MRI a lesion was seen in the right insula, extending into the frontal and temporal lobes. The tumor was classified by pathology as a mixed oligoastrocytoma, WHO Grade II/IV.

Figure 7-5 shows the 2D visualization of the patient, while Figure 7-6 demonstrates the 3D tractography visualization with the same colors as applied in Case 1. The shape of the corona radiata is quite different from left to right, and the uncinate fasciculus appears to curve around the tumor. On the right side, fewer paths belonging to the inferior occipitofrontal fasciculus (dark purple) were found, indicating lower anisotropy in this region.

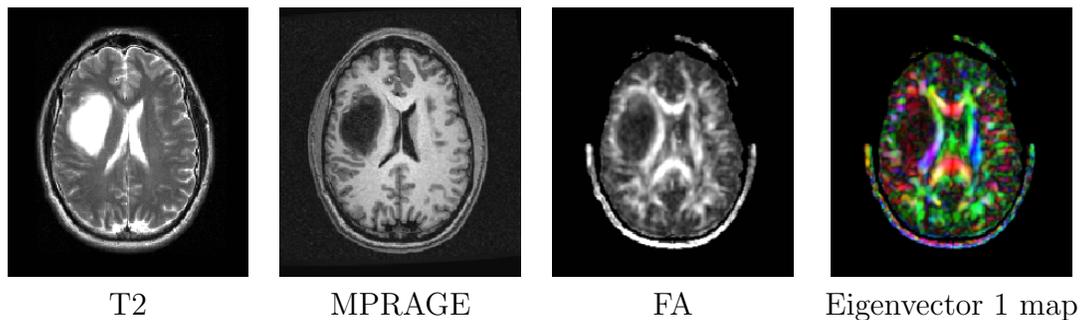


Figure 7-5: Slice-based DTI visualization of patient 2.

Discussion

Using the same settings (and the same voxel size and similar number of slices) nearly 50,000 paths were initially produced in patient 2 (28 years old), versus close to 20,000 in patient 1 (54 years old). This difference may be caused by the overall anisotropy decrease with age, and shows that anisotropy thresholds for tractography (generally held constant in the literature at 0.2 FA) could benefit from further investigation, especially in cases of pathology or when comparing different age groups.

As in patient 1, the bilateral tract coloring allows visual comparison of the two hemispheres. The tract clustering indicates that the uncinate fasciculus (light purple in Figure 7-6) and the internal capsule (yellow) have been displaced by the tumor.

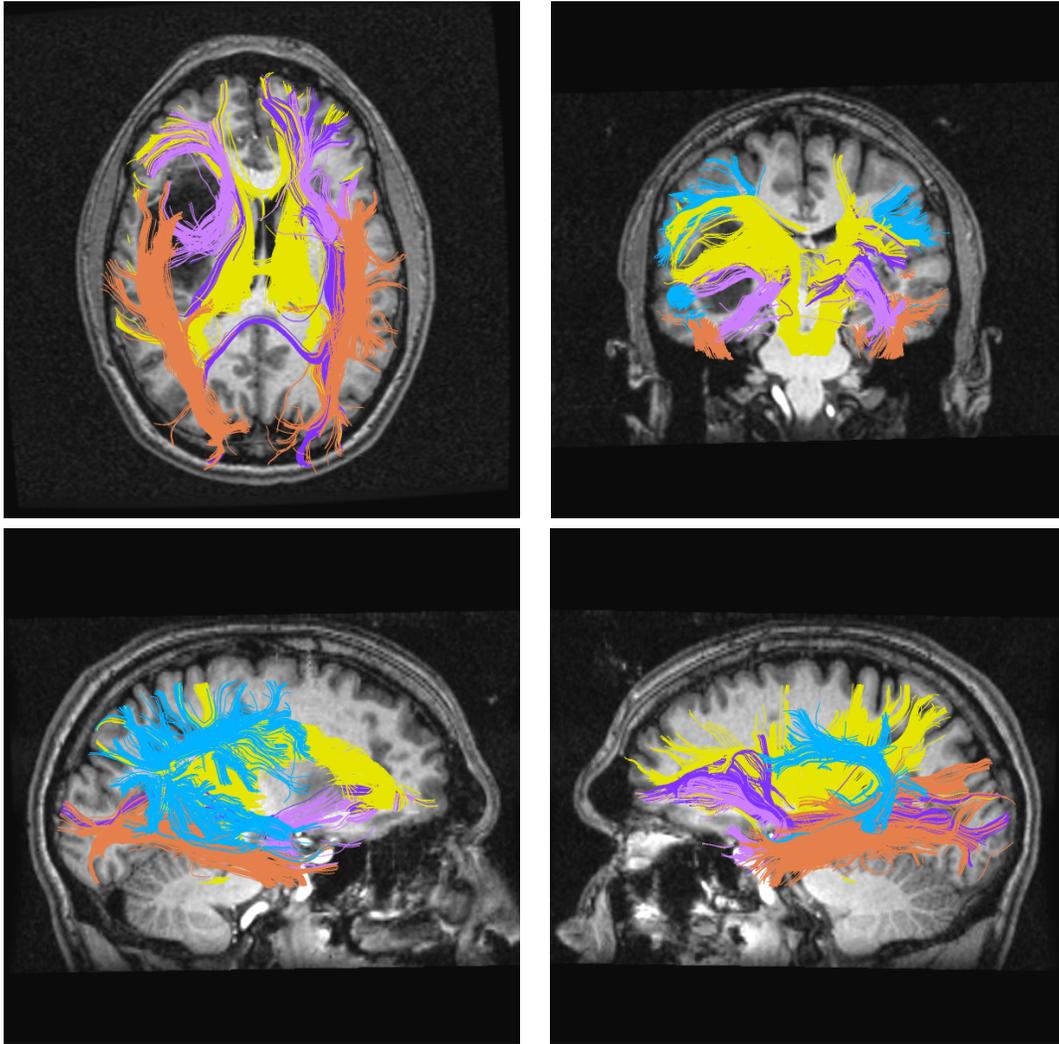


Figure 7-6: Three-dimensional DTI visualization of patient 2, with consistent coloring on both sides to facilitate left/right hemisphere comparison. The top left image is an inferior view, not superior as used for patient 1. The colors correspond to anatomy as follows: light blue, arcuate fasciculus and (on right side) additional fibers bordering tumor; orange, inferior longitudinal fasciculus; dark purple, inferior fronto-occipital fasciculus; light purple, uncinate fasciculus; yellow, corona radiata. Note the differences in the uncinate fasciculus and the corona radiata across hemispheres. There is a possible error in tractography in the top left image, where several paths (dark purple) traverse part of the corpus callosum and part of the occipitofrontal fasciculus.

7.3 Summary

We presented an investigation into the application of tractography and tract clustering to neurosurgical visualization. First we created a list of neurosurgically relevant questions that could potentially be answered by DTI processing and visualization. Then we proposed methods to address the questions. Finally, we created retrospective visualizations of two surgical cases, applying the clustering method presented in this thesis and some of the methods proposed in this chapter. The procedure for creating the visualizations included an automatic part: whole-brain tractography and tractography clustering; and a manual interactive part: anatomical cluster labeling and bilaterally symmetric coloring of clusters/anatomical structures. We further investigated how to visually provide information about the presence of a white matter tract inside (or near) a tumor, implementing one method for “painting” the tractography with information from an MRI scan which clearly illustrates the tumor.

The tract clustering experiments presented in this chapter demonstrate the potential clinical utility of the methods presented in this thesis. The visualizations with bilaterally matching colors were found to be more useful by the surgeon than the original clusters (and this has motivated research into the automatic bilateral cluster matching presented in Chapter 5). The method of painting MRI intensity value information onto a tract cluster was also considered by the surgeon to be potentially useful in surgical planning. Finally, the clustering supported the hypothesis that the arcuate fasciculus was involved in the tumor region in patient 1. The location of the arcuate fasciculus cluster near the tumor border was consistent with the speech arrest that took place during surgical tumor resection.

Chapter 8

Reproducibility and Parameter Settings

In this chapter we investigate two types of reproducibility: reproducibility of the embedding, clustering, and automatic segmentation; and reproducibility of the measurement of scalar invariants. Our choices for some of the parameter settings used for spectral clustering have been motivated by this reproducibility analysis. In this chapter we also show results of experiments investigating the sensitivity of our method to changes in parameter settings. All of the parameters of our method are listed in Table 8.1. We organize the experiments in this chapter as in the table: first we present experiments related to distance/affinity calculation; then spectral embedding, clustering, and atlas generation; and finally measurement of scalar invariants using atlas-segmented regions of interest.

8.1 Distance/Affinity Calculation

We investigated the effect of the sigma (σ) parameter used in conversion from pairwise fiber distance to affinity. For relatively local neighborhoods of paths from tractography, the embedding is stable over a range of values of σ as shown in Figures 8-1 and 8-2.

