Abstract of "**Computational Brain Connectivity Using Diffusion MRI**" by Çağatay Demiralp, Ph.D., Brown University, May 2013.

This dissertation shows that qualitative and quantitative characterization of patterned structures in brain connectivity data obtained using diffusion MRI not only improves the exploration of the intricate space of brain connectivity but also provides clinically meaningful measures, quantifying normal and pathological variation in the brain. To this end, we introduce a set of computational and mathematical models, algorithms, and interactive tools to explore, understand, and characterize diffusion-derived structural brain connectivity.

We contribute to all stages of modeling, visualization, and analysis of brain connectivity. In modeling, our contributions are twofold. First, we model the joint distribution of local neural fiber configurations with Markov random fields and infer the most likely configuration with maximum a posteriori estimation. We demonstrate this framework's use in resolving fiber crossings. Second, we introduce new planar map representations of three-dimensional neural tract datasets. These planar representations improve the exploration of brain connectivity by reducing visual and interaction complexity.

In visualization, we contribute to structure-preserving color mappings. First, we introduce Boy's surface as a model for coloring 3D line fields and show results from its application in visualizing orientation in diffusion MRI brain datasets. This coloring method is smooth and one-to-one except on a set of measure zero. Second, we propose a general coloring method based on manifold embedding that conveys spatial relations among neural fiber tracts perceptually. We also introduce a new bivariate coloring model, the flat torus, that allows finer adjustments of coloring arbitrarily.

We contribute to both local and global analysis of brain connectivity. In local analysis, we introduce a geometric slicing-based coherence measure for clusters of neural tracts. Clustering refinement based on this measure leads to a significant improvement in clustering quality that is not possible directly with standard methods. We also introduce tract-based probability density functions and demonstrate their effective use in nonparametric hypothesis testing and classification. In global analysis, we propose computing the ranks of persistent homology groups in the neural tract space. This captures the effects of diffuse axonal dropout and provides a global descriptor of structural brain connectivity.

Computational Brain Connectivity Using Diffusion MRI

by
Çağatay Demiralp
B. Sc., Ege University, 2000
Sc. M., Brown University, 2004

A dissertation submitted in partial fulfillment of the
requirements for the Degree of Doctor of Philosophy
in the Department of Computer Science at Brown University

Providence, Rhode Island
May 2013

This dissertation by Çağatay Demiralp is accepted in its present form by
the Department of Computer Science as satisfying the dissertation requirement
for the degree of Doctor of Philosophy.

Date _____ _____

David H. Laidlaw, Director

Recommended to the Graduate Council

Date _____ _____

John F. Hughes, Reader

Date _____ _____

David B. Mumford, Reader

Approved by the Graduate Council

Date _____ _____

Peter M. Weber, Dean of the Graduate School

Vita

Education

Ph.D., Computer Science, Brown University, Providence, RI (2006 - present)
Thesis Title: *Computational Brain Connectivity Using Diffusion MRI*
Advisor: *David H. Laidlaw*
Committee: *John F. Hughes and David B. Mumford*

Sc.M., Computer Science, Brown University, Providence, RI (2004)

B.Sc., Computer Engineering, Ege University, İzmir, Turkey (2000)

Research Experience

Research Intern, Machine Learning and Perception Group, Microsoft Research Cambridge (Fall 2011)

Graduate Research Assistant, Computer Science Dept., Brown University (2006-present)

Graduate Research Fellow, NLM at the National Institutes of Health (Nov-Dec 2004)

Visiting Research Fellow, Centre for Neuroimaging Sciences, King's College London (Summer 2004)

Graduate Research Assistant, Computer Science Dept., Brown University (2002-2004)

Staff Researcher, Computer Science Dept., Brown University (2000-2002)

Visiting Undergraduate Research Assistant, Computer Science Dept., Brown University (May-Sep 1999)

Ugraduate Research Intern, Fraunhofer Center for Research in Computer Graphics (Nov 1998-May 1999)

Awards and Honors

Best Poster Award at IEEE Visualization Conference (2010)

Travel Award to attend Machine Learning Summer School/Workshop at the University of Chicago (2009)

Brain Science Graduate Research Award (2008)

Best Scientific Content and Best Layout and Presentation awards at the American Society for Surgery of the Hand 56th Annual Meeting (2001)

Book Chapters

1. Exploring Brain Connectivity with Two-dimensional Maps. **Ç. Demiralp**, R. Jianu, and D. H. Laidlaw. *New Developments in the Visualization and Processing of Tensor Fields*. LNCS-Springer, 2012.

Journal Publications

7. Developing a Generative Theory of Visualization. **Ç. Demiralp**, C. Scheidegger, G. L. Kindlmann, and David H. Laidlaw. *IEEE Computer Graphics & Applications*, 2012 (in review)
6. Exploring Brain Connectivity with Two-dimensional Neural Maps. R. Jianu, **Ç. Demiralp**, and D. H. Laidlaw. *IEEE Trans. Vis. Comput. Graphics*, 2011.
5. Coloring 3D Line Fields Using Boy's Real Projective Plane Immersion. **Ç. Demiralp**, J. F. Hughes, and D. H. Laidlaw. *IEEE Trans. Vis. Comput. Graphics (Proc. IEEE Visualization)*, 2009.
4. Exploring 3D DTI Fiber Tracts with Linked 2D Representations. R. Jianu, **Ç. Demiralp**, and D. H. Laidlaw. *IEEE Trans. Vis. Comput. Graphics (Proc. IEEE Visualization)*, 2009.
3. A Qualitative and Quantitative Comparison of CAVE and Fishtank Virtual-Reality Displays. **Ç. Demiralp**, C. D. Jackson, D. B. Karelitz, S. Zhang, and D. H. Laidlaw. *IEEE Trans. Vis. Comput. Graphics*, 2006.
2. In-vivo Measurement of Contact Areas and Ligament Lengths in the Distal Radioulnar Joint. G. E. Marai, D. H. Laidlaw, **Ç. Demiralp**, S. Andrews, C. M. Grimm, and J. J. Crisco. *IEEE Trans. Biomed. Eng.*, 2004.
1. Visualizing Diffusion Tensor MR Images Using Streamtubes and Streamsurfaces. S. Zhang, **Ç. Demiralp**, and D. H. Laidlaw. *IEEE Trans. Vis. Comput. Graphics*, 2003.

Conference Publications

10. Decision Forests for Tissue-specific Segmentation of High-grade Gliomas in Multi-channel MR. D. Zikic, B. Glocker, E. Konukoglu, A. Criminisi, **Ç. Demiralp**, J. Shatton, O. Thomas, T. Das, R. Jena, S. Price. *Proc. Med. Image. Comput. Comput. Assist. Interv. (MICCAI)*, 2012.
9. Context-sensitive Classification Forests for Segmentation of Brain Tumor Tissues. D. Zikic, B. Glocker, E. Konukoglu, A. Criminisi, J. Shatton, D. H. Ye, **Ç. Demiralp**, O. Thomas, T. Das, R. Jena, S. Price. *MICCAI 2012 Challenge on Multimodal Brain Tumor Segmentation*, 2012.
8. Generalizing Diffusion Tensor Model Using Probabilistic Inference in Markov Random Fields. **Ç. Demiralp** and David H. Laidlaw. *Proc. Med. Image. Comput. Comput. Assist. Interv. (MICCAI) Workshop on Computational Diffusion MRI*, 2011.

7. Tract-based Probability Densities of Diffusivity Measures in DT-MRI. **Ç. Demiralp** and D. H. Laidlaw. *Proc. Med. Image. Comput. Comput. Assist. Interv. (MICCAI)*, 2010.
6. Surface Deformations Driven by Vector-Valued 1-Forms. G. Taubin and **Ç. Demiralp**. *Proc. Shape Modeling International*, 2010.
5. Similarity Coloring of DTI Fiber Tracts. **Ç. Demiralp** and D. H. Laidlaw. *Proc. Med. Image. Comput. Comput. Assist. Interv. (MICCAI) Workshop on Diffusion Modeling and the Fibre Cup*, 2009.
4. Slicing-based Coherence Measure for Refining Clusters of 3D curves. **Ç. Demiralp**, G. Shakhnarovich, S. Zhang, and D. H. Laidlaw. *Proc. Med. Image. Comput. Comput. Assist. Interv. (MICCAI)*, 2008.
3. Connectivity-aware Sectional Visualization of 3D DTI Volumes Using Perceptual Flat-Torus Coloring and Edge Rendering. **Ç. Demiralp**, S. Zhang, D. F. Tate, S. Correia, D. H. Laidlaw. *Eurographics*, 2006.
2. An Immersive Virtual Environment for DT-MRI Volume Visualization Applications: A Case Study. S. Zhang, **Ç. Demiralp**, D. F. Keefe, M. J. da Silva, D. H. Laidlaw, B. D. Greenberg, P. J. Basser, E. A. Chiocca, C. Pierpaoli, T. S. Deisboeck. *Proc. IEEE Visualization*, 2001.
1. Application of Virtual Reality to Visualization of DT-MRI Volumes. S. Zhang, **Ç. Demiralp**, D. F. Keefe, M. J. da Silva, D. H. Laidlaw, B. D. Greenberg, P. J. Basser, E. A. Chiocca,C. Pierpaoli, T. S. Deisboeck. *Proc. Med. Image. Comput. Comput. Assist. Interv. (MICCAI)*, 2001.

Conference Abstracts and Posters

13. Generalizing Diffusion Tensor Model Using Probabilistic Inference in Markov Random Fields. **Ç. Demiralp** and David H. Laidlaw. *Proc. Intl. Soc. Mag. Reson. Med. (ISMRM)*, 2011.
12. Exploring Brain Connectivity with Two-Dimensional Neural Maps. R. Jianu, **Ç. Demiralp**, and D. H. Laidlaw. *Poster Compendium IEEE Visualization* (***Best Poster Award***), 2010.
11. Visualizing and Exploring Tractograms via Two-Dimensional Connectivity Maps. R. Jianu, **Ç. Demiralp**, and D. H. Laidlaw. *Proc. Intl. Soc. Mag. Reson. Med. (ISMRM)*, 2010.
10. Quantitative Tractography Metrics of White Matter Integrity in Diffusion-Tensor MRI Using Diffusivity Scalars. E. Halilaj, **Ç. Demiralp**, S. Correia, and D. H. Laidlaw. *Proc. Intl. Soc. Mag. Reson. Med. (ISMRM)*, 2009.
9. Using Boy's Real Projective Plane Immersion for Coloring DT-MRI Slices. **Ç. Demiralp**, J. F. Hughes and D. H. Laidlaw. *Proc. Intl. Soc. Mag. Reson. Med. (ISMRM)*, 2008.
8. Visualizing Spatial Relations Between 3D-DTI Integral Curves Using Texture Patterns. D. Jianu, W. Zhou, **Ç. Demiralp**, and D. H. Laidlaw. *Poster Compendium IEEE Visualization*, 2007.
7. JointViewer – An Interactive System for Exploring Orthopedic Data. G. E. Marai, **Ç. Demiralp**, S. Andrews, D. H. Laidlaw. *Poster Compendium IEEE Visualization*, 2004.
6. Subjective Usefulness of CAVE and Fishtank VR Display Systems for a Scientific Visualization Application. **Ç. Demiralp**, D. H. Laidlaw, C. Jackson, D. Keefe, and S. Zhang. *Poster Compendium IEEE Visualization*, 2003.

5. Contact Areas and Ligament Lengths Are Abnormal in Patients with Malunited Distal Radius Fracture Despite Normal Radioulnar Kinematics. G. E. Marai, D. H. Laidlaw, **Ç. Demiralp**, C. Grimm, J. J. Crisco, D. C. Moore, and E. Akelman. *Proc. 4th World Congress of Biomechanics*, 2002.
4. Visualizing the Differences Between Diffusion Tensor Volume Images. M. J. da Silva, S. Zhang, **Ç. Demiralp**, and D. H. Laidlaw. *Proc. Intl. Soc. Mag. Reson. Med. (ISMRM) Workshop on Diffusion MRI*, 2002.
3. Modeling and Visualization of Inter-Bone Distances in Joints. **Ç. Demiralp**, G. E. Marai, S. Andrews, D. H. Laidlaw, J. J. Crisco, and C. Grimm. *Proc. IEEE Visualization Work in Progress*, 2001.
2. Visualizing Diffusion Tensor Volume Differences. M. J. DaSilva, S. Zhang, **Ç. Demiralp**, and D. H. Laidlaw. *Proc. IEEE Visualization Work in Progres*, 2001.
1. Interactive Visualization of 3D Carpal Kinematics and Bony Anatomy. J. J. Crisco, **Ç. Demiralp**, D. H. Laidlaw, A-P. C. Weiss, E. Akelman, and S. W. Wolfe. *Scientific Exhibit at the American Surgery of the Hand 56th Annual Meeting* (***Best Scientific Content*** and ***Best Layout and Presentation*** awards), 2001.