Parameter	Function
T_{seed}	start tractography where $C_L > T_{\text{seed}}$
T_{stop}	stop tractography where $C_L < T_{\text{stop}}$
T_{length}	cull paths whose length $< T_{\text{length}}$
σ	Std. deviation of Gaussian kernel for distance to affinity conversion
D	Distance measure chosen to compare paths
S	Method used to symmetrize the distance measure, if needed
N_E	Number of eigenvectors used to define embedding space; also length of embedding vectors (dimension of embedding space)
N_A	Number of randomly selected paths used to calculate the \mathbf{A} matrix
k	Number of clusters
N_P	Number of paths per subject used in clustering for atlas creation
N_S	Number of subjects used in atlas creation
R	Alignment method used before computing tract affinities

Table 8.1: Parameters of the method. First listed are parameters for tractography, then for distance/affinity calculation, then for spectral embedding and clustering, and finally for atlas generation.

8.2 Spectral Embedding, Clustering and Atlas Generation

In this section we present quantitative results measuring embedding reproducibility with single-subject and atlas data. We also present results measuring reproducibility of clustering versus the number of clusters; this gives an idea of the number of clusters inherent in the data. In addition we show images demonstrating that atlas labeling results from the same subject are similar across scans. Finally, we give quantitative results showing that the segmented clusters are very similar if a subject is clustered as part of the atlas, or labeled using the atlas.

The DTI dataset used for the single-subject experiments (on embedding reproducibility and number of clusters) was from Population 1, and the mean closest point distance measure was used with mean symmetrization and σ of 60. The atlas dataset consisted of the 10 subjects from Population 1 which were used in atlas generation in Chapter 6. For the atlas experiments the mean closest point distance measure was used with minimum symmetrization and σ of 30.

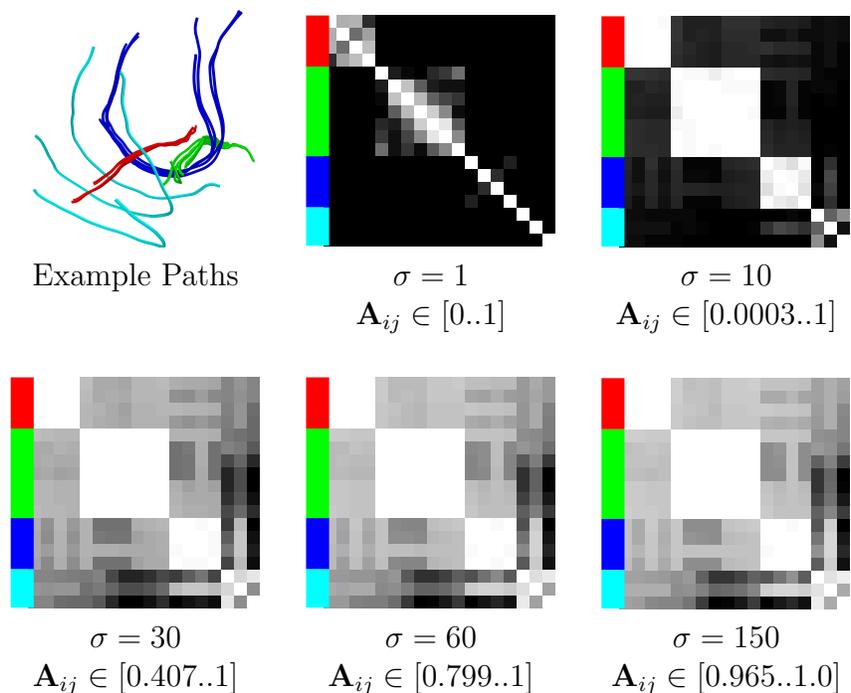


Figure 8-1: Effect of varying σ on affinity matrix. The affinity matrices were calculated from the pairwise fiber distances for the tractography example (upper left). Each matrix was created with a different value of σ which is shown below the matrix along with the range of values in the matrix. The affinity matrix images have been displayed with the lowest value as black and the highest as white. The minimum was used to symmetrize the pairwise distances but similar results are obtained when the mean is used. This example shows that for a range of values of σ , the affinity matrices are very similar (up to a scale factor which will be removed during matrix normalization). This observation is generally useful because if this small clustering example were placed within a larger whole-brain clustering problem, these specific (local) pairwise affinity values would not change.

8.2.1 Reproducibility of Embedding Subspace

The reproducibility of clustering is directly related to the reproducibility of the embedding. Considering the embedding problem as a Nystrom eigenvector approximation, successive eigenvector approximations can be compared to test if the embedding is stable. (The reason all of the eigenvector approximations aren't identical is that they are calculated by extending the eigenvector solution computed for a small random subset of the data.) If the eigenvectors span a similar subspace, the embedding is considered similar (the eigenvector basis can rotate when eigenvalues are similar, so

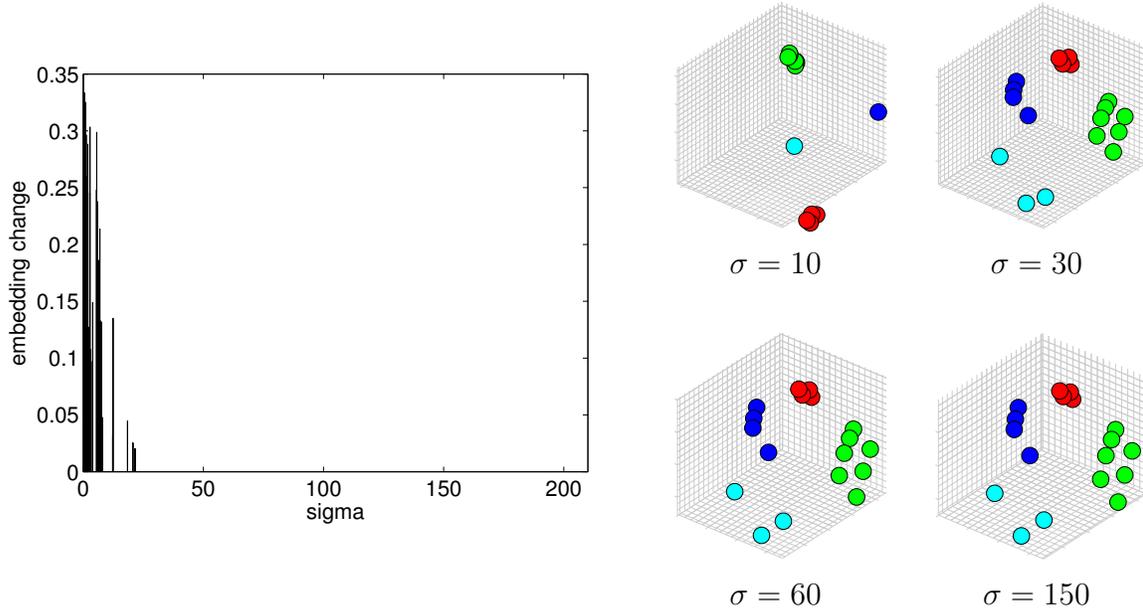


Figure 8-2: Effect of varying σ on embedding. For each value of σ , an affinity matrix (shown in Figure 8-1) and an embedding (images at right) were computed. The embedding is relatively stable for values of σ above 22mm in this example. The plot (left) shows how much the embedding vectors move on average when σ changes. To create the plot, for each value of σ the average distance (from that embedding to the previous embedding) was computed between corresponding embedding vectors. This average distance and the standard deviation of the distance were approximately zero for $\sigma \geq 22$. The values of σ used for clustering in this thesis are within this stable range.

the important thing is that the same subspace is described).

The normalized Frobenius matrix norm

$$\frac{1}{N_E} \|\mathbf{U}^T \mathbf{V}\|_F, \quad (8.1)$$

where \mathbf{U} and \mathbf{V} are two matrices containing N_E eigenvectors, is high when the N_E selected eigenvectors span the same subspace [40]. As described in [40], we have performed spectral embedding for various sizes of the \mathbf{A} matrix, for several random selections of \mathbf{A} per size, to investigate how large \mathbf{A} should be for repeatable embedding. In addition we performed the experiment using several different numbers of eigenvectors (the number of eigenvectors determines the dimensionality of the em-

bedding space). We averaged the norm over unique pairs of embedding trials. We performed this experiment in a single subject and with the atlas dataset.

In Single Subject

Figure 8-3 shows the result of the embedding reproducibility experiment in the single subject dataset. 10 embeddings were performed for each size of \mathbf{A} . For this dataset, reproducibility is highest when using approximately 20 eigenvectors, indicating that the cluster information is well-represented in 20 dimensions. This has motivated our use of 20 eigenvectors for other clustering results presented in this thesis. Reproducibility for any number of eigenvectors is highest when using 1,500 or more paths when creating the \mathbf{A} matrix. For this reason, 1,500 is our recommended minimum number of paths for creation of \mathbf{A} when performing single-subject clustering.

In Atlas

Figure 8-4 shows results of the embedding reproducibility experiment when performed on an atlas dataset. In this experiment the paths selected for affinity computation (random sample used to form \mathbf{A}) came from all subjects, and embedding was performed for all subjects as described in Chapter 6. There were 10 subjects used in atlas embedding, with 3,000 trajectories each. Embedding was performed 5 times for each size of the \mathbf{A} matrix, and the mean and standard deviation of the reproducibility value (matrix norm) was computed using all unique pairs of embeddings.

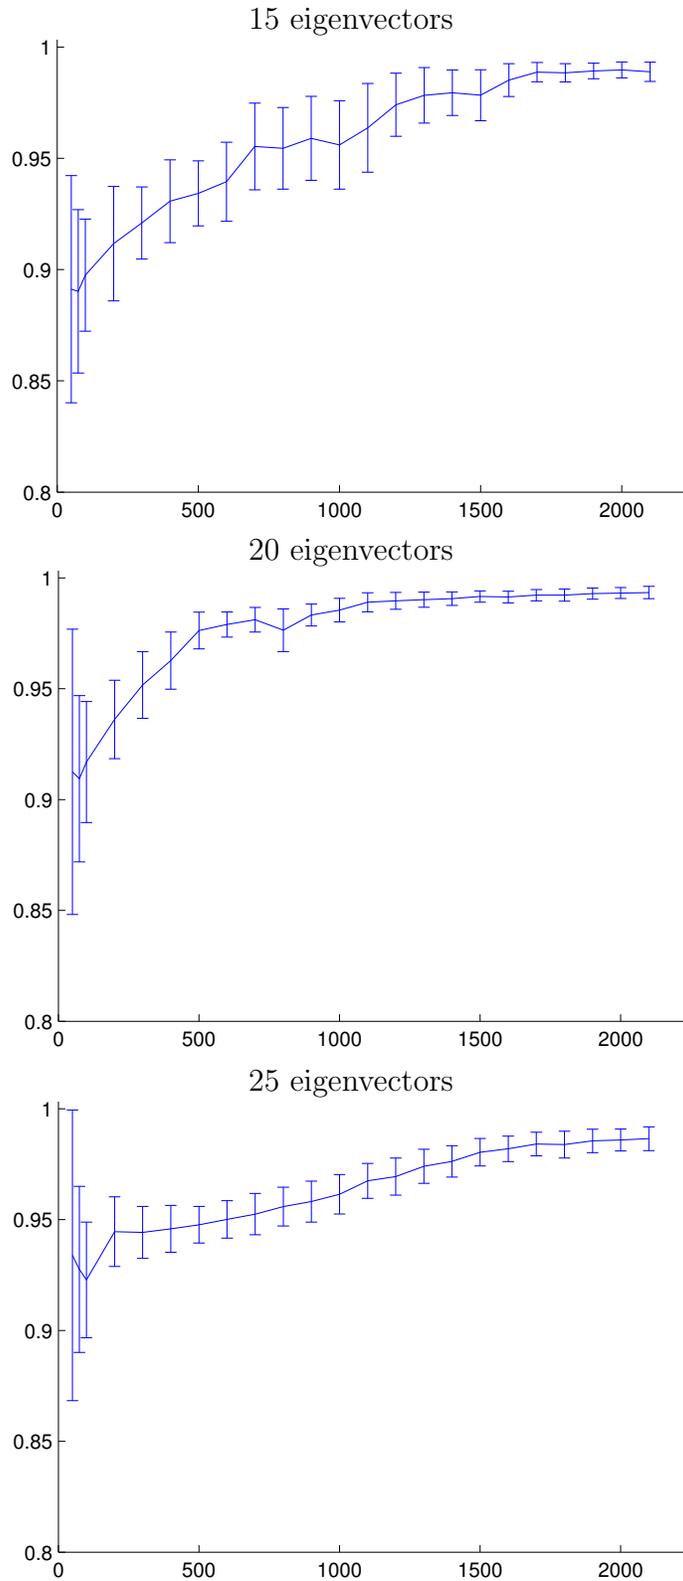


Figure 8-3: Embedding reproducibility. Matrix norm (reproducibility of embedding subspace) versus the size of the \mathbf{A} matrix, for different numbers of eigenvectors (dimensionalities of the embedding subspace). \mathbf{A} of size 1500x1500 or greater with 20 eigenvectors is optimal for this single-subject dataset.

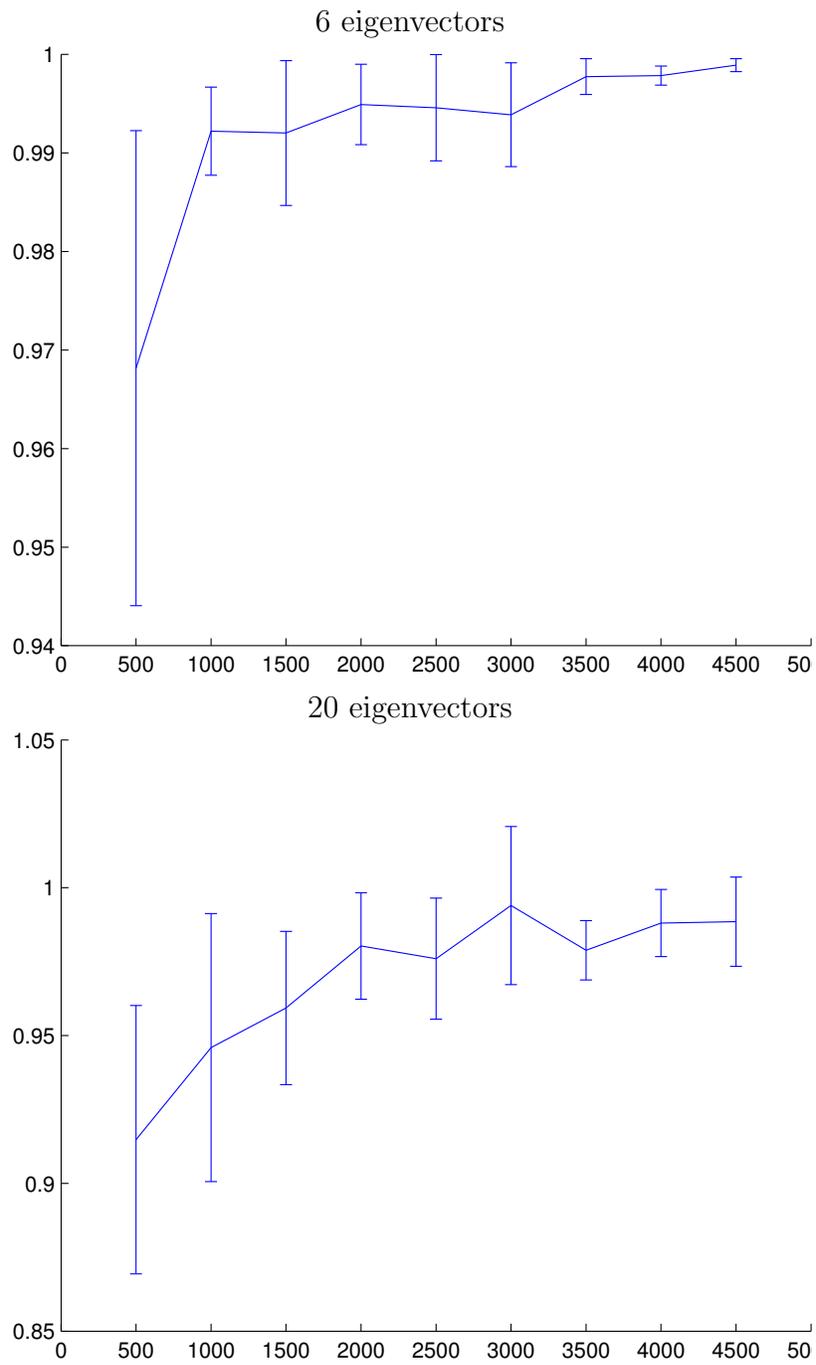


Figure 8-4:
 Embedding reproducibility for atlas data. Matrix norm (reproducibility of embedding subspace) versus the size of the \mathbf{A} matrix.

8.2.2 Number of Clusters

The reproducibility of clustering is also related to the number of clusters inherent in the data. If too few clusters are used, different partitions may be found, leading to low reproducibility. If too many clusters are used, data points which are near the border between nearby clusters may be sometimes assigned to one cluster and sometimes another. We quantified this clustering variability by testing whether pairs of data points (paths from tractography) clustered always in the same cluster, never in the same cluster, or sometimes in the same and sometimes in different clusters. The fraction of pairs that sometimes were in the same cluster, and sometimes in different clusters, was measured for various numbers of clusters. For each number of clusters, we ran clustering 10 times. The fraction of pairs that were not always either in the same or different clusters (for all 10 runs) is plotted in Figure 8-5. The lowest values in the plot show the best performance (highest consistency across pairs) and this would indicate that for this dataset, either using around 75-100 clusters or around 250 clusters would give the most repeatable results. In any case, the fraction of pairs that are inconsistent is less than 5 percent for any number of clusters greater than about 75.

The quality of the clustering solution can also be measured using the distances from embedded points to cluster centroids. A worst-case measure, the maximum sum of point-to-centroid distances over all clusters, was measured for various numbers of clusters. The result is shown in Figure 8-6. 100 or more clusters give the best result according to this measure.

8.2.3 Segmentation Results From Repeat Scans

Selected clusters from atlas labeling of two scans of subject 11 are shown in Figures 8-7 and 8-8.