Panels

Theories of Visualization–Are There Any? *IEEE VisWeek* (2011)

Invited Talks

· Computational Brain Connectivity Using Diffusion MRI, *A. Martinos Center, Harvard Medical School* (Sep 2012)

· Computational Brain Connectivity Using Diffusion MRI, *University of Maryland, Baltimore County* (Sep 2012)

· Cycles of Brain White Matter, *Dagstuhl Seminar* (Dec 2011)

· Exploring Brain Connectivity with Two-dimensional Neural Maps, *University College London* (Nov 2011)

· Coloring 3D Line Fields Using Boy's Real Projective Plane Immersion, *University of New Hampshire* (2010)

· Manifold Ways of Coloring, *Dagstuhl Seminar* (Jun 2009)

Teaching Experience

Teaching Assistant, Computational Topology, Brown University (Spring 2011)

Teaching Assistant, Discrete Mathematics, University of Utah (Fall 2005)

Teaching Assistant, Advanced Algorithms & Data Structures, University of Utah (Spring 2005)

Service

Program Committee Member, *MICCAI Workshop on Medical Computer Vision'12*

Organizer, VisWeek'11 Panel, *Theories of Visualization–Are There Any?*

Reviewer, *CHI, EuroVis, PacificVis, IEEE Visualization, IEEE Trans. Vis. Comput. Graphics, MICCAI*

Organizer, Brown SciVis Seminars by Mert R. Sabuncu, Ofer Pasternak, Miriah Meyer, Won-Ki Jeong, Peter Savadjiev

PhD Admission Committee Member, Computer Science Dept., Brown University (2011)

Faculty Search Czar, Computer Science Dept., Brown University (2009 - 2010)

Department LaTeX Administrator, Computer Science Dept., Brown University (2007 - 2011)

Outreach and Media Coverage

Exploring Brain Connectivity with Two-dimensional Neural Maps. **Ç. Demiralp**. *Visual Strategies: A Practical Guide to Graphics for Scientists and Engineers* by Felice C. Frankel and Angela H. DePace. Yale University Press, 2012.

"Researchers Map, Measure Brain's Neural Connections" at *ScienceDaily et al.*

"A New Way to Look at the Brain Using Google Maps" at *Gizmodo et al.*

Acknowledgments

This dissertation work is result of a long journey that could not have been undertaken without the support of many people. I would like to start with thanking my advisor David H. Laidlaw for the opportunity, freedom and support he provided. I thank also my committee members John F. Hughes and David B. Mumford for their time and feedback. Their genuine curiosity and intellectual clarity will always stay with me as inspiring examples.

Over the years, I have enjoyed collaborating with many people on numerous interesting problems. I would like to thank all my collaborators including Stuart Andrews, Stephen Correa, Antonio Criminisi, Joseph Crisco, Ben Glocker, Benjamin Greenberg, Cindy Grimm, Cullen Jackson, Dan F. Keefe, Radu Jianu, Ender Konukoğlu, Gregory Shakhnarovich, Marco da Silva, David Tate, Gabriel Taubin, and Song Zhang.

I have been lucky to be surrounded with role models, mentors, and friends alike at Brown and beyond. I would like to thank Andy van Dam for all his support when it really mattered, from finding a dentist after my debilitating bicycle accident to searching for a job. I also thank Mert Akdere, Thomas F. Banchoff, Uğur Çetintemel, Andrew Furnas, Stuart Geman, Kathryn A. Good, Elizabeth Grossman, James Hays, Hilal Kazan, Onur Keskin, Emrah Orhun, Franco P. Preparata, Mert Sabuncu, John Savage, Rosemary M. Simpson, Lawrance Stanley, Nesime Tatbul, Rebecca Watson, Frank Wood and Battalgazi Yıldırım. I should also thank past and present visualization group members for their feedback, support and putting up with my particular sense of humor.

Brown University in general and the computer science department in particular have been a home for me since my first visit as an undergraduate student from Turkey. I would like to thank all the administrative and technical staff for making the department such a welcoming place. Katrina Avery's editorial help has been key to success of many publications associated with this thesis. I thank her for the excellent job she has done.

In many ways we are products of our parents. I want to thank my parents for their trust and support. I also thank my brother for his friendship. The singular event of my PhD years by far was to become a father. It has been a life-altering experience. I am indebted to my lovely family for all their love and understanding. I dedicate this thesis to my beautiful daughter Alexandra and look forward to making up for the time we have been apart.

To Alexandra the great

Contents

List of Tables

List of Figures

Chapter 1

Introduction

The history of neuroscience is a testament to the fact that size and structure matter and are deeply related to function. In many ways, however, this relation is less understood in the human brain than in other organs [104]. The human brain, the largest part of the central nervous system, has billions of axons and trillions of synaptic connections and regulates all bodily functions. Its health is central to the length and quality of human life. According to the World Health Organization (WHO), neurological disorders ranging from epilepsy to dementia affect approximately a billion people worldwide and that number is expected to increase with aging populations [159]. Among these disorders, for example, cerebral strokes, multiple sclerosis (MS), Parkinson's disease (PD) and traumatic brain injuries cause an estimated 6.8 million deaths each year, accounting for 12% of global deaths.

Due to the complexity and sensitivity of the brain, its relative inaccessibility to drugs because of the blood-brain barrier, and the combinatorial nature of the causes, neurological disorders present an important scientific and public health challenge. There has thus been a vast amount of research into the health of the brain in general and into understanding the brain's neural architecture in particular. The latter is important not merely because the structural connectivity constrains and often determines what is functionally possible but also because most neurological disorders leave a footprint on brain connectivity.

Diffusion MRI, a relatively recent radiological imaging modality, provides an indirect probe into tissue microarchitecture and makes possible the noninvasive estimation of neural pathways (tracts) in the living brain. Its sensitivity to microstructural changes and ability to capture the geometry of the neural pathways opens new possibilities for revealing the neural architecture of the brain as well as for understanding, diagnosis and treatment of neurological disorders or injuries. The relation between diffusion-derived parameters in the brain and brain health has been established with several studies. Among the clinical applications of diffusion MRI are dementia, including AD and HIV-associated dementia, MS, amyotrophic lateral sclerosis (ALS), ischemic stroke, epilepsy, and brain tumor [22]. Its combination with other neuroimaging techniques, particularly functional imaging, is likely to expand these possibilities further. It is thus of great interest to develop effective tools for modeling, visualizing, and analyzing diffusion MRI data.

In this context, this dissertation establishes that qualitative and quantitative characterization of patterned

structures in diffusion MRI-derived brain connectivity data not only improves our understanding of the intricate space of brain connectivity but also provides us clinically meaningful measures that quantify normal and pathological variation in the brain. For this, we introduce a set of computational and mathematical models, algorithms, and interactive tools to explore, understand, and characterize structural brain connectivity obtained using diffusion-weighted magnetic resonance imaging. We contribute to all stages of modeling, visualization, and analysis of diffusion-derived brain connectivity. Chapter 2 gives an overview of neuroanatomy and related concepts to provide a background for the rest of the thesis. The chapters after Chapter 2 are presented in three parts: Modeling (I), Visualization (II), and Analysis (III). Of course, this division into three parts is not always clear cut: for example, an effective modeling of the data space often presumes a good visualization.

Inadequacy in modeling mixed tissue orientations in a given voxel is a well known limitation of the tensor model of diffusion. In Chapter 3, we give a proof of concept for efficiently modeling tensor configuration distributions with Markov random fields (MRFs) and inferring the most likely tensor configurations with maximum a posteriori (MAP) estimations. We demonstrate the plausibility of our method by resolving fiber crossings in a synthetic dataset, experimenting with three different MAP estimation methods on a grid MRF model.

As discussed above, effective data modeling and representation is often necessary for visualization and interactive exploration. Chapter 4 describes and compares two low-dimensional representations, 2D point and 2D path, for neural tracts. The goal is to facilitate exploration of dense tractograms (collection of curves representing neural tracts) by reducing visual complexity both in static representations and interaction. The proposed planar maps have desirable properties such as visual clarity, easy tract-of-interest selection, and multiscale hierarchy. The 2D path representation also conveys the anatomical familiarity of 3D brain models and cross-sectional views. We demonstrate the utility of the representations in two interactive systems in which the views and interactions of the standard 3D streamtube representation are linked to those of the planar representations. We also demonstrate a web interface that integrates precomputed neural-path representations into a geographical digital-maps framework with associated labels, metrics, statistics, and linkouts. We compare the two representations both anecdotally and quantitatively with the help of expert input. Results indicate that the planar path representation is more intuitive and easier to use and learn. Similarly, users are faster and more accurate in selecting bundles using the path representation than the 2D point representation. Finally, expert feedback on the web interface suggests that it can be useful for collaboration as well as quick exploration of data.

Visualization can facilitate exploration of brain connectivity by qualitatively characterizing structures, and color mapping is a basic technique in visualization. Intuitively, a good coloring of data points should map similar data points to similar colors and dissimilar data points to dissimilar colors. If the data points change smoothly, the color mapping should change smoothly as well. Chapters 5 and 6 present coloring techniques that formalize this intuition. In Chapter 5, we introduce a new method for coloring 3D line fields and show results from its application in visualizing neural tract and local tissue orientations in diffusion MRI brain data sets. The method uses Boy's surface, an immersion of the real projective plane in 3D. This coloring method is smooth and one-to-one except on a set of measure zero, the double curve of Boy's surface.

Chapter 6 introduces a coloring method that conveys spatial relations among neural fiber tracts effectively;

similar tracts are assigned similar colors and different tracts are assigned different colors in a smooth and continuous manner. To this end, we combine a standard spectral approach with a mass-spring heuristic to embed fiber tracts into a perceptually uniform color space, L*a*b*. We also introduce a new geometric bivariate coloring model, the flat torus, that allows finer adjustment of coloring arbitrarily. Results indicate that our method allows quick evaluation of tract projections (endpoints).

Clustering is a common technique for exploring and analyzing structures in neural tract datasets. Furthermore, the success of many analysis methods depends on the quality of the clustering results. Chapter 7 presents a slicing-based coherence measure for clusters of neural tracts. We probe samples from a given cluster by slicing it with a plane at regularly spaced locations parametrized by curve arc lengths. Then we compute a stability measure based on the spatial relations between the projections of the curve points in individual slices and their change across the slices. We demonstrate its use in refining agglomerative hierarchical clustering results of neural pathways. Expert evaluation shows that refinement based on our measure can lead to improvement of clustering that is not possible directly by using standard methods.

The effects of neurological disorders are either localized or diffuse (widespread, global). In Chapter 8, we introduce a method for global analysis of neural tracts. Topology offers a set of invariants that can help characterize data globally (qualitatively). The number of independent p-dimensional cycles, the rank of the pth homology group, in the space is one such invariant and can be computed efficiently using linear-algebraic methods. We present initial results from a multiscale computation of this topological invariant for neural pathways in the brain, providing a global characterization of brain connectivity not available through existing, essentially local descriptors. We show results from computation of this descriptor on brain datasets from patients with vascular cognitive disorder (VCI) and tumor.

Analysis of individual and group differences is a typical method in clinical research on neurological disorders. In Chapter 9, we introduce, as biomarkers, tract-based probability density functions, including the joint density of tract arc length and scalar diffusivity measures. We demonstrate their simple and effective use in detecting individual and group differences. We also show that probability density entropy itself is a useful biomarker.

In the last chapter, we summarize our contributions, provide conclusions and point to possible future directions.

Chapter 2

Background

This dissertation involves a set of computational methods applied to structural human brain connectivity. Therefore, we give an overview of neuroanatomy and relevant imaging techniques here to introduce the necessary background and context to make reading the ensuing chapters easier. We focus on major neural pathways and diffusion imaging due to their particular relevance to our work.

2.1 Structure of the Central Nervous System

The central nervous system (CNS) is one of the two anatomical subdivisions of the nervous system (NS). With the peripheral nervous system (PNS), it makes up the NS. The central nervous system (CNS) has two parts: the brain and the spinal cord. The brain is enclosed within the cranium–with the mandible, the cranium constitutes the skull–and is the largest and, in many ways, the most important component of the human CNS. The brain is not just center of the regulation bodily activities, including processing sensory information and controlling the muscles and organs; it is where consciousness, thought, memory, and emotion are formed.

2.1.1 Coordinate System for the CNS

When studying neuroanatomy (or general anatomy, for that matter), we need a spatial framework and associated terminology, a coordinate system, to refer locations of structures in a consistent manner. This need arises from the fact that daily terms of direction or location are often insufficient to describe structures of interest regardless of the orientation. Up and down, for example, makes sense only if the orientation of the head or body is known. We present here the location/direction axes used in the neuroanatomy. As customary, we use the rat CNS to illustrate the axes (Figure 2.2); these correspond to the human CNS axes with minor variations.

The axis that goes from the rat's nose to its spinal cord is called the rostrocaudal axis. The end of this axis towards the nose is called the rostral (or anterior) direction and the opposite direction is referred to as caudal (or posterior) direction. The axis perpendicular to the rostrocaudal axis and going from the back to the rat's abdomen is called the dorsoventral axis. The direction pointing up to the rat's back is called the dorsal

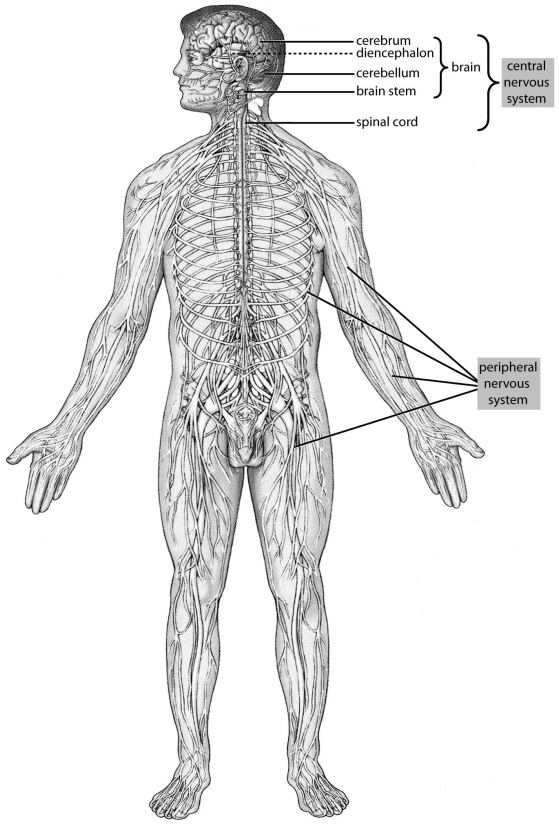

Figure 2.1: The central nervous system (CNS), with the peripheral nervous system, constitutes the nervous system (adapted from [14]).