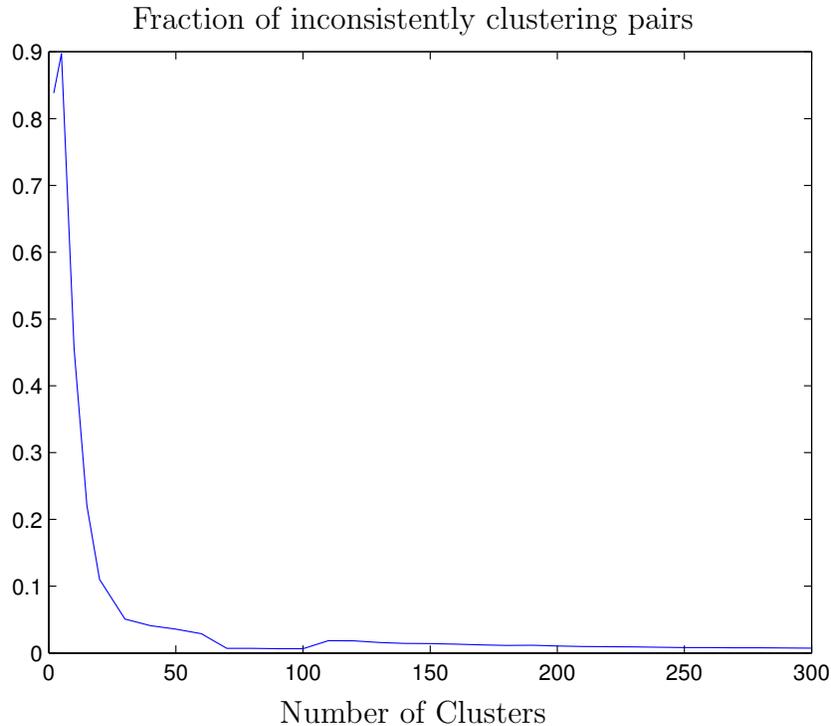


Figure 8-5: Fraction of pairs of paths that clustered sometimes in the same cluster and sometimes in different clusters, for 10 clusterings of a single-subject dataset. The plot indicates that there are more than around 75 clusters in this single-brain dataset.

8.2.4 Consistency of Atlas Clustering and Atlas Labeling

We created two atlases, one using all 10 subjects, and the other using 9 subjects. The subject which was left out was then labeled using the second atlas. We compared the clustering result (when the subject was part of atlas clustering) to the labeling result (when the subject was automatically segmented using the leave-one-out atlas). To quantify clustering consistency, we counted the number of pairs of trajectories that were clustered together in both results, that were clustered separately in both results, and that were clustered once together and once separately. The percent of pairs that was consistent (both in the same or both in different clusters) was 99.38%.

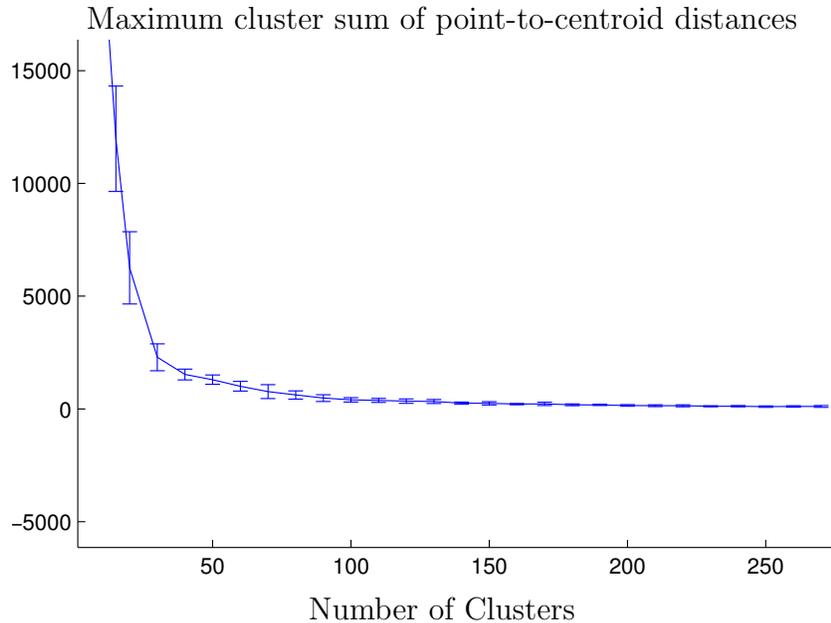


Figure 8-6: Maximum cluster centroid distance sum versus number of clusters, indicating that 100 or more clusters should be used for this single-subject dataset.

8.3 Stability of Measured Scalar Invariants

8.3.1 Comparison of FA Values From Repeat Scans

Using the atlas, we automatically segmented tractography and then measured FA in two subjects (10 and 11) who had each been scanned three times. The quantitative measurements of FA demonstrate reproducibility and are shown in Figure 8-9.

8.4 Summary

In this chapter we demonstrated two types of reproducibility: reproducibility of the embedding, clustering, and automatic segmentation; and reproducibility of the measurement of scalar invariants. We showed images from repeat scans, and quantitative results regarding cluster stability and FA measurement, indicating that our automatic segmentation method is reproducible. In addition we specified all of the parameters of our method. The experiments presented here have motivated our use of parameter settings for σ , the number of eigenvectors, and the number of clusters.

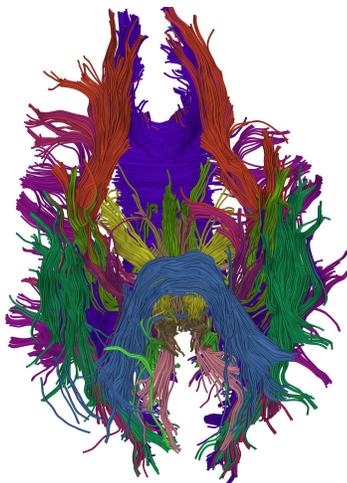


S11

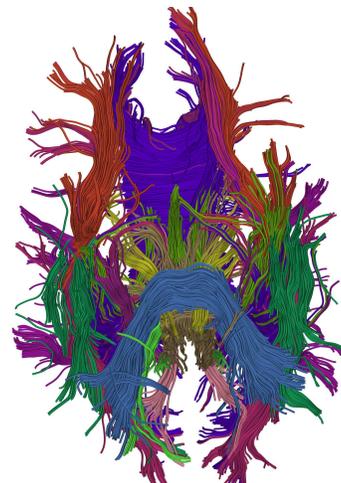


S11 (repeat scan)

Figure 8-7: Result of repeat automatic segmentation of novel subject (view from left). Selected regions are shown as follows: navy blue, corpus callosum; yellow, corticospinal fibers; purple, arcuate fasciculus/SLF region; orange, uncinate fasciculus; green, occipitotemporal (inferior longitudinal) fasciculus; sky blue, middle cerebellar peduncle; light pink, superior cerebellar peduncle; hot pink, occipitofrontal fasciculus.



S11



S11 (repeat scan)

Figure 8-8: Result of repeat automatic segmentation of novel subject (inferior view). Subject numbers are below the images. Selected regions are shown as follows: navy blue, corpus callosum; yellow, corticospinal fibers; purple, arcuate fasciculus/SLF region; orange, uncinate fasciculus; green, occipitotemporal (inferior longitudinal) fasciculus; sky blue, middle cerebellar peduncle; light pink, superior cerebellar peduncle; hot pink, occipitofrontal fasciculus.

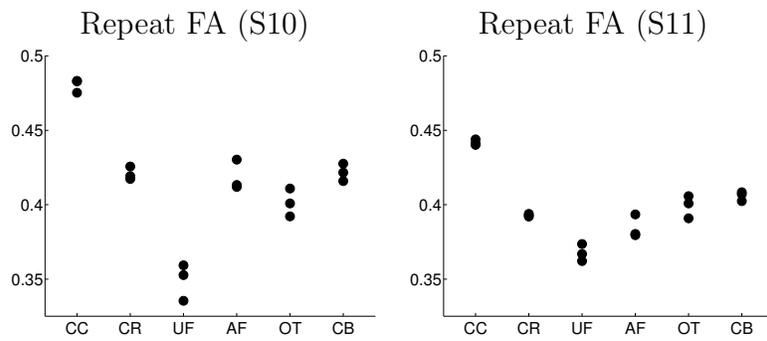


Figure 8-9: Mean FA measurements in regions, from three repeat scans of two subjects (10 and 11). Abbreviations are: CC, corpus callosum; CR, corona radiatae; UF, uncinate fasciculi; AF, arcuate fasciculi; OT, occipitotemporal fasciculi; and CB, cingulum bundles. The arcuate measurements were made in one bilateral cluster containing the traditional C-shaped fibers, and did not include additional clusters contained in the full arcuate/SLF region shown in Figure 8-7.

Chapter 9

Conclusion

The work presented in this thesis addresses the problem of automatic organization and anatomical labeling of the white matter using diffusion tensor MRI tractography. We present an automatic spectral clustering method which finds corresponding white matter regions across subjects and performs matching of tractography across hemispheres. We then show how to use spectral embedding and clustering to learn a fiber tract model which is represented in a high-dimensional embedding space. By annotating each cluster in the space with anatomical information, we create a high-dimensional tractography atlas.

We demonstrate that by embedding tractography from novel subjects into the atlas space, the atlas can be used for automatic segmentation of tractography. The automatic segmentation provides a consistent way to organize tractography across subjects, and the anatomical information allows visualization or analysis of entire structures (made up of many clusters) as well as the clusters themselves. We present results demonstrating the cluster/anatomical region correspondence across subjects, as well as an application of our method to parcellation of the region of the corpus callosum in the midsagittal plane. Our results indicate that our automatic segmentation method repeatably defines regions of interest that enable measurement of diffusion values such as FA; the results are similar across normal subjects and reproducible across scans of the same subject. We also demonstrate the use of the single-subject clustering method for neurosurgical visualization.

To our knowledge, this work represents the first automatic whole-brain tractography segmentation method, as well as the first method for automatically finding tract correspondences across subjects, and the first method for automatically finding tract correspondences across hemispheres. In addition, we show how to embed novel data in the spectral clustering framework (Normalized Cuts with the Nystrom method) of Fowlkes and Malik. This is the first application of spectral clustering to the problem of learning an anatomical model or atlas.

Our work enables statistical white matter analysis using tractography. We find corresponding white matter regions across subjects and across hemispheres, allowing neuroscientific hypotheses to be tested regarding group differences and also questions of symmetry.

Appendix A

Data Acquisition

Population I

We thank Dr. Martha Shenton's Psychiatry Neuroimaging Laboratory (PNL) at Brigham and Women's Hospital, Harvard Medical School, for this data. MR diffusion scans were performed with a quadrature head coil on a 1.5 T GE Echospeed system (General Electric Medical Systems, Milwaukee, WI), which permits maximum gradient amplitudes of 40 mT/m. Coronal LSDI scans were acquired perpendicular to both the AC-PC line and interhemispheric fissure. To increase the precision of the acquisition alignment, instead of one 3D localizer, a set of three 2D T1-weighted localizers (sagittal, axial oblique aligned to the anterior commissure-posterior commissure (AC-PC) line, and another sagittal oblique aligned to the interhemispheric fissure) were acquired. Finally, the last sagittal oblique T1W image served as the localizer for the LSDI coronal scans. For each section, six images with high (1000 s/mm²) diffusion-weighting along six non-collinear directions [e.g., relative amplitudes, $(G_x, G_y, G_z) = (1,1,0), (0,1,1), (1,0,1), (1,1,0), (0, 1,1)(1,0,1)$] and two with low (5 s/mm²) diffusion-weighting have been collected. The following scan parameters were used: rectangular FOV (field of view) 220 165 mm²; 128 128 scan matrix (256 256 image matrix); slice thickness 4 mm; interslice distance 1 mm; receiver bandwidth T4 kHz; TE (echo time) 64 ms; effective TR (repetition time) 2592 ms; scan time 60 s/slice section. A total of 31-35 coronal slices covering the entire brain (depending

upon brain size) were acquired. The total scan time was 31–35 min. Henceforth this data will be referred to as population I.

Population II

We thank Dr. Susumu Mori of Johns Hopkins University for this data, which was acquired as follows. DTI images were acquired using a SENSE head coil on 1.5 Tesla whole-body MR scanners (Philips Medical Systems, gyroscan NT) equipped with explorer gradients (40 mT/m). For acquisition, an eight-element arrayed RF coil, converted to six-channel to be compatible with the six-channel receiver, was used. For DTI acquisitions, a single-shot spin echo –echo planar sequence (SE-EPI) was used, with diffusion gradients applied in 30 non-collinear directions and $b = 700$ s/mm². Five additional reference image with least diffusion weighting ($b = 33$ s/mm²) was also acquired. Fifty to sixty axial slices were acquired to cover the entire hemisphere and the cerebellum, parallel to the AC-PC line. The field of view, the size of the acquisition matrix, and the slice thickness were 240 x 240 mm/96x96/2.5 mm. X and Y resolutions were then zero-filled to 256x256. Other imaging parameters were: TR 7,000 ms and TE = 80 ms; and SENSE reduction factor = 2.5. To improve the signal-to-noise ratio, three datasets were acquired, leading to a total acquisition time of 12–15 minutes. T2-weighted images with TE = 40 and 100 ms were separately acquired with the same EPI acquisition scheme to ensure the accurate co-registration. For anatomical image, MPRAGE with higher resolution (1.25 mm isotropic) was also acquired. Henceforth this data will be referred to as population II.

Bibliography

- [1] *Brain and Spinal Cord Tumors*. National Institute of Neurological Disorders and Stroke, 1993.
- [2] P. Basser. Inferring microstructural features and the physiological state of tissues from diffusion-weighted images. *NMR in Biomedicine*, 8:333–344, 1995.
- [3] P. Basser, S. Pajevic, C. Pierpaoli, J. Duda, and A. Aldroubi. In vivo fiber tractography using DT–MRI data. *Magnetic Resonance in Medicine*, 44:625–632, 2000.
- [4] P. Basser and C. Pierpaoli. Microstructural and physiological features of tissues elucidated by quantitative-diffusion-tensor MRI. *J. Magn. Reson. Ser. B*, 111:209–219, 1996.
- [5] P. J. Basser and D. K. Jones. Diffusion-tensor mri: theory, experimental design and data analysis – a technical review. *NMR in Biomedicine*, 15:456–467, 2002.
- [6] P. J. Basser, J. Mattiello, and D. LeBihan. Estimation of the effective self-diffusion tensor from the nmr spin echo. *Journal of Magnetic Resonance, Series B* 103:247–254, 1994.
- [7] P. J. Basser, J. Mattiello, and D. LeBihan. Mr diffusion tensor spectroscopy and imaging. *Biophysical Journal*, 66:259–267, 1994.
- [8] P. Batchelor, D. Hill, F. Calamante, and D. Atkinson. Study of connectivity in the brain using the full diffusion tensor from MRI. In *Information Processing in Medical Imaging, 17th International Conference (IPMI)*, pages 121–133, 2001.

- [9] C. Beaulieu. The basis of anisotropic water diffusion in the nervous system - a technical review. *NMR IN BIOMEDICINE*, 15:435–455, 2001.
- [10] T. Behrens. *MR Diffusion Tractography: Methods and Applications*. PhD thesis, University of Oxford, 2004.
- [11] T. Behrens, H. Johansen-Berg, M. Woolrich, S. Smith, C. Wheeler-Kingshott, P. Boulby, G. Barker, E. Sillery, K. Sheehan, O. Ciccarelli, A. Thompson, J. Brady, and P. Matthews. Non-invasive mapping of connections between human thalamus and cortex using diffusion imaging. *Nature Neuroscience*, 6:750–757, 2003.
- [12] T. Behrens, M. Woolrich, M. Jenkinson, H. J.-B. and R.G Nunes, S. Clare, P. Matthews, J. Brady, and S. Smith. Characterisation and propagation of uncertainty in diffusion weighted mr imaging. *Magn. Reson Med*, 50:1077–1088, 2003.
- [13] M. Belkin and P. Niyogi. Laplacian eigenmaps and spectral techniques for embedding and clustering. In T. G. Dietterich, S. Becker, , and Z. Ghahraman, editors, *Advances in Neural Information Processing Systems*, volume 14, Cambridge, MA, 2002. MIT Press.
- [14] S. Belongie, C. Fowlkes, F. Chung, and J. Malik. Spectral partitioning with indefinite kernels using the nystrom extension. In A. Hayden, editor, *Proc. European Conf. Comp. Vision*, volume LNCS 2352, Berlin Heidelberg, 2005. Springer-Verlag.
- [15] Y. Bengio, J.-F. Paiement, , P. Vincent, O. Delalleau, N. Le Roux, and M. Ouimet. Out-of-sample extensions for lle, isomap, mds, eigenmaps, and spectral clustering. In *Advances in Neural Information Processing Systems*, volume 16, 2004.

- [16] M. Björnemo, A. Brun, R. Kikinis, and C.-F. Westin. Regularized stochastic white matter tractography using diffusion tensor mri. In *MICCAI*, Tokyo, Japan, September 2002.
- [17] M. Brand and K. Huang. A unifying theorem for spectral embedding and clustering. In C. M. Bishop and B. J. Frey, editors, *Proceedings of the Ninth International Workshop on Artificial Intelligence and Statistics*, 2003.
- [18] A. Brun, H. Knutsson, H.-J. Park, M. E. Shenton, and C.-F. Westin. Clustering fiber traces using normalized cuts. In *Proc. Medical Image Conference and Computer Assisted Interventions (MICCAI)*, pages 368–375, 2004.
- [19] A. Brun, H.-J. Park, H. Knutsson, and C.-F. Westin. Coloring of DT-MRI fiber traces using Laplacian eigenmaps. In *European Conference on Computer Aided Systems Theory*, pages 564–572, 2003.
- [20] C. J. C. Burges. *Advances in Kernel Machines*, chapter 7 Geometry and Invariance in Kernel Based Methods. The MIT Press, 1999.
- [21] J. T. Bushberg, J. A. Seibert, J. Edwin M. Leidholt, and J. M. Boone. *The Essential Physics of Medical Imaging*. Williams and Wilkins, Baltimore, 1994.
- [22] J. Campbell, K. Siddiqi, B. Vemuri, and G. Pike. A geometric flow for white matter fibre tract reconstruction. In *IEEE International Symposium on Biomedical Imaging*, pages 505–508, 2002.
- [23] J. S. Campbell, K. Siddiqi, V. V. Rymar, A. F. Sadikot, and G. B. Pike. Flow-based fiber tracking with diffusion tensor and q-ball data: Validation and comparison to principal diffusion direction techniques. *NeuroImage*, 27:725–736, 2005.
- [24] M. Catani, R. J. Howard, S. Pajevic, and D. K. Jones. Virtual in vivo interactive dissection of white matter fasciculi in the human brain. *NeuroImage*, 17:77–9, 2002.