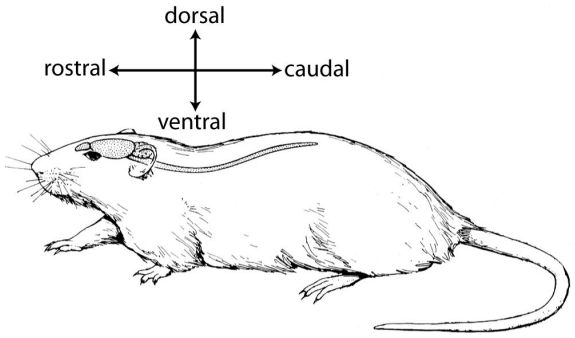

Figure 2.2: The anatomical coordinate system illustrated on the rat CNS (from [106]).

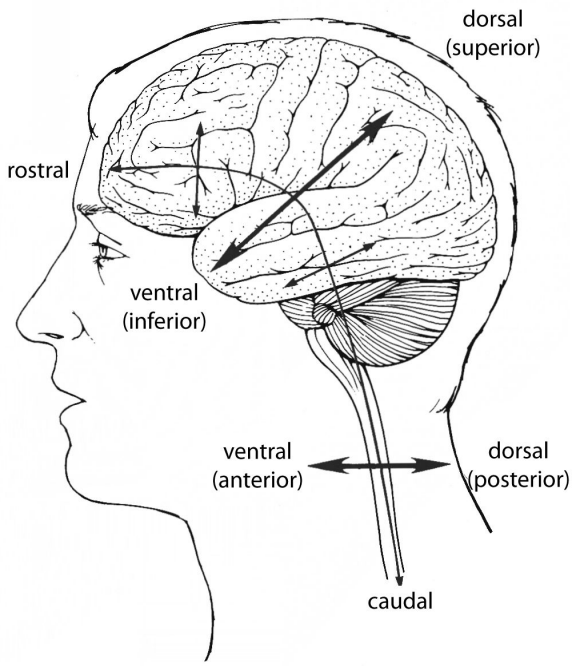

Figure 2.3: The anatomical coordinate system of the human CNS. The bend in the human CNS midline leads to terminological variations (from [106]).

direction and the direction pointing down to the rat's abdomen is called the ventral direction. Perpendicular to the rostrocaudal-dorsalventral plane is the left-right axis; either direction on this axis towards is called the lateral direction. Conversely, the direction in the left-right axis towards the the midline, the hypothetical line going through the middle of the central nervous system, is called the medial direction. Therefore, either half of the left-right axis is also referred to as the mediolateral axis.

Figure 2.3 illustrates this coordinate system in the human CNS. Note that the midline (or the longitudinal axis) in humans is not straight but is bent at the midbrain. This bend is the source of variations in terminology describing the axes in the human brain. Between the rostral end and the bend (flexure), the dorsal and ventral directions are called as superior and inferior, respectively.

2.1.2 Anatomical Planes

The customary way of viewing internal structure of the CNS is to use slices (sections) along the major anatomical reference planes. The three main anatomical planes are horizontal (axial), coronal, and sagittal planes. The sagittal plane is parallel to the rostrocaudal-dorsoventral plane, the horizontal plane is parallel to the left-right axis, and the coronal section is parallel to the plane determined by the left-right and superior-inferior axes (see Figure 2.4).

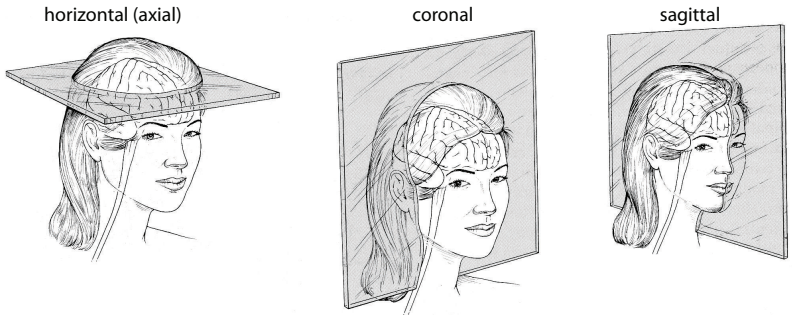

Figure 2.4: Anatomical reference planes illustrated (adapted from [106]).

2.1.3 Gross Divisions of the CNS

The CNS can be divided into five major structures that are common to all mammals: the cerebrum, the diencephalon, the cerebellum, the brain stem, and the spinal cord. Figure 2.5 illustrates these structures and their further subdivisions. We briefly discuss these major divisions below.

- The cerebrum, the largest structure of the brain, is made of two almost symmetrical hemispheres, called the cerebral hemispheres, separated by an opening called the sagittal fissure. The cerebrum consists of the basal ganglia and the cortex, the layer covering the cerebrum. Both of these structures handle higher-level perceptual, cognitive, and motor functions and are made of neural cell bodies. The cortex (cerebral cortex) has four major divisions named after corresponding cranial regions: the frontal lobe, the parietal lobe, the temporal lobe, and the occipital lobe. A closer look at the surface of cortex (cortical surface) reveals a folded structure: the grooves on the cortical surface are called gyri (singular gyrus) and the ridges are called sulci (singular sulcus) or fissures (the term fissure is generally used to refer to particularly deep sulci). The basal ganglia are essentially a network of eight deep brain nuclei (collection of neuron cell bodies in the CNS, singular nucleus) whose function vary from reward-based learning and action selection/gating to motor preparation. The functionality of this structure, however, is not completely understood.

- The diencephalon is the region containing two major relay structures of the forebrain, the thalamus and hypothalamus, squeezed between the frontal lobe and the midbrain.

- The cerebellum lies behind the cerebrum and is the center of movement control. The cerebrum and cerebellum are strongly connected through several major neural pathways.

- The brain stem is the remaining posterior part of the brain after the cerebrum, diencephalon, and cerebellum are left out. It is continuous with the spinal cord (Figure 2.5). The functionality of the brain stem is twofold: it acts as a relay center between the cerebrum/cerebellum and the spinal cord and is also where vital functions such as breathing, consciousness, and body temperature are regulated.

- The spinal cord lies within the bony vertebral column, is continuous with the brain stem, and relays information from the skin, joints, and muscles to the brain and vice versa. With the brain, it makes up

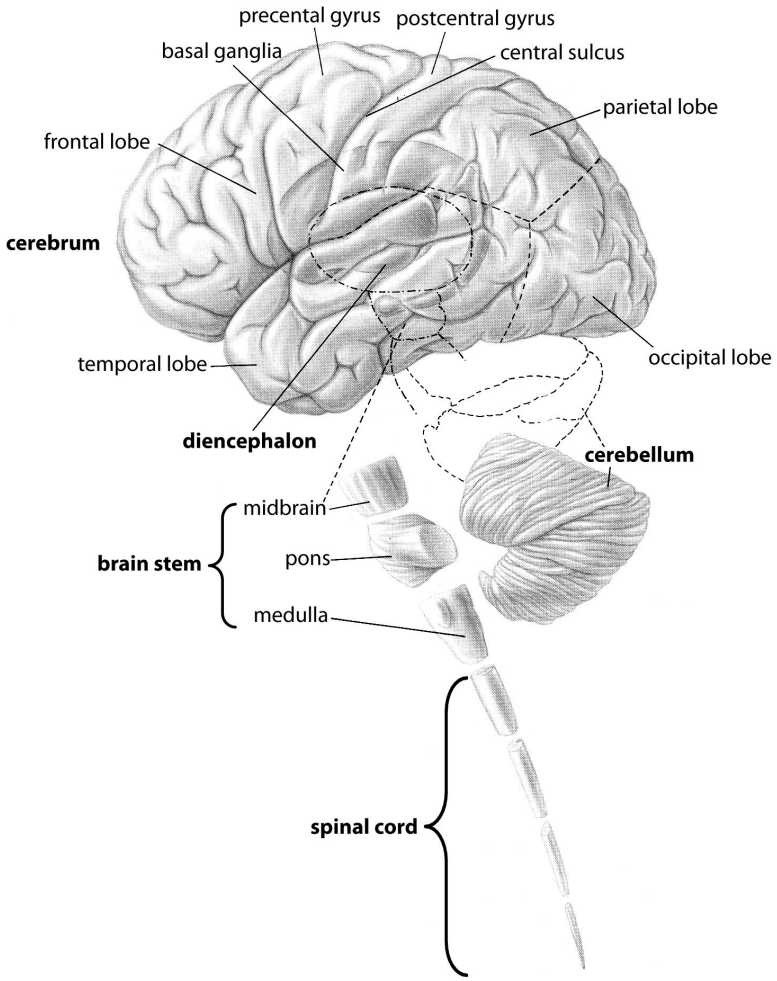

Figure 2.5: Gross CNS divisions (from [94]).

the CNS.

2.2 Cellular Structure of the Central Nervous System

The nerve cell, or neuron, is the main cellular unit of the CNS. It is supported structurally and metabolically by the other major cellular constituent of the nervous system, the neuroglial cell (or glia). There are approximately 80 billion neurons and an equal number of glial cells in the human brain, although their ratio varies across the brain [81].

While there are many different types of nerve cells, all share the same structural units that are unique to nerve cells: dendrites, cell body, axon, and axon terminals (Figure 2.6).

Dendrites receive information from other neurons. The cell body contains the nucleus and cellular organelles that are vital for neuron functioning. It performs integrative functions on information received from

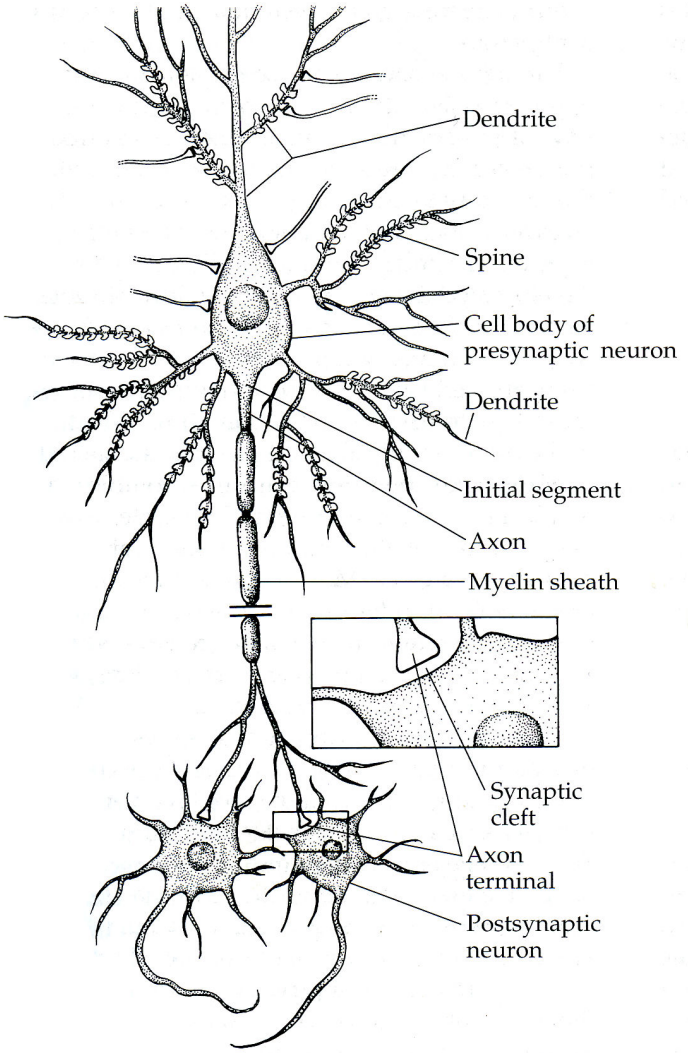

Figure 2.6: Neuron (from [106]).

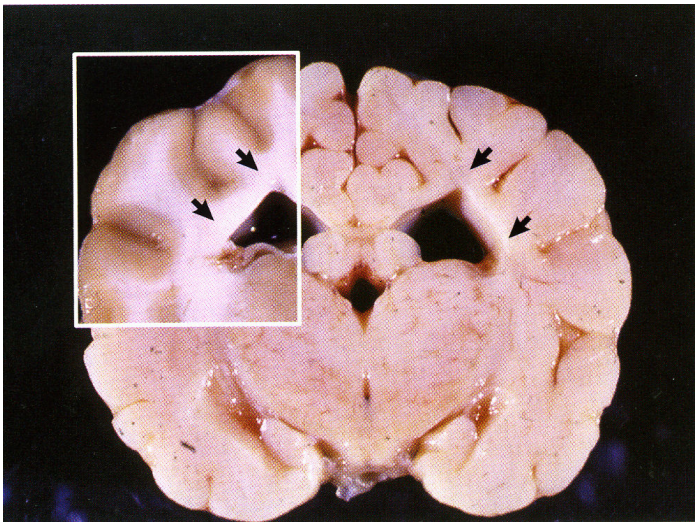

Figure 2.7: Freshly dissected coronal slice of a brain. Dark areas correspond to the gray matter while pinkish-white areas, indicated with arrows, correspond to the white matter (from [62]).

other neurons. The axon carries information encoded in the form of action potentials to the axon terminal. Many axons are covered with a myelin sheath that functions as an insulator (Figure 2.6). Dendrites and axons as extensions of the cell body (soma) are collectively called neurites. Neurons can be classified into different types on the basis of different morphological criteria. With respect to the number of neurites, for example, they can be classified as unipolar, bipolar, and multipolar. As can be guessed, unipolar neurons have a single neurite, bipolar neurons have two, and multipolar neurons have more than two. Most neurons are multipolar neurons.

With respect to dendrite configurations, there are roughly two classes of neurons: Stellate (having star-shaped dendrites) and pyramidal (having pyramid-shaped dendrites) neurons. Note that the dendrite configuration evolved to facilitate the functional behavior of these neurons. With respect to connection type, there are three classes of neurons: Primary sensory neurons, motor neurons, and interneurons; most neurons are of interneuron type. With respect to axon length, on the other hand, we have projection neurons (Golgi type I) and local circuit neurons (Golgi type II). Projection neurons synapse beyond their vicinity while local circuit neurons project to nearby, short-distance regions.

The difference between the appearance of neural cell bodies and that of axons leads to the well-known tissue division of the CNS into white matter (WM) and gray matter (GM). Specifically, parts of the CNS such as the cerebral cortex and the diencephalon, which contain cell bodies and appear gray in a freshly dissected brain is called gray matter. The parts that contain myelinated axons and appear white-pinkish in a freshly dissected brain, are called white matter (Figure 2.7).

White-matter fiber tracts (coherent collection of axons, bundles of axons) can be classified into three major categories based on the regions they connect:

- Association fibers: Fibers that connect regions within the same cerebral hemisphere. Among examples

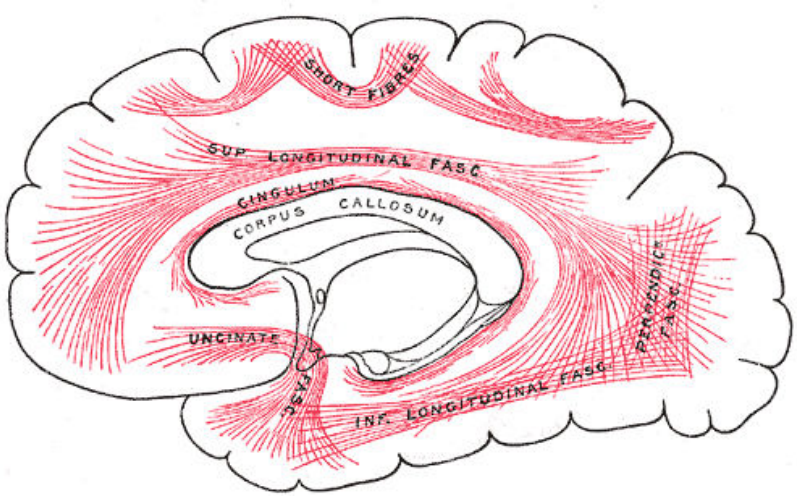

Figure 2.8: Association tracts.

are the inferior longitudinal fasciculus and the superior longitudinal fasciculus. Tracts that connect the limbic system structures (e.g., the hippocampus, amygdala, septum, etc.) with each other or other structures in the cerebrum constitute an important subset of the association fibers in the brain. Examples of the limbic system association fibers include the cingulum bundle, the fornix, and the stria terminalis.

- Projection fibers: fibers that connect the cortex with the lower parts (noncortical areas) of the brain, including the thalamus, brain stem, and the spinal cord. Examples include descending and ascending pathways (e.g., the corticospinal tract, the corticobulbar tract, the medial lemniscus), and the fibers of the thalamocortical loop (the reciprocal connection between the thalamus and the cortex).

- Commisural fibers: fibers that connect the opposite hemispheres of the brain. The corpus callosum and the anterior commisure are two primary examples of the commisural fibers in the brain.

2.3 Major Neural Pathways

Since we refer throughout this thesis to several major brain pathways, we briefly discuss them here. Note that the terms neural tract, neural pathway, fiber pathway, and fiber tract all refer loosely to collection of axons that follow a coherent trajectory and have a common site of origin and destination.

- Internal capsule: A fan-shaped planar white-matter structure formed by the convergence of various projection fibers including the corticopinal tract, the corticobulbar tract, the medial lemniscus, and the thalamocortical fibers. The dorsal part of the internal capsule is called the corona radiata.

- Corticospinal tract: A projection fiber that descends from the cerebral cortex and terminates on the motor neurons in the spinal cord. It is the most important pathway for controlling voluntary movement.

- Corpus callosum: The largest of the white-matter commisural fibers that connect left and right hemispheres of the brain.
- Cerebellar peduncles: Three sets of paired U-shaped bundles of the hindbrain (superior, middle, and inferior), connecting the cerebellum to the brain stem (midbrain, pons, and medulla oblongata).
- Uncinate fasciculus: A limbic system association pathway interconnecting the limbic system nuclei such as the hippocampus and amygdala in the anterior portions of the temporal lobe with the ones in the frontal lobe such as the orbifrontal cortex.
- Fornix: Limbic system association pathway projecting from the hippocampus into the hypothalamus. It appears C-shaped in a sagittal view and has important function in memory. While the exact nature of the effects is not clear, recent studies have consistently found that any damage to the fornix causes memory loss.
- Cingulum bundle: Another limbic system association fiber pathway that encircles the corpus callosum, interconnecting the cingulate gyrus and the entorhinal cortex of the medial temporal lobe in the brain and thus allowing communication between different components of the limbic system.
- Superior longitudinal fasciculus: Association fiber pathway that interconnects the association areas of the parietal and temporal lobes with those of the the frontal lobe. While its functional structure is still subject of ongoing research, there is no dispute on its centrality to the higher order and associative functioning of the brain.

2.4 Imaging the CNS

Different techniques aim to expose different structural or functional characteristics of the CNS at different scales, cellular or macroscopic. Not surprisingly, gross dissection is the oldest “imaging” technique, going back to the influential Greek physician Galen (130-200 AD). It was Galen who coined the term, among many others, “autopsy.” While dissection is limited in scale and requires fine manual skills from dissectors, gross dissection paved the way for exploring and tracing the internal structures of the CNS up until 18th and 19th centuries. And it is still a prevalent method used in conjunction with newer techniques.

Microscopy is by far the most common method for studying the brain at the cellular level. In the microscopic study of neural cells, staining is often used to highlight cellular structures of interest on living tissues (in-vivo staining) as well as postmortem tissues (in-vitro staining). For example, myelin stains (e.g., Weil’s stain) use dyes that bind to the myelin sheath wrapping the axons, enabling visualization of myelinated axonal connections. On the other hand, cell stains (e.g., Nissl stain) use dyes that bind to the nucleus and the surrounding components of the neuronal cell body. Some staining techniques, such as silver staining (Golgi staining), color both the cell body and neurites (Figure 2.9). Staining is an active research area and ever-improving techniques make possible selective labeling of individual neurons and their specific parts. Genetic labels, for example, have been developed to selectively visualize specific neurons and neuronal structures. One well publicized method using genetically encoded tracers is Brainbow (Figure 2.10). As usual, structure

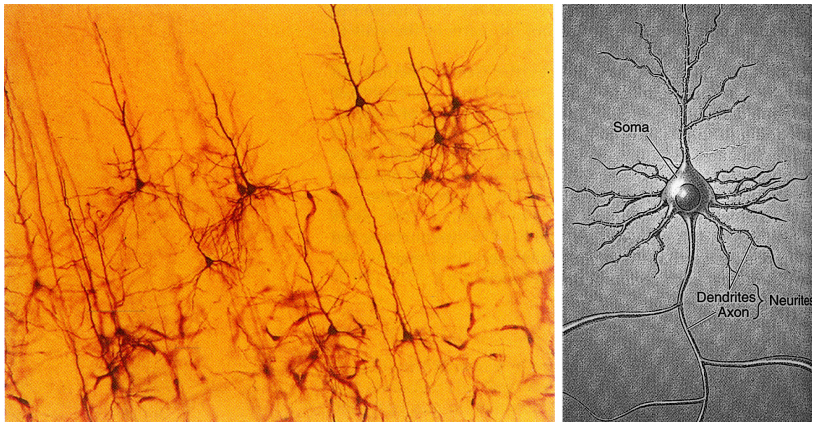

Figure 2.9: Golgi staining, named after the Italian physician Camillo Golgi (1843-1926), is a general name for any staining method that changes the appearance of the neurons in a brain tissue sample by immersing the tissue in potassium dichromate and silver nitrate solutions. The main characteristic of this method is that it stains only a random subset of the neurons in the tissue, which is both a strength and a limitation of the technique (adapted from [14]).

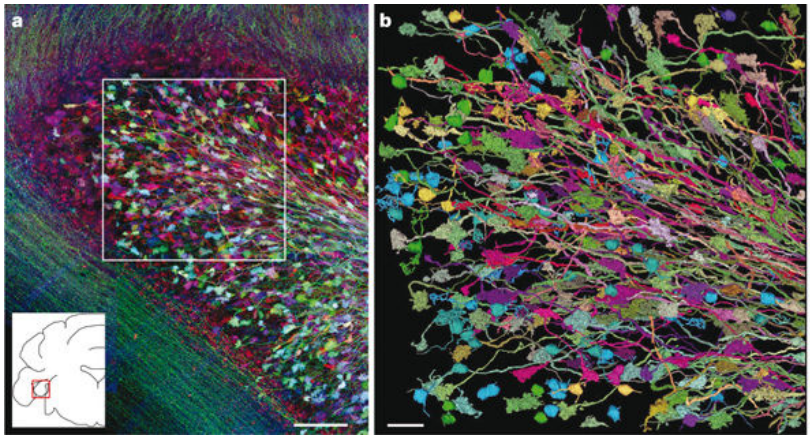

Figure 2.10: Individual neurons are colored with different colors using a genetic labeling technique called Brainbow, based on expression of multiple fluorescent proteins from a single transgene (image from [105]).

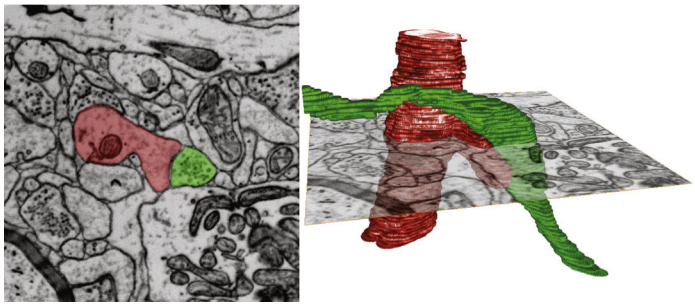

Figure 2.11: Nanoscale electron microscopy captures the geometry of individual axons (from [84]).

is only one part, often the beginning, of the story; there are also various anterograde (terminal-to-cell-body) and retrograde (cell-body-to-terminal) tracing methods that help reveal the directionality of the connectivity using axonal transport mechanisms.

Recent advances in microscopy, such as electron microscopy (EM), have enabled the individual axons to be imaged, overcoming the resolution limitation of light microscopy. EM has become one of the driving forces behind the emerging field of connectomics. In this context, combining EM with other neuroimaging techniques, structural and functional alike, is a promising approach for research into neural connectivity [21, 26]. Revealing long distance axonal connections with EM remains, however, to be a challenge.

There are several radiological imaging techniques for structural and functional imaging of the brain. These techniques are particularly important in visualizing the brain noninvasively. Compared to microscopy, radiological images in general give a coarse-scale view of the brain. Here, we briefly discuss two common structural imaging techniques with their functional counterparts: computerized tomography (CT) and magnetic resonance imaging (MRI). CT makes it possible to reconstruct an image of the brain based on the degree to which different tissues absorb transmitted X-rays. On the other hand, MRI uses information about how hydrogen atoms in tissues respond to perturbations of a strong magnetic field. While CT is often used for visualizing bony and calcified structures, MRI is good for distinguishing soft tissues. MRI is typically used to distinguish fatty (such as myelin containing) tissues from water-rich ones (Figure 2.12).

Both CT and MRI have "functional" variants that are used to detect functional changes in the brain; positron emission tomography (PET) and functional magnetic resonance imaging (fMRI). While the underlying working mechanisms of these two techniques are different, they both detect changes in the local blood flow and metabolism. This is useful because, when active, neurons demand more glucose and oxygen. The brain vasculature (the network of ventricles and arteries that provide the brain's blood supply) responds to this neural activity by directing more blood to the active regions. Therefore, by detecting blood activity, PET and fMRI provide a picture of the functionally active brain regions at a given time. Such indirect probing and measurement is a recurring theme in radiological imaging.