- [25] M. Catani, D. K. Jones, R. Donato, and D. H. ffytche. Occipito-temporal connections in the human brain. *Brain*, 126:2093–2107, 2003.
- [26] M. Catani, D. K. Jones, and D. H. ffytche. Perisylvian language networks of the human brain. *Annals of Neurology*, 57:8–16, 2004.
- [27] O. Ciccarelli, A. T. Toosy, G. J. M. Parker, C. A. M. Wheeler-Kingshott, G. J. Barker, D. H. Miller, and A. J. Thompson. Diffusion tractography based group mapping of major white-matter pathways in the human brain. *Neuroimage*, 19(4):1545–1555, 2003.
- [28] L. Concha, C. Beaulieu, and D. W. Gross. Bilateral limbic diffusion abnormalities in unilateral temporal lobe epilepsy. *Annals of Neurology*, 57(2):188 – 196, 2004.
- [29] L. Concha, D. W. Gross, , and C. Beaulieu. Diffusion tensor tractography of the limbic system. *American Journal of Neuroradiology*, 26:2267–2274, 2005.
- [30] T. E. Conturo, N. F. Lori, T. S. Cull, E. Akbuda, A. Z. Snyder, J. S. Shimony, R. C. McKinstry, H. Burton, and A. E. Raichle. Tracking neuronal fiber pathways in the living human brain. *Neurobiology*, 96:10422–10427, 1999.
- [31] P. A. Cook, H. Z. 0005, B. B. Avants, P. A. Yushkevich, D. C. Alexander, J. C. Gee, O. Ciccarelli, and A. J. Thompson. An automated approach to connectivity-based partitioning of brain structures. In J. S. Duncan and G. Gerig, editors, *Proc. Medical Image Conference and Computer Assisted Interventions (MICCAI)*, Lecture Notes in Computer Science, pages 164–171. Springer, 2005.
- [32] I. Corouge, P. T. Fletcher, S. Joshi, J. H. Gilmore, and G. Gerig. Fiber tract-oriented statistics for quantitative diffusion tensor mri analysis. In J. S. Duncan and G. Gerig, editors, *Proc. Medical Image Conference and Computer Assisted Interventions (MICCAI)*, Lecture Notes in Computer Science, pages 131–139. Springer, 2005.

- [33] I. Corouge, S. Gouttard, and G. Gerig. Towards a shape model of white matter fiber bundles using diffusion tensor MRI. In *International Symposium on Biomedical Imaging*, pages 344–347, 2004.
- [34] T. Delmarcelle and L. Hesselink. Visualization of second order tensor fields and matrix data. In *Proceedings of IEEE Visualization '92*, pages 316–323, 1992.
- [35] Z. Ding, J. C. Gore, and A. W. Anderson. Classification and quantification of neuronal fiber pathways using diffusion tensor MRI. *Magnetic Resonance in Medicine*, 49:716–721, 2003.
- [36] R. F. Dougherty, M. Ben-Shachar, R. Bammer, A. A. Brewer, and B. A. Wandell. Functional organization of human occipital-callosal fiber tracts. *PNAS*, 102(20):7350–7355, 2005.
- [37] P. Drineas and M. W. Mahoney. On the nystrom method for approximating a gram matrix for improved kernel-based learning. *Journal of Machine Learning Research*, 6:2153–2175, 2005.
- [38] D. B. Ennis and G. Kindlmann. Orthogonal tensor invariants and the analysis of diffusion tensor magnetic resonance images. *Magnetic Resonance in Medicine*, 55:136–146, 2006.
- [39] D. A. Forsyth and J. Ponce. *Computer Vision a Modern Approach*. Prentice Hall, 2003.
- [40] C. Fowlkes, S. Belongie, F. Chung, and J. Malik. Spectral grouping using the Nystrom method. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 26(2):214–225, 2004.
- [41] O. Friman, G. Farneback, and C.-F. Westin. A bayesian approach for stochastic white matter tractography. *TMI*, 2006.
- [42] O. Friman and C.-F. Westin. Uncertainty in fiber tractography. In *Eighth International Conference on Medical Image Computing and Computer-Assisted*

- Intervention (MICCAI'05)*, Lecture Notes in Computer Science 3749, pages 107–114, Palm Springs, CA, USA, October 2005.
- [43] G. Gerig, S. Gouttard, and I. Corouge. Analysis of brain white matter via fiber tract modeling. In *IEEE Engineering in Medicine and Biology Society Conference*, page 426, 2004.
- [44] D. Gering, A. Nabavi, R. Kikinis, N. Hata, L.O'Donnell, W. E. L. Grimson, F. Jolesz, P. Black, and W. W. III. An integrated visualization system for surgical planning and guidance using image fusion and an open mr. *Journal of Magnetic Resonance Imaging*, 13:967–975, 2001.
- [45] G. Gong, T. Jiang, C. Zhu, Y. Zang, F. Wang, S. Xie, J. Xiao, and X. Guo. Asymmetry analysis of cingulum based on scale-invariant parameterization by diffusion tensor imaging. *Human Brain Mapping*, 24(2):92 – 98, 2004.
- [46] H. Gray. *Gray's Anatomy*. Lea and Febiger, 20th u.s. edition edition, 1918. online edition: bartelby.com, 2000.
- [47] K. M. Hall. An r-dimensional quadratic placement algorithm. *Management Science*, 17(3):219–229, 1970.
- [48] J. R. Highley, M. A. Walker, M. M. E. Esiri, T. J. C. Crow, and P. J. H. Harrison. Asymmetry of the uncinate fasciculus: A post-mortem study of normal subjects and patients with schizophrenia. *Cerebral Cortex*, 12(11):1218–1224, 2002.
- [49] H. Huang, J. Zhang, H. Jiang, S. Wakana, L. Poetscher, M. I. Miller, P. C. van Zijl, A. E. Hillis, R. Wytik, and S. Mori. Dti tractography based parcellation of white matter: Application to the mid-sagittal morphology of corpus callosum. *NeuroImage*, 26(1):195–205, 2005.
- [50] B. J. Jellison, A. S. Field, J. Medow, M. Lazar, M. S. Salamat, and A. L. Alexander. Diffusion tensor imaging of cerebral white matter: A pictorial review of physics, fiber tract anatomy, and tumor imaging patterns. *American Journal of Neuroradiology*, 25:356–369, 2004.

- [51] H. Johansen-Berg, T. Behrens, M.D.Robson, I. Drobnjak, M. Rushworth, J. Brady, S. Smith, D. Higham, and P. Matthews. Changes in connectivity profiles define functionally distinct regions in human medial frontal cortex. *Proceedings of the National Academy of Sciences*, 101(36):13335–40, 2004.
- [52] L. Jonasson, P. Hagmann, J.-P. Thiran, and V. J. Wedeen. Fiber tracts of high angular resolution diffusion MRI are easily segmented with spectral clustering. In *International Society for Magnetic Resonance in Medicine*, 2005.
- [53] D. K. Jones. Determining and visualizing uncertainty in estimates of fiber orientation from diffusion tensor mri. *Magnetic Resonance in Medicine*, 49:7–12, 2003.
- [54] D. K. Jones, M. Catani, C. Pierpaoli, S. J. Reeves, S. S. Shergill, M. O’Sullivan, P. Golesworthy, P. McGuire, M. A. Horsfield, A. Simmons, S. C. Williams, and R. J. Howard. Age effects on diffusion tensor magnetic resonance imaging tractography measures of frontal cortex connections in schizophrenia. *Human Brain Mapping*, 2005.
- [55] D. K. Jones, A. R. Travis, G. Eden, C. Pierpaoli, and P. J. Basser. Pasta: Point-wise assessment of streamline tractography attributes. *Magnetic Resonance in Medicine*, 53:1462–1467, 2005.
- [56] R. A. Kanaan, S. S. Shergill, G. J. Barker, M. Catani, V. W. Ng, P. K. McGuire, and D. K. Jones. Tract-specific anisotropy measurements in diffusion tensor imaging,. *Psychiatry Research: Neuroimaging*, 146:73–82, 2006.
- [57] E. R. Kandel, J. H. Schwartz, and T. M. Jessel. *Principles of Neural Science*, chapter 2. McGraw-Hill, 2000.
- [58] G. Kindlmann. Superquadric tensor glyphs. In *IEEE Transactions on Visualization and Computer Graphics/EG Symposium on Visualization*, pages 147–154, 2004.