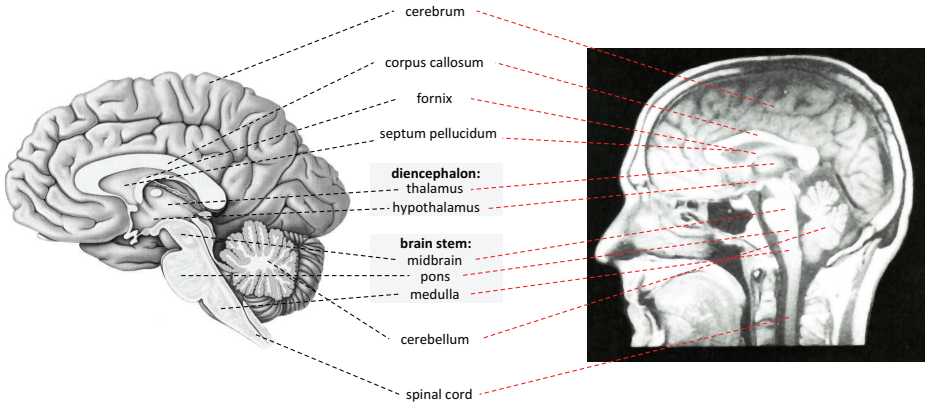

Figure 2.12: The midsagittal crossection of the human brain with major structures identified: (left) drawing, (right) T_1-weighted MRI image, in which fat appears brighter and water darker (adapted from [106]).

2.5 Diffusion MRI

The source of structural brain connectivity in this dissertation is diffusion-weighted magnetic resonance imaging Compared to conventional MRI, diffusion-weighted imaging provides much richer anatomical information about the white-matter structure in the brain (the white matter looks homogeneous on conventional MRIs). It is also unique among imaging techniques in being noninvasive and inherently three-dimensional. In biological structures, the shape of the diffusion is determined by the macromolecules, cell membranes, geometry of the surrounding microstructure, etc. Therefore, by measuring diffusion, diffusion-weighted imaging opens up possibility of revealing the underlying microstructure. Note that the terms diffusion MRI, diffusion-weighted MRI (DW-MRI), diffusion-weighted imaging (DWI) and diffusion imaging all refer to diffusion-weighted magnetic resonance imaging.

2.5.1 What Is Diffusion?

Diffusion is a physical process that arises from mixing or migration of molecules or particles through Brownian motion due to their internal thermal energy. This process does not require bulk motion of the particles. The phenomenon was first observed by the Scottish botanist Robert Brown (1773-1858), whence the name Brownian motion. While characterizing an individual particle's trajectory may not be particularly useful, as it is random, doing so for the average displacement of an ensemble of particles is. Einstein (1879 - 1955) in his 1905 article showed that, given a sufficiently large number of particles within a constant temperature environment in which they can diffuse freely, the average squared displacement of the particles is proportional to the observation time t:

$$r^2 = 6Dt$$

where D is the constant diffusion coefficient. The probability distribution of the mean squared ensemble displacements is a Gaussian distribution (by the central limit theorem), where the probability of zero displacement is the highest and the probability of the displacement r in every direction is the same (i.e., the

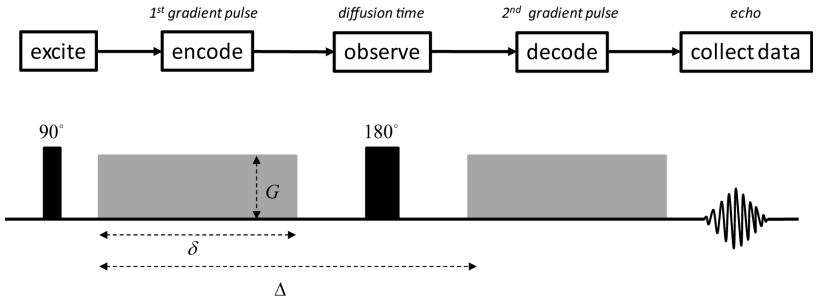

Figure 2.13: Pulsed gradient spin echo (PGSE) experiment.

distribution is isotropic). If we observe water molecules diffusing for 50 milliseconds at body temperature, for which $D \approx 3 \times 10^{-3}\text{mm}^2\text{s}^{-1}$, their average displacement will be 10 μm in all directions in three-space. This average motion is comparable to the size of many tissue structures and cells in the brain. Note that the diffusion coefficient is constant only if the temperature is fixed. Otherwise, it is a function of the temperature. For example, the diffusion coefficient for water is approximately 1×10^{-3} mm^{2}s^{-1} at 1 °C while it is 8.6×10^{-3} mm^{2}s^{-1} at 100 °C.

2.5.2 How to Encode the MR Signal for Diffusion

As noted earlier, MR imaging works by measuring the response of water molecules to perturbations of a strong magnetic field. Specifically, it measures the signal from hydrogen nuclei, which have magnetic dipoles called spins with corresponding north and south poles. Spins can be viewed as microscopic needles that can be aligned with a strong magnetic field. The Larmor equation, which is the basic equation in MR, states that the precessional frequency of spins in a magnetic field is proportional to the strength of the magnetic field. For example, in a typical MRI magnet of 1.5 tesla (T), the rate of precession is about 64 MHz. Because the strength of the magnetic field is kept as homogeneous as possible, this precession rate is also very homogeneous across the magnet. This homogeneity can be disturbed linearly by using a so-called pulsed field gradient. The strength (slope) of the gradient, its direction, and the time period can be controlled.

While there are many diffusion-sensitizing MR acquisition protocols, the most common is the pulsed gradient spin-echo (PGSE) experiment (Figure 2.13), proposed by Stejskal and Tanner [146]. It is rather simple. First, all the spins are aligned with magnetic field of the MR coil with a 90-degree RF excitation pulse. Then, a magnetic field gradient is applied for duration δ. This gradient is reversed, as we discussed, by applying another RF pulse, this time at 180-degrees, followed by a magnetic gradient identical to the first. In this context, MR sequences are made sensitive to diffusion by the addition of magnetic field gradients, i.e., a magnetic field that varies linearly in the volume of interest. The magnetic field gradient imposes phase differences in spin precessions that are proportional to the magnetic field gradient. Therefore, there will be a gradient of phase increases spatially varying along the magnetic field. This is the foundation of spatial encoding in MR. This gradient of phase acquisitions is typically reversed by applying a 180- degree refocusing pulse first, then a gradient that is exactly the same as the first magnetic field gradient. If spins

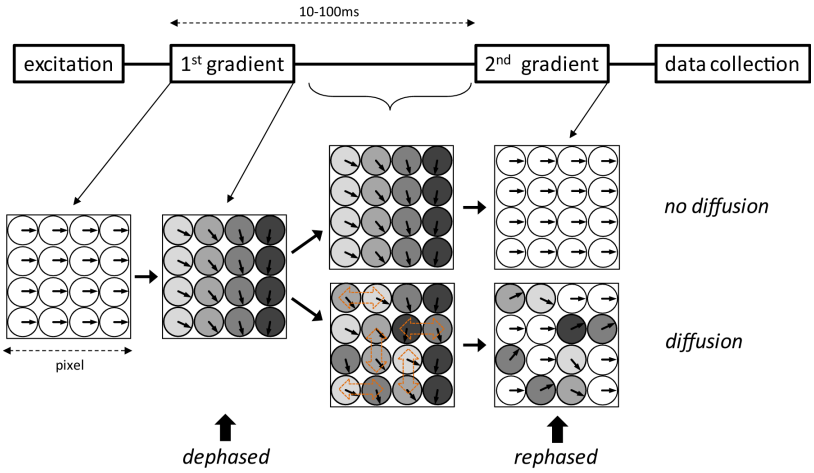

Figure 2.14: Spin behavior in diffusion-weighted MR experiment (adapted from [153]).

were stationary (i.e., not diffusing) then we would expect spins to be in the same phase as when the first RF pulse was applied, so that the strength of the field would be the same. However, we know that the water molecules to which the spins are attached are in constant random motion. So, there will be a loss in field strength that is proportional to the amount of the diffusion. Therefore, due to diffusion, we get a distribution of displacements that leads to a loss in the field coherence and, hence, a decrease in the amplitude of the signal. Figure 2.14 illustrates spin behavior in a PGSE experiment.

The amount of phase dispersion depends on the strength and duration of the magnetic field gradient and the time difference ("mixing" or "diffusion" time) between the two gradients:

$$b = \gamma^2 G^2 \delta^2 \left(\Delta - \frac{\delta}{3} \right)$$

where γ is the gyromagnetic ratio, G is the amplitude of the magnetic field gradient pulses, and δ and Δ are the duration and temporal separation of the gradient pulses.

2.5.3 Diffusion Models

As discussed above, the current standard method for measuring diffusion in tissues with MRI is based on the Stejskal-Tanner PGSE sequence [146]. This method samples from the Fourier transform of molecular displacement that is characterized by a probability density function $p(\mathbf{r}|\tau)$, where $\mathbf{r}$ is the displacement vector and τ is the diffusion time. In this context, the relation between the MRI signal $S(\mathbf{q},\tau)$ measured in the direction of a diffusion gradient $\mathbf{G}$ and $p(\mathbf{r}|\tau)$ in a Stejskal-Tanner spin-echo sequence experiment is given by

$$S(\mathbf{q},\tau) = S_0 \int_{R^3} p(\mathbf{r}|\tau) e^{-2\pi i \mathbf{q}^T \mathbf{r}} d\mathbf{r} \tag{2.1}$$

$$= S_0 \int_{R^3} p(\mathbf{r}|\tau) e^{-i\gamma\delta|\mathbf{G}|\mathbf{g}^T \mathbf{r}} d\mathbf{r} \tag{2.2}$$

where $\mathbf{q} = \gamma\delta\mathbf{G}/2\pi$, S_0 is a reference signal obtained without a diffusion gradient, γ and δ are as above, and $\mathbf{g}$ is the unit vector in the direction of the diffusion gradient $\mathbf{G}$ (i.e., $\mathbf{g} = \mathbf{G}/|\mathbf{G}|$).

What we want to recover from the equation above is the probability density function $p(\mathbf{r}|\tau)$ of molecular displacement. In this sense, modeling diffusion locally can be seen as a probability density estimation problem. There are essentially two basic approaches, parametric and nonparametric.

Scalar Diffusion Model

If we assume the probability density function $p(\mathbf{g}|\tau)$ to be an isotropic Gaussian and take the temporal difference between the gradients into account, then the original equation reduces to

$$S_1 = S_0 e^{-bD}$$

This equation is known as the Stajskel-Tanner equation. Note that the assumption of an isotropic Gaussian distribution implicitly assumes that the signal is invariant to the gradient direction, reducing p to an essentially one-dimensional Gaussian. Note also that, given the mean fixed at the center of the voxel, there is only one free parameter, the variance, that determines the distribution in this case. D is called the apparent diffusion coefficient (ADC). It is important to understand that what is measured here is the amount of phase dispersion of the spins, which is effected by the cytoarchitecture and cytochemistry of the tissue. Hence, the term "apparent".

While using a single scalar value is sufficient in characterizing diffusion in non-coherent tissue structures such as the gray matter, it is not valid for aligned fibrous structures such as the white matter, muscles, and heart. The primary cause of the anisotropic diffusion in the CNS is the axonal membranes, although myelin plays a secondary role in its modulation. Note that, however, since both myelin and axon packing density correlate with reduced perpendicular diffusion (hence increased anisotropy), a clear link has not yet been established [15].

Therefore, removing the isotropic Gaussianity assumption leads to the tensor model of diffusion.

The Diffusion Tensor Model

Generalization of the Stejskal-Tanner equation using a single second-order tensor to quantify diffusion is the most widely used model in diffusion MRI [12]. In this case, equation 2.1 becomes

$$S(\mathbf{g}, \tau) = S_0 e^{-b\mathbf{g}^T\mathbf{D}\mathbf{g}} \tag{2.3}$$

where $b = \tau|\mathbf{q}|^2 = \tau(\gamma\delta\,|\mathbf{G}|\,/2\pi)^2$ and $\mathbf{D}$ is a three-dimensional, second-order symmetric tensor. In this case, $p(\mathbf{r}|\tau)$ has the following form:

$$p(\mathbf{r}|\tau) = \frac{e^{-\frac{\mathbf{r}^T\mathbf{D}^{-1}\mathbf{r}}{4\tau}}}{\sqrt{(4\pi\tau)^3|\mathbf{D}|}} \tag{2.4}$$

Since the mean of this three-dimensional normal distribution is assumed to be $(0,0,0)$, the covariance matrix (or the second-order symmetric tensor) $\mathbf{D}$ determines the distribution. A second-order diffusion tensor in

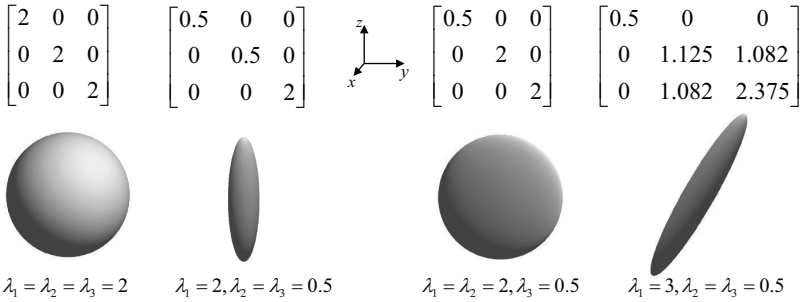

Figure 2.15: Tensor examples.

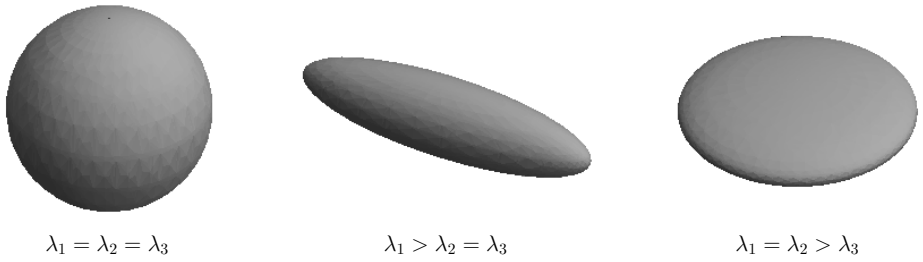

Figure 2.16: Three configurations of tensor ellipsoids. $\lambda_1, \lambda_2, \lambda_3$ correspond to lengths of the longest, median and shortest semiaxes of the ellipsoid. They also represent the eigenvalues of the corresponding diffusion tensor. Similarly, the directions of the semiaxes of the ellipsoid represent the eigenvectors of the corresponding diffusion tensor.

three-space is a 3-by-3 symmetric positive definite matrix (see Figure 2.15 for examples):

$$D = \begin{pmatrix} D_{xx} & D_{xy} & D_{xz} \\ D_{yx} & D_{yy} & D_{yz} \\ D_{zx} & D_{zy} & D_{zz} \end{pmatrix}$$

Diffusion MRI based on the tensor model of diffusion is called diffusion tensor imaging (DTI) or diffusion tensor magnetic resonance imaging (DT-MRI). Note that the three eigenvalues, λ_1,λ_2, and λ_3, of the diffusion tensor $\mathbf{D}$ correspond to the diffusivities along the principal axes $\mathbf{e_1}$, $\mathbf{e_2}$, and $\mathbf{e_3}$. D also uniquely determines an ellipsoid (i.e., $x^t Dx = 1$), so we can visualize diffusion tensors with the corresponding ellipsoids. Consequently, the major axes of the ellipsoid and corresponding radii also represent the direction of the major diffusion axes (eigenvectors of D) and diffusion rates (eigenvalues of D; see Figure 2.16).

The most important limitation of the tensor model is that it cannot capture mixtures of tissue orientations. To address this limitation, multi-tensor models have been proposed that turn the single Gaussian density estimation problem into a non-linear Gaussian mixture density estimation problem. (e.g., [3, 4]). An n-tensor

model of diffusion is given by

$$S_i = S_0 \sum_{j=1}^{n} \alpha_j e^{-b\mathbf{g_i^T D_j g_i}} \tag{2.5}$$

where α is the mixing coefficient (or volume fraction) with $\sum_j \alpha_j = 1$.

There are other model-based approaches that aim to address the shortcomings of the tensor model. Separating the diffusion in extracellular and intracellular spaces to directly relate the diffusion processes in the underlying tissue compartments to the MRI signal is a common theme in most of these works (e.g., [125, 17, 9, 163]). Still, many of these approaches assume some form of Gaussian diffusion for different compartments modeled while requiring more complex MRI acquisition protocols.

Nonparametric Models

It is possible to estimate $p(\mathbf{r}|\tau)$ without assuming a particular shape for the distribution. All existing nonparametric approaches try to reconstruct a probability density function defined on the sphere. One approach is to use the Fourier relationship between the diffusion function and the signal echo attenuation in q-space. Diffusion spectrum imaging (DSI) samples diffusion on a volumetric grid [157], while q-ball imaging samples the density function on a sphere [150].

2.5.4 Estimating the Diffusion Tensor

In DTI, a single second-order diffusion tensor D is fit at each voxel to diffusion MRI signals measured at different gradient field directions. $\mathbf{D}$ above is an unknown symmetric tensor with six independent elements. Therefore, in addition to S_0, we need at least six measurements[1] taken at different gradient directions to estimate $\mathbf{D}$ at each voxel:

$$S_i(\mathbf{g}_i, \tau) = S_0 e^{-b\mathbf{g}_i^t \mathbf{D} \mathbf{g}_i} \tag{2.6}$$

$$\ln S_i(\mathbf{g}_i, \tau) = \ln S_0 - b\mathbf{g}_i^t \mathbf{D} \mathbf{g}_i \tag{2.7}$$

For m measurements, equation 2.7 constitutes a system of linear equations. First, observe that given $\mathbf{g}_i = (x_i y_i z_i)$,

$$\begin{aligned}
\mathbf{g}_i^t \mathbf{D} \mathbf{g}_i &= \begin{pmatrix} x_i & y_i & z_i \end{pmatrix} \begin{pmatrix} D_{xx} & D_{xy} & D_{xz} \\ D_{yx} & D_{yy} & D_{yz} \\ D_{zx} & D_{zy} & D_{zz} \end{pmatrix} \begin{pmatrix} x_i \\ y_i \\ z_i \end{pmatrix} \\
&= \begin{pmatrix} x_i^2 & y_i^2 & z_i^2 & 2x_i y_i & 2x_i z_i & 2y_i z_i \end{pmatrix} \begin{pmatrix} D_{xx} \\ D_{yy} \\ D_{zz} \\ D_{xy} \\ D_{xz} \\ D_{yz} \end{pmatrix}
\end{aligned}$$

[1] In practice, more than six measurements are used for better results.

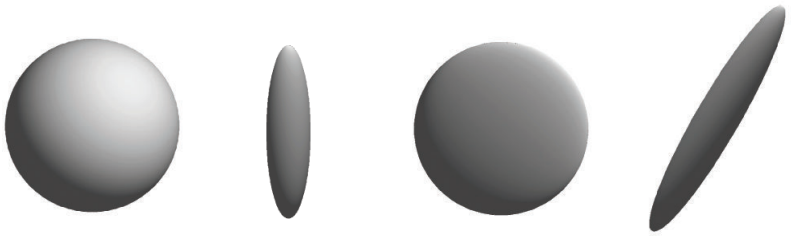

Figure 2.17: Tensors with various FA and MD values.

Now we incorporate S_0 and express m equations of the form 2.8 as a matrix multiplication:

$$\begin{pmatrix} x_1^2 & y_1^2 & z_1^2 & 2x_1y_1 & 2x_1z_1 & 2y_1z_1 & -1/b \\ x_2^2 & y_2^2 & z_2^2 & 2x_2y_2 & 2x_2z_2 & 2y_2z_2 & -1/b \\ \vdots & & & & & & \\ x_m^2 & y_m^2 & z_m^2 & 2x_my_m & 2x_mz_m & 2y_mz_m & -1/b \end{pmatrix} \begin{pmatrix} D_{xx} \\ D_{yy} \\ D_{zz} \\ D_{xy} \\ D_{xz} \\ D_{yz} \\ \ln S_0 \end{pmatrix} = -\frac{1}{b} \ln \begin{pmatrix} S_1 \\ S_2 \\ S_3 \\ \vdots \\ S_m \end{pmatrix}$$

This linear system of equations can be easily solved using, for example, least squares minimization.

2.5.5 Parameters Derived from the Diffusion Tensor

The most common measures are the mean diffusivity (MD) and fractional anisotropy (FA). Both can be seen as simple statistics over the eigenvalues of D. While MD is the mean of the eigenvalues, FA corresponds to the standard deviation of the normalized eigenvalues and its value ranges from 0 to 1 (see Figure 2.17). FA is an anisotropy measure, quantifying the extent with which the diffusion tensor deviates from being isotropic [18]:

$$\mathrm{MD} = \frac{\lambda_1 + \lambda_2 + \lambda_3}{3}$$

$$\mathrm{FA} = \sqrt{\frac{3}{2} \frac{(\lambda_1 - \mathrm{MD})^2 + (\lambda_2 - \mathrm{MD})^2 + (\lambda_3 - \mathrm{MD})^2}{\lambda_1^2 + \lambda_2^2 + \lambda_3^2}}$$

Further details on the theoretical and practical underpinnings of DTI data acquisition, processing, and applications can be found in any good review (e.g., [13, 116]).

2.5.6 Tractography

Models of diffusion that capture anisotropic diffusion as measured MR enable neural pathways in the *in vivo* brain to be estimated as a collection of space curves called a tractogram. In the diffusion tensor model,

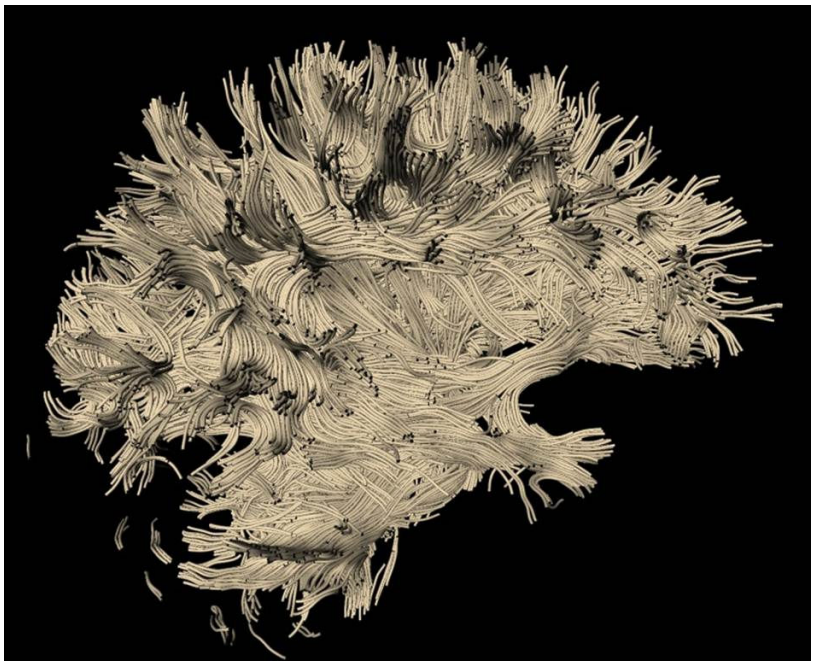

Figure 2.18: Whole brain tractogram.

tractograms are essentially solutions to the first-order differential equation

$$\frac{dC}{ds} = \vec{v_1}(C(s))$$

where s parametrizes the curve and v_1 is the principal eigenvector at the point $C(s) = (x(s), y(s), z(s))$. We compute the integral curve $C(s)$ passing through a given seed point $C(0)$ (initial conditions) by integrating the above equation for $s > 0$ and $s < 0$ (i.e., in both directions from the seed point). The resulting integral curves are often represented as polylines. Tractography, the study of revealing neural pathways with tractograms, has important applications in both clinical and basic neuroscience research on the brain.

Tractograms provide a valuable noninvasive window into the internal structure of white matter. One interesting question is what is the scale of the anatomical information provided? The radius of axons varies across the CNS. For example, given a 1 mm^3 voxel size, a radial crossection of an individual tractogram curve packs approximately $300,000$-$400,000$ axons in the human corpus callosum [62]. It is also important to understand that diffusion MRI is a purely structural imaging modality and as such cannot distinguish afferent pathways, which carry information from CNS to PNS, from efferent pathways, which bring information from PNS to CNS.

2.5.7 Applications of Diffusion MRI

The clinical relevance of diffusivity in the brain has been established with several studies. One of the earliest and arguably the most successful clinical application of diffusion MRI has been in brain ischemia. Moseley *et al.* found that diffusion drops at a very early stage of the ischemia in the cat brain [118]. Diffusion

MRI gives some patients the opportunity to receive suitable treatment at a stage when brain tissue might still be salvageable. In general, diffusion imaging has found applications in all neurodegenerative diseases affecting white-matter integrity such as multiple sclerosis [83, 158], CADASIL [35, 113], ALS [64, 76, 137], Alzheimer's disease [78, 57, 135], and HIV [66, 129, 133]. In addition to clinical research, diffusion MRI is increasingly used in neuroscience research. Examples include phenotype characterization [165, 155], brain development [119, 164], and white-matter segmentation and visualization [32, 16].

2.6 Conclusions

We have given an overview of the structural and cellular anatomy of the CNS and related concepts. We have also discussed examples of neuroimaging techniques while focusing on diffusion-weighted MRI, which is the radiological imaging modality used to collect the brain connectivity datasets here. Our aim in this overview was to make reading through the following chapters easier. Clearly, this short chapter cannot do justice to such a broad and deep set of topics, most of which are active research areas and hence are evolving rapidly. Bear *et al.* include chapters that provide a friendly introduction to the subject [14]. We refer interested readers to [106] and [94] for an in-depth study of the cellular and gross neuroanatomy. Similarly, on the diffusion-imaging side, two recent edited books give a good up-to-date review of the field and relevant research literature [90, 93].

Part I

MODELING

Chapter 3

Generalizing Diffusion Tensor Model Using Probabilistic Inference in Markov Random Fields

3.1 Introduction

It is well known that the standard single second-order tensor model of diffusion falls short in quantifying fiber geometries such as crossing, kissing, and branching. This limitation has led to the development of different acquisition techniques along with more complex models of diffusion. One strategy for recovering the underlying complex fiber architecture is to estimate nonparametrically the displacement probability distribution function, which characterizes the diffusion locally, using the Fourier relationship between diffusion and the signal echo attenuation in q-space [157, 150]. These q-space methods, however, require a large number of scans and hence long acquisition times, making them impractical in a clinical setting. Therefore, various regularization techniques have been proposed to sample the q-space efficiently and thereby reduce the number of scans required (e.g., [54, 108, 110]).