- [59] G. Kindlmann and D. Weinstein. Hue balls and lit-tensors for direct volume rendering of diffusion tensor fields. In *IEEE Visualization 1999, VIS1999*, Salt Lake City, Utah, October 8–13 1999.
- [60] M. Kinoshita, K. Yamada, N. Hashimoto, A. Kato, S. Izumoto, T. Baba, M. Maruno, T. Nishimura, and T. Yoshimine. Fiber-tracking does not accurately estimate size of fiber bundle in pathological condition: initial neurosurgical experience using neuronavigation and subcortical white matter stimulation. *Neuroimage*, 25(2):424–429, 2005.
- [61] V. E. Kouby, Y. Cointepas, C. Poupon, D. Rivière, N. Golestani, J.-B. Poline, D. L. Bihan, and J.-F. Mangin. Mr diffusion-based inference of a fiber bundle model from a population of subjects. In J. S. Duncan and G. Gerig, editors, *Proc. Medical Image Conference and Computer Assisted Interventions (MICCAI)*, Lecture Notes in Computer Science, pages 196–204. Springer, 2005.
- [62] M. Kubicki, H. Park, C. Westin, P. Nestor, R. Mulkern, S. Maier, M. Niznikiewicz, E. Connor, J. Levitt, M. Frumin, R. Kikinis, F. Jolesz, R. McCarley, and M. Shenton. Dti and mtr abnormalities in schizophrenia: Analysis of white matter integrity. *Neuroimage*, 26(1109–1118), 2005.
- [63] M. Kubicki, C.-F. Westin, S. E. Maier, M. Frumin, P. G. Nestor, D. Salisbury, R. Kikinis, F. A. Jolesz, R. W. McCarley, and M. E. Shenton. Uncinate fasciculus findings in schizophrenia: A magnetic resonance diffusion tensor imaging study. *AmJPsych*, 159:813–820, 2002.
- [64] M. Kubicki, C.-F. Westin, S. E. Maier, H. Mamata, M. Frumin, H. Ernst-Hirschfeld, R. Kikinis, F. A. Jolesz, R. W. McCarley, and M. E. Shenton. Diffusion tensor imaging and its application to neuropsychiatric disorders. *Harvard Rev of Psychiatry*, 10:234–336, 2002.
- [65] M. Kubicki, C.-F. Westin, S. E. Maier, H. Mamata, M. Frumin, H. Ernst-Hirschfeld, R. Kikinis, F. A. Jolesz, R. W. McCarley, and M. E. Shenton. Cin-

- gulate fasciculus integrity disruption in schizophrenia: A magnetic resonance diffusion tensor imaging study. *Biological Psychiatry*, 54:1171–1180, 2003.
- [66] M. Kubicki, C.-F. Westin, R. McCarley, and M. E. Shenton. The application of dti to investigate white matter abnormalities in schizophrenia. *Ann NY Acad Sci*, 1064:134–148, 2005.
- [67] M. Lazar, A. S. Field, J. H. Lee, and A. L. Alexander. Lateral asymmetry of superior longitudinal fasciculus: A white matter tractography study. In *International Society for Magnetic Resonance in Medicine*, page 1290, 2004.
- [68] M. Lazar, D. Weinstein, J. Tsuruda, K. Hasan, K. Arfanakis, M. Meyerand, B. Badie, H. Rowley, V. Haughton, A. Field, and A. Alexander. White matter tractography using diffusion tensor deflection. *Human Brain Mapping*, 18(4):306–321, 2003.
- [69] D. Le Bihan. Looking into the functional architecture of the brain with diffusion mri. *Nature Reviews. Neuroscience*, 4(6):469–480, 2003.
- [70] D. Le Bihan, J.-F. Mangin, C. Poupon, C. A. Clark, S. Pappata, N. Molko, and H. Chabriat. Diffusion tensor imaging: Concepts and applications. *Journal of Magnetic Resonance Imaging*, 13:534–546, 2001.
- [71] D. LeBihan, E. Breton, D. Lallemand, P. Grenier, E. Cabanis, and M. Laval-Jeantet. Mr imaging of intravoxel incoherent motions: application to diffusion and perfusion in neurologic disorders. *Radiology*, 161:401–407, 1986.
- [72] M. Maddah, W. Grimson, and S. Warfield. Statistical modeling and em clustering of white matter fiber tracts. In *Third IEEE International Symposium on Biomedical Imaging*, 2006.
- [73] M. Maddah, A. Mewes, S. Haker, W. E. L. Grimson, and S. Warfield. Automated atlas-based clustering of white matter fiber tracts from dtmri. In *Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, pages 188 – 195, 2005.

- [74] N. Makris, J. W. Meyer, J. F. Bates, E. H. Yeterian, D. N. Kennedy, and V. S. Caviness. Mri-based topographic parcellation of human cerebral white matter and nuclei ii. rationale and applications with systematics of cerebral connectivity. *Neuroimage*, 9(1):18–45, 1999.
- [75] A. D. McNaught and A. Wilkinson. *IUPAC Compendium of Chemical Terminology*. Blackwell Science, 1997.
- [76] A. Mead. Review of the development of multidimensional scaling methods. *The Statistician*, 1:27–39, 1992.
- [77] M. Meila and J. Shi. A random walks view of spectral segmentation. In *AI and Statistics (AISTATS)*, 2001.
- [78] B. Moberts, A. Vilanova, and J. J. van Wijk. Evaluation of fiber clustering methods for diffusion tensor imaging. In *IEEE Transactions on Visualization and Computer Graphics*, page 9, 2005.
- [79] S. Mori, B. Crain, V. Chacko, and P. van Zijl. Three dimensional tracking of axonal projections in the brain by magnetic resonance imaging. *Ann Neurol*, 45:265–269, 1999.
- [80] S. Mori, S. Wakana, L. M. Nagae-Poetscher, and P. C. van Zijl. *MRI Atlas of Human White Matter*. Elsevier, 2005.
- [81] M. Nakamura, R. W. McCarley, M. Kubicki, C. C. Dickey, M. A. Niznikiewicz, M. M. Voglmaier, L. J. Seidman, S. E. Maier, C.-F. Westin, R. Kikinis, and M. E. Shenton. Fronto-temporal disconnectivity in schizotypal personality disorder: A diffusion tensor imaging study. *Biol Psychiatry*, 58:468–478, 2005.
- [82] A. Ng, M. Jordan, and Y. Weiss. On spectral clustering: Analysis and an algorithm. In *Neural Information Processing Systems*, volume 14, 2001.
- [83] C. Nimsky, O. Ganslandt, P. Hastreiter, R. Wang, T. Benner, A. Sorensen, and R. Fahlbusch. Preoperative and intraoperative diffusion tensor imaging-based fiber tracking in glioma surgery. *Neurosurgery*, 56(1):130–137, 2005.

- [84] J. Nolte. *The Human Brain. An Introduction to Its Functional Anatomy*. Mosby, Inc., 2002.
- [85] L. O’Donnell, S. Haker, and C.-F. Westin. New approaches to estimation of white matter connectivity in diffusion tensor MRI: Elliptic PDEs and geodesics in a tensor-warped space. In T. Dohi and R. Kikinis, editors, *Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, pages 459–466, Tokyo, Japan, 2002.
- [86] L. O’Donnell, M. Kubicki, M. E. Shenton, W. E. L. Grimson, and C.-F. Westin. A method for clustering white matter fiber tracts. *American Journal of Neuroradiology*, 27(5):1032–1036, 2006.
- [87] L. O’Donnell and C.-F. Westin. White matter tract clustering and correspondence in populations. In *Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, pages 140–147, 2005.
- [88] L. O’Donnell and C.-F. Westin. A high-dimensional fiber tract atlas. In *International Society for Magnetic Resonance in Medicine*, 2006.
- [89] L. O’Donnell and C.-F. Westin. High-dimensional white matter atlas generation and group analysis. In *Proc. Medical Image Conference and Computer Assisted Interventions (MICCAI)*, 2006.
- [90] S. Pajevic and C. Pierpaoli. Color schemes to represent the orientation of anisotropic tissues from diffusion tensor data: application to white matter fiber tract mapping in the human brain. *Magnetic Resonance in Medicine* 2000, 42(3):526–540, 1999.
- [91] H. J. Park, C.-F. Westin, M. Kubicki, S. E. Maier, M. Niznikiewicz, A. Baer, M. Frumin, R. Kikinis, F. A. Jolesz, R. W. McCarley, and M. E. Shenton. White matter hemisphere asymmetries in healthy subjects and in schizophrenia: A diffusion tensor MRI study. *Neuroimage*, 24:213–223, 2004.