Some of earlier work has also proposed multi-tensor as well as simple, restricted geometric models to address the issue of mixed tract geometries (e.g., [4, 126]). In this context, the work presented here can be viewed as a regularized multi-tensor model. Given signal values (i.e., DWI sequences) and a set of prior constraints, we would like to know the most likely tensor configuration that "explains" the signal. This is a typical Bayesian inference problem. Therefore, we model the probability distribution of configurations with a Markov random field (MRF) and obtain the most probable configuration using the maximum a priori (MAP) estimate [71]. While we apply the MAP-MRF framework to extend the tensor model here, as a general spatial probabilistic model selection approach, it can easily be applied to other diffusion models. Note that, all or parts of this chapter have appeared in substantially similar form at [50, 51].

3.2 Background and Related Work

The single second-order tensor model of diffusion is the most widely used model quantifying anisotropic diffusion with MRI [12]. As remarked above, one of the limitations of this model is that it cannot handle mixtures of different tissue orientations. To address the issue, multi-tensor models have been proposed that turn the single Gaussian density estimation problem into the Gaussian mixture density estimation. (e.g., [3, 4]). An n-tensor model of diffusion is given by

$$S_i = S_0 \sum_{j=1}^{n} \alpha_j e^{-b\mathbf{g_i^T D_j g_i}} \tag{3.1}$$

where α is the mixing coefficient (or volume fraction) with $\sum_j \alpha_j = 1$.

Our method can be considered a multi-tensor model in which the number of tensors, and possibly tensor parameters, is estimated using Bayesian model selection. In the demonstration here, we use a two-tensor model with $\alpha_1 = \alpha_2 = 0.5$.

It is possible to estimate $p(\mathbf{r}|\tau)$ without assuming a particular shape for the distribution. All existing non-parametric approaches try to reconstruct a probability density function defined on the sphere. One approach is to use the Fourier relationship between the diffusion function and the signal echo attenuation in q-space. Diffusion spectrum imaging (DSI) samples diffusion on a volumetric grid [157], while q-ball imaging samples the density function on a sphere [150].

3.2.1 Markov Random Fields

A Markov random field is a conditional probability distribution with a Markov property over a set of random variables described by an undirected graphical model $G(V,E)$ with a set V of vertices and a set E of edges [71]. The vertex set V of the graph corresponds to the random variables and the edge set E determines the conditional dependencies (i.e., Markov properties). One of the advantages of MRFs is that it is easy to model local dependencies (or interactions) with arbitrary energy functions while ensuring that the probability distribution remains a proper probability distribution. Note that the Hammersley-Clifford theorem indicates that all MRFs have Gibbs representations, given $P(x) > 0$ [77].

MRFs have been used in the past to regularize diffusion tensor fields [130, 131]. The main difference between the two earlier works using MRFs and the present effort is that while we estimate the most likely configurations that can include multiple tensors for some voxels (crossing, branching, etc.), they estimate only the most likely single-tensor configurations. In other words, the MRF-based regularization method in [130, 131] does not resolve mixed tissue orientations. Also, at the technical level, we exploit more recent developments in MAP estimation algorithms: In addition to the greedy algorithm iterated conditional mode (ICM) used in [130, 131], we also experiment with belief propagation (BP) and tree-reweighted message passing (TRW) algorithms. Our results suggest that, even in the simple setting used, BP and TRW are more accurate and faster than ICM.

3.3 Methods

We express the problem of finding the most likely tensor configuration for a given sequence of diffusion images as a Bayesian inference problem. If x is a random variable representing the hidden tensor configuration and y is another random variable corresponding to the diffusion signal, then Bayes' rule suggests that

$$P(x|y) \propto P(y|x)P(x)$$
$$\text{posterior} \propto \text{likelihood} \times \text{prior}$$

where we model the prior distribution $P(x)$ using an MRF. In this context, the most likely tensor configuration is an assignment of x that maximizes the posterior probability $P(x|y)$.

3.4 MAP Estimation and Equivalent Energy Formulation

The maximum a posteriori (MAP) estimate is a choice of x that maximizes $P(x|y)$; MAP estimation in MRFs has an equivalent energy formulation [147] that we have used in our experiments. Let $V = \{v_1, v_2, \ldots\}$ be a set of random variables (e.g., corresponding to the voxels in a DWI sequence) and L be a finite set of labels (e.g., number of tensors to fit). Let x_i denote a labeling of v_i; then

$$E(\mathbf{x}) = \underbrace{\sum_{v_i \in V} D(x_i)}_{\text{data cost}} + \underbrace{\sum_{(i,j) \in N} W(x_i, x_j)}_{\text{smoothness cost}} \tag{3.2}$$

where $D(x_i)$ is the cost of assigning label x_i to v_i and $W(x_i, x_j)$ is the cost of assigning label x_i and x_j to two neighboring v_i and v_j. Minimizing $E(\mathbf{x})$ means finding a labeling assignment vector $\mathbf{x}$ that minimizes the cost.

3.5 Experiment

We demonstrate the use of the MAP-MRF framework on a synthetic fiber-crossing dataset.

3.5.1 Synthetic Data and MRF Model

We generate a synthetic fiber crossing diffusion MRI sequence using 12 gradient directions with $b = 1500$ s mm^{-2} [11]. We use a two-dimensional 4-connected grid MRF.

For our experiment, we use a simple smoothness function $W(x_i, x_j)$, where x_i and x_j are the random variables corresponding to the two neighboring nodes (voxels) i and j in the graphical model (MRF). Note that $x_i, x_j \in L = \{1, 2\}$ and, hence, $\mathbf{x} \in \{1, 2\}^{32 \times 32}$.

We first define a "distance" between two tensors $\mathbf{D}_i$ and $\mathbf{D}_j$ using their eigenvectors:

$$\rho(\mathbf{D}_i, \mathbf{D}_j) = \sum_{k=1}^{3} \theta(\mathbf{e}_i^k, \mathbf{e}_j^k) \tag{3.3}$$

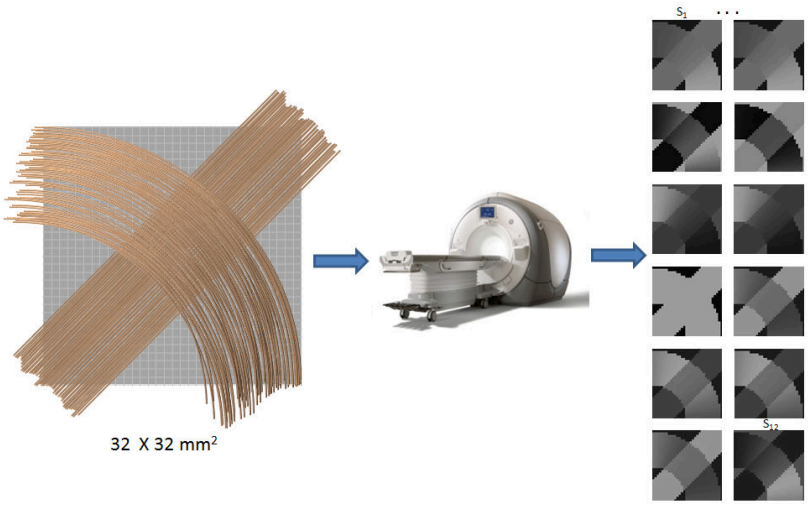

Figure 3.1: Synthetic diffusion MRI data of resolution 32×32 generated using $b = 1500$ s/mm^2 and 12 diffusion gradient directions.

where $\mathbf{e}_l^k$ is the kth eigenvector of the tensor $\mathbf{D}_l$, and $\theta(\mathbf{u},\mathbf{v}) = \arccos(|\mathbf{u}^T\mathbf{v}|)$ (the smaller angle between the two vectors). Therefore, ρ is basically the sum of the angles between corresponding eigenvectors of the two tensors.

Our smoothness function tries to capture how well the principal axes of tensors associated with one node are aligned with those of the other node:

$$W(x_i,x_j) = \begin{cases} \rho(\mathbf{D}_i^1,\mathbf{D}_j^1) & \text{if } x_i = 1 \text{ and } x_j = 1 \\ \min(\rho(\mathbf{D}_i^1,\mathbf{D}_j^1),\rho(\mathbf{D}_i^1,\mathbf{D}_j^2)) & \text{if } x_i = 1 \text{ and } x_j = 2 \\ W(x_j,x_i) & \text{if } x_i = 2 \text{ and } x_j = 1 \\ \rho(\mathbf{D}_i^l,\mathbf{D}_j^m) + \rho(\mathbf{D}_i^{\{1,2\}-l},\mathbf{D}_j^{\{1,2\}-m}) & \text{if } x_i = 2 \text{ and } x_j = 2 \end{cases}$$

where $(l,m) = \arg\min_{(l,m)} \rho(\mathbf{D}_i^l,\mathbf{D}_j^m)$.

The last case, where $x_i = 2$ and $x_j = 2$, needs some explanation. Here we have exactly two tensors $\{\mathbf{D}_i^1,\mathbf{D}_i^2\}$ and $\{\mathbf{D}_j^1,\mathbf{D}_j^2\}$ associated with x_i and x_j, respectively. First, we find the pair $(\mathbf{D}_i^l,\mathbf{D}_j^m)$ for which ρ is the minimum. Then we define W to be the sum of this minimum and the distance between the remaining pair.

The data cost term used in this case is also very simple:

$$D(x_i) = \begin{cases} 1 & \text{if } x_i = 1 \\ 2 & \text{otherwise} \end{cases}$$

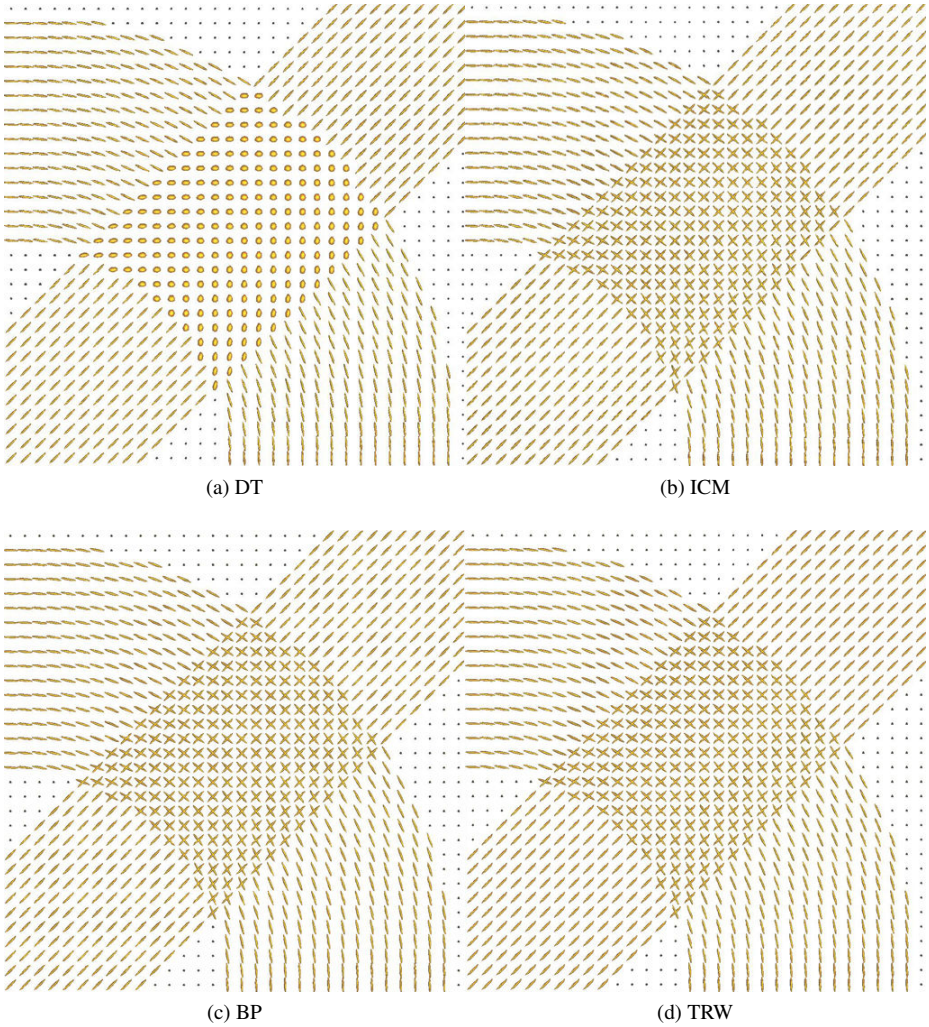

(a) DT (b) ICM

(c) BP (d) TRW

Figure 3.2: Results for a 2D synthetic fiber-crossing data set. While the single-tensor fitting (DT) cannot resolve the crossing, all the MAP estimation algorithms used (ICM, BP, and TRW) "select" the underlying orientation configuration correctly.

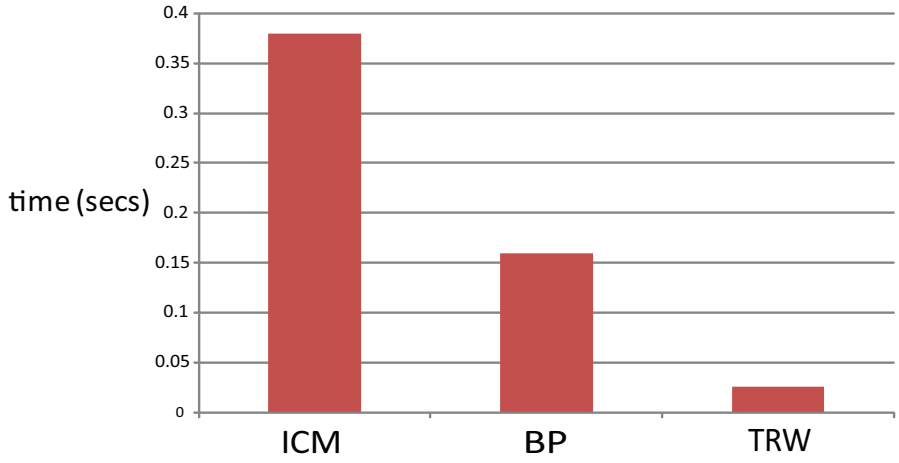

Figure 3.3: Running times for the MAP estimation algorithms.