- [92] G. J. Parker, C. A. Wheeler-Kingshott, and G. J. Barker. Estimating distributed anatomical connectivity using fast marching methods and diffusion tensor imaging. *IEEE Transactions on Medical Imaging*, 21(5):505–512, 2002.
- [93] S. C. Partridge, P. Mukherjee, J. I. Berman, R. G. Henry, S. P. Miller, Y. Lu, O. A. Glenn, D. M. Ferriero, A. J. Barkovich, and D. B. Vigneron. Tractography-based quantitation of diffusion tensor imaging parameters in white matter tracts of preterm newborns. *Journal of Magnetic Resonance in Medicine*, 22(4):467–474, 2005.
- [94] C. Pierpaoli, P. Jezzard, P. J. Basser, A. Barnett, and G. D. Chiro. Diffusion tensor MR imaging of the human brain. *Radiology*, 201:637, 1996.
- [95] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery. *Numerical Recipes in C: the Art of Scientific Computing*, chapter 16.1 Runge-Kutta Method. Cambridge University Press, 1992.
- [96] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery. *Numerical Recipes in C: the Art of Scientific Computing*, chapter 18.1 Fredholm Equations of the Second Kind. Cambridge University Press, 1992.
- [97] D. Salat, A. Ward, J. A. Kaye, and J. S. Janowsky. Sex differences in the corpus callosum with aging. *Neurobiology of Aging*, 18(2):191–197, 1997.
- [98] D. H. Salat, D. S. Tuch, D. N. Greve, A. J. van der Kouwe, N. D. Hevelone, A. K. Zaleta, B. R. Rosen, B. Fischl, S. Corkin, H. D. Rosas, and A. M. Dale. Age-related alterations in white matter microstructure measured by diffusion tensor imaging. *Neurobiol Aging*, 26:1215–1227, 2005.
- [99] D. H. Salat, D. S. Tuch, N. D. Hevelone, A. K. Zaleta, S. Corkin, J. H. Growdon, and H. D. Rosas. White matter changes in ad measured with diffusion tensor imaging. In *Human Brain Mapping*, 2005.

- [100] G. Schlaug, B. Siewert, A. Benfield, R. R. Edelman, and S. Warach. Time course of the apparent diffusion coefficient (ADC) abnormality in human stroke. *Neurology*, 49(1):113–119, 1997.
- [101] B. Scholkopf, A. Smola, and K.-R. Muller. *Advances in Kernel Machines*, chapter 20 Kernel Principal Component Analysis. The MIT Press, 1999.
- [102] A. Schwartzman, R. F. Dougherty, and J. E. Taylor. Cross-subject comparison of principal diffusion direction maps. *Magnetic Resonance in Medicine*, 53:1423–1431, 2005.
- [103] P. N. Sen and P. J. Basser. A model for diffusion in white matter in the brain. *Biophysical Journal*, 89:2927–2938, 2005.
- [104] J. Shi and J. Malik. Normalized cuts and image segmentation. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 22(8):888–905, 2000.
- [105] S. M. Smith, M. Jenkinson, H. Johansen-Berg, D. Rueckert, T. E. Nichols, C. Mackay, K. E. Watkins, O. Ciccarelli, M. Z. Cader, P. M. Matthews, and T. E. Behrens. Tract-based spatial statistics: Voxelwise analysis of multi-subject diffusion data. Technical Report TR05SS1, Oxford (FMRIB Analysis Group), 2005.
- [106] G. J. Stanley Lu, Daniel Ahn and S. Cha. Peritumoral diffusion tensor imaging of high-grade gliomas and metastatic brain tumors. *American Journal of Neuroradiology*, 24:937–941, 2003.
- [107] G. Strang. *Introduction to Applied Mathematics*. Wellesley-Cambridge Press, 1986.
- [108] G. Strang. *Introduction to Linear Algebra*. Wellesley-Cambridge Press, 1998.
- [109] M. A. Styner, I. Oguz, R. G. Smith, C. Cascio, and M. Jomier. Corpus callosum subdivision based on a probabilistic model of inter-hemispheric connectivity.

- In J. S. Duncan and G. Gerig, editors, *Proc. Medical Image Conference and Computer Assisted Interventions (MICCAI)*, pages 765–772, 2005.
- [110] J. Talairach and P. Tournoux. *Co-planar Stereotaxic Atlas of the Human Brain: 3-Dimensional Proportional System - an Approach to Cerebral Imaging*. Thieme Medical Publishers, New York, NY, 1988.
- [111] I.-F. Talos, L. O’Donnell, C.-F. Westin, S. K. Warfield, W. W. III, S.-S. Yoo, L. P. Panych, A. Golby, H. Mamata, S. S. Maier, P. Ratiu, C. R. Guttmann, P. M. Black, F. A. Jolesz, and R. Kikinis. Diffusion tensor and functional MRI fusion with anatomical MRI for image-guided neurosurgery. In *Proc. Medical Image Conference and Computer Assisted Interventions (MICCAI)*, pages 407–415, Toronto, Canada, 2003.
- [112] D. S. Tuch. *Diffusion MRI of Complex Tissue Structure*. PhD thesis, Division of Health Sciences and Technology, Massachusetts Institute of Technology, 2002.
- [113] D. S. Tuch. High angular resolution diffusion imaging reveals intravoxel white matter fiber heterogeneity. *Magn. Reson. Med.*, 48:577–582, 2002.
- [114] D. S. Tuch. Q-ball imaging. *Magnetic Resonance in Medicine*, 52:1358–1372, 2004.
- [115] D. S. Tuch, D. H. Salat, J. J. Wisco, A. K. Zaleta, N. D. Hevelone, and H. D. Rosas. Choice reaction time performance correlates with diffusion anisotropy in white matter pathways supporting visuospatial attention. *Proc Natl Acad Sci USA*, 102:12212–12217, 2005.
- [116] S. Wakana, H. Jiang, L. M. Nagae-Poetscher, P. C. M. van Zijl, and S. Mori. Fiber tract-based atlas of human white matter anatomy. *Radiology*, 230:77–87, 2004.
- [117] X. Wang, K. Tieu, and E. Grimson. Learning semantic scene models by trajectory analysis. Technical report, MIT CSAIL, 2006.

- [118] D. Weinstein, G. Kindlmann, and E. Lundberg. Tensorlines: Advection-diffusion based propagation through diffusion tensor fields. In *Proceedings of IEEE Visualization '99*, pages 249–253, 1999.
- [119] Y. Weiss. Segmentation using eigenvectors: A unifying view. In *ICCV (2)*, pages 975–982, 1999.
- [120] E. W. Weisstein. Injection. From MathWorld—A Wolfram Web Resource. <http://mathworld.wolfram.com/Injection.html>.
- [121] E. W. Weisstein. Osculating circle. From MathWorld—A Wolfram Web Resource. <http://mathworld.wolfram.com/OsculatingCircle.html>.
- [122] C.-F. Westin, S. E. Maier, B. Khidhir, P. Everett, F. A. Jolesz, and R. Kikinis. Image processing for diffusion tensor magnetic resonance imaging. In *Medical Image Computing and Computer-Assisted Intervention*, Lecture Notes in Computer Science, pages 441–452, September 19–22 1999.
- [123] C.-F. Westin, S. E. Maier, H. Mamata, A. Nabavi, F. A. Jolesz, and R. Kikinis. Processing and visualization of diffusion tensor MRI. *Medical Image Analysis*, 6(2):93–108, 2002.
- [124] C.-F. Westin, S. Peled, H. Gudbjartsson, R. Kikinis, and F. A. Jolesz. Geometrical diffusion measures for MRI from tensor basis analysis. In *ISMRM 97*, Vancouver, Canada, 1997.
- [125] Wikipedia. Brownian motion — wikipedia, the free encyclopedia, 2006. [Online; accessed 27-March-2006].
- [126] Wikipedia. Fick’s law of diffusion — wikipedia, the free encyclopedia, 2006. [Online; accessed 27-March-2006].
- [127] C. K. I. Williams and M. Seeger. Using the nystrom method to speed up kernel machines. In *Advances in Neural Information Processing Systems*, 2001.

- [128] T. H. Williams, N. Gluhbegovic, and J. Y. Jew. *The Human Brain: Dissections of the Real Brain*. Virtual Hospital, The University of Iowa, College of Medicine, 1997. <http://www.vh.org/adult/provider/anatomy/BrainAnatomy/BrainAnatomy.html>.
- [129] S. F. Witelson. Hand and sex differences in the isthmus and genu of the human corpus callosum a postmortem morphological study. *Brain*, 112(3):799–835, 1989.
- [130] B. Witwer, R. Moftakhar, K. Hasan, P. Deshmukh, V. Haughton, A. Field, K. Arfanakis, J. Noyes, C. Moritz, M. Meyerand, H. Rowley, A. Alexander, and B. Badie. Diffusion-tensor imaging of white matter tracts in patients with cerebral neoplasm. *J Neurosurg*, 2002.
- [131] Y. Xia, U. Turken, S. L. Whitfield-Gabrieli, and J. D. Gabrieli. Knowledge-based classification of neuronal fibers in entire brain. In J. S. Duncan and G. Gerig, editors, *Proc. Medical Image Conference and Computer Assisted Interventions (MICCAI)*, Lecture Notes in Computer Science, pages 205–212. Springer, 2005.
- [132] D. Xu, S. Mori, M. Solaiyappan, P. C. M. van Zijl, and C. Davatzikos. A framework for callosal fiber distribution analysis. *NeuroImage*, 17:1131–1143, 2002.
- [133] S. Zhang, S. Correia, D. F. Tate, and D. H. Laidlaw. Correlating dti fiber clusters with white matter anatomy. In *International Society for Magnetic Resonance in Medicine*, 2006.
- [134] S. Zhang and D. H. Laidlaw. DTI fiber clustering and cross-subject cluster analysis. In *International Society for Magnetic Resonance in Medicine*, 2005.
- [135] L. Zollei, E. Learned-Miller, W. E. L. Grimson, and W. M. Wells III. Efficient population registration of 3D data. In *ICCV 2005, Computer Vision for Biomedical Image Applications*, 2005.