3.6 Results

We use three different MAP estimation algorithms: iterated conditional mode (ICM), belief propagation (BP), and tree-reweighted message passing (TRW) [147]. All the three algorithms, even one as simple as ICM, provide good results, effectively selecting the configuration that represents the fiber crossing accurately (see Figure 3.2).

3.7 Discussion and Conclusions

While diffusion is a local phenomenon, connectivity is global. Therefore, a model fitting for diffusion should consider goodness of fit in both aspects. We have provided a proof of concept for efficiently modeling tensor configuration distributions with MRFs and their practical MAP estimations. We believe that the MAP-MRF framework can address the two main limitations of multi-compartment models: nonlinear optimization and model selection. Note that it is easy to incorporate a noise (degradation) factor into our model. Having shown the plausibility of the approach, we plan to test it on synthetic and real data sets with different noise levels.

Chapter 4

Exploring Brain Connectivity with Two-dimensional Maps

4.1 Introduction

Diffusion-weighted MRI (DWI) enables neural pathways in the *in vivo* brain to be estimated as a collection of space curves, called a tractogram. The study of tractograms (i.e., tractography) has important applications in both clinical and basic neuroscience research on the brain. Tractograms have visual complexity proportional to the intricacy of the axonal brain connectivity and, with increasing DWI resolutions, this complexity is becoming greater and greater. It is thus often difficult for practitioners to see tract projections clearly or identify anatomical and functional structures easily in these dense curve collections. This is important because, for example, a clinical study into a neurodegenerative disease typically involves selecting more than 30 tract of interests (TOIs) manually across different datasets. Therefore, it is necessary for tractography visualization tools to provide means to reduce and help cope with visual complexity at the interaction and representation levels. We believe two concepts, *abstraction* and *filtration*, applied to representation of datasets can help users overcome the difficulties of visual complexity. While abstraction involves simplification and generalization, filtration here entails clustering and hierarchization. With these ideas in mind, we proposed planar point representations for better interaction with fiber tracts in [86], along with [37]. Two-dimensional embedding of fiber tracts provides an interesting window into the manifold space of neural connectivity. More importantly, the resulting planar point representation improves fine selection of tracts when it is coordinated with the standard three-dimensional streamtube or streamline representation in tractography tools.

A drawback of point representations is, however, that coordinate axes in the low-dimensional space lack an anatomical interpretation. It is clear from evaluations in [37, 86] that having a frame of reference, anatomical or otherwise, is important to users. Motivated by this problem, we introduced in [87] two-dimensional neural paths that have the desirable properties of low-dimensional representations while preserving meaningful and familiar coordinates. Here, we provide a unified discussion of our earlier work on planar point and path representations of tractograms. Note that, all or parts of this chapter have appeared in substantially

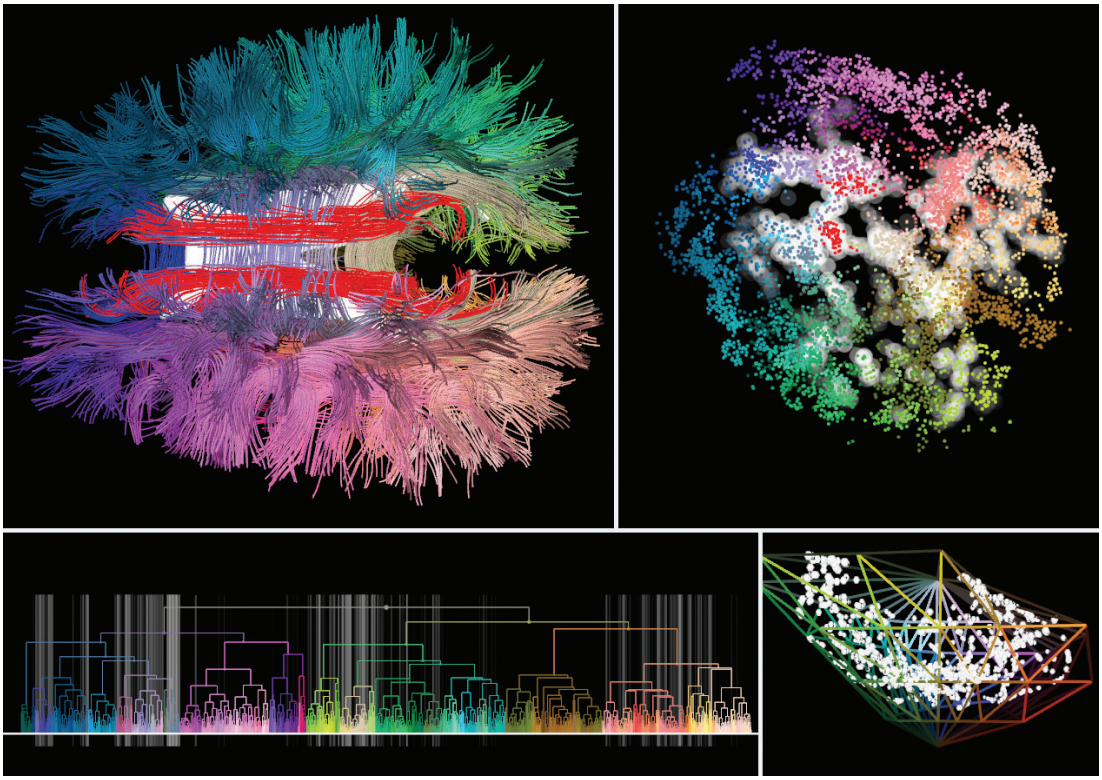

Figure 4.1: *2D point* representation linked with a streamtube representation in an interactive tractography visualization tool.

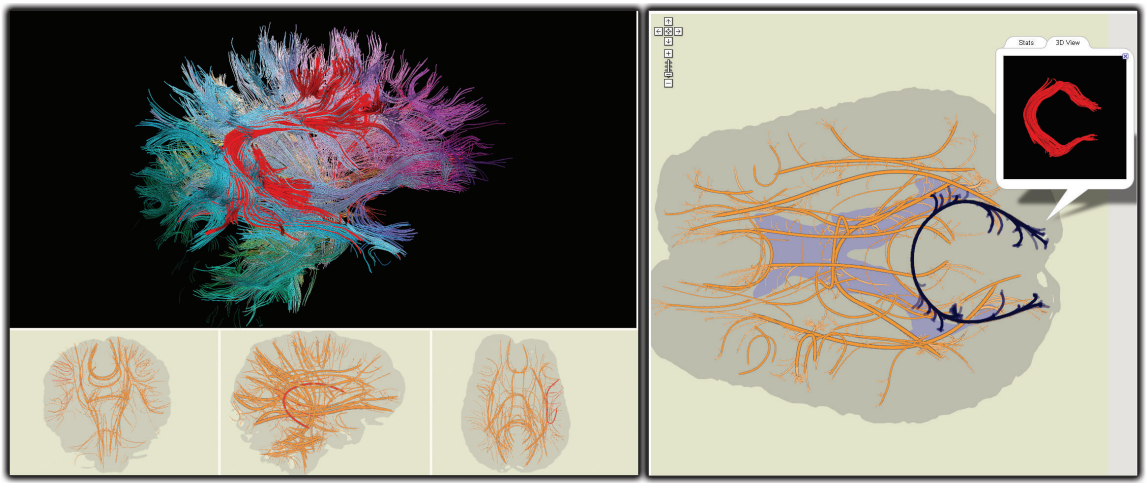

Figure 4.2: Schematic *2D path* projections of tractograms as part of a standalone interactive system (left) and as a web-accessible digital map (right). The digital map interface easily incorporates any tract-associated information, including labels, links, metrics, and statistics. Shown in the pop-up window on the right is the "brain view" of the selected tract.

similar form at different venues [47, 89, 87, 88, 86].

DWI is the source of all the data used in the work presented here. We first give an elementary introduction to DWI and then briefly discuss related work on tractography and web-based visualizations.

4.2 Related Work

Diffusion-weighted magnetic resonance imaging (DWI) measures the diffusion rate of water molecules in biological tissues *in vivo* [148]. Since tissue characteristics, geometric or otherwise, at a given point affect the diffusion rate, measured diffusion-rate information is an indicator of the tissue characteristics at the point. In particular, water in fibrous tissues such as brain white matter (i.e., a collection of myelinated axons) diffuses faster along fibers than orthogonal to them. Therefore, it is possible to estimate fiber trajectories computationally using diffusion models such as the tensor model that quantify anisotropic diffusion. Diffusion imaging based on fitting second-order tensors to DWI sequences is known as diffusion-tensor magnetic resonance imaging (DTI) [12]. Fiber trajectories are computed from DTI data by integrating bidirectionally along the principal eigenvector of the underlying tensor field. This process, called fiber tracking, yields a dense collection of integral curves (i.e., a tractogram). All the tractograms used our work were obtained by fiber-tracking in DTI volumes.

Tractograms are often visualized with streamlines or variations of streamlines in 3D [114, 166]. Reflecting the intricacy of the connectivity in the brain, these 3D models are generally visually dense. Consequently, typical interaction tasks over tracts, such as fine bundle selection, are often difficult to perform and have been a focus of recent research [2, 1]. We proposed planar point representations to improve interaction with DTI fiber tracts [86], concurrently with Chen et al. [37]. There are, however, differences between these two works. First, we use hierarchical clustering to create multiscale representations, which makes the exploration of large datasets easier, both visually and computationally. Second, while our work uses a simple force-based embedding method, Chen *et al.* uses the SMACOF algorithm, an iterative method that minimizes the metric stress of multidimensional scaling (MDS) using majorization [44]. And third, we use the embedding procedure also to create a "nice" coloring such that colors of data points perceptually reflect the structure in data points.

To address the expert concern about the insufficient anatomical context in the 2D point representation, we recently introduced 2D path representations [87], which are also projections of fiber tracts into a plane, but as planar curves rather than points.

One of the advantages of the 2D path representation is that it can be naturally integrated into web-based digital geographic map frameworks. Basic data visualization has been available on the web for many years but was usually limited to traditional techniques such as bar graphs and charts. More recently, however, visualization research started targeting this environment and advanced applications have emerged. ManyEyes [152] paved the way for everyday data visualization, with subsequent studies such as [151] and [43] proving the need for accessible web visualization. While web-development toolkits such as [23] greatly aid web visualization development, large-scale web-visualization is limited by inherent browser capabilities, as demonstrated in [91]. Alternatively, stand-alone systems have been made available in applet form or can be run

as client applications directly from websites. However, users still must control the parameters involved in producing visualizations, specify their data queries and learn the system features. This often constitutes an undesirable overhead. Yet another approach, most similar to our work from an implementation standpoint, is to use Ajax (asynchronous JavaScript and XML) technology to perform the rendering on the server side and serve images asynchronously to the client browser. The essential difference between the present work and traditional offline visualization systems is that the former separates interaction and display from rendering and computation. Our brain maps differ by eliminating users' efforts in creating visualizations and assigning this task to experts and by using the Google Maps API, which is an Ajax framework for interactive display of pre-rendered images. Closest to our work are X:MAP [161] and Genome Projector [8], which present genome browser tools implemented using the Google Maps API. We extend this idea to a new domain and demonstrate its usefulness for tractography datasets.

4.3 Methods

Both point and path representations are projections of fiber tracts onto the plane: Each tract is represented with a 2D point in the former and with a 2D curve in the latter. Formally, given a polyline form of a fiber tract $C \in \mathbb{R}^{n\times 3}$ with n vertices in 3-space, its point representation is obtained with a map $\pi_{point} : \mathbb{R}^{n\times 3} \rightarrow \mathbb{R}^2$ and its path representation is obtained with another map $\pi_{path} : \mathbb{R}^{n\times 3} \rightarrow \mathbb{R}^{m\times 2}$.

Generation of these two representations shares three common steps. First, we obtain a whole-brain tractogram by fiber tracking in a diffusion-tensor volume fitted to a given DWI brain sequence. Second, we compute similarities between all pairs of tracts within the tractogram, obtaining a similarity (or affinity) matrix. Third, using the similarity matrix from the previous step, we run a hierarchical clustering algorithm on the tractogram, obtaining a clustering tree (i.e., dendrogram).

We create the 2D point representation of the tractogram by embedding the tracts in the plane with respect to the similarity matrix, using a simple iterative force-directed method. We use the hierarchical clustering tree to create multiscale point representations.

In the case of the path representation, we first pick a cut on the clustering tree and obtain a clustering. Then, by treating cluster centroids as pivots, we create projections of tractograms onto the major orthogonal planes as curves. We render these 2D curves stylistically using heuristics determined by the topology and geometry of the corresponding tracts and tract clusters.

We give details of these steps in subsequent sections.

4.3.1 Image Acquisition and Fiber Tract Generation

DWI brain datasets used in this chapter were acquired from healthy volunteers on a 1.5T Siemens Symphony scanner with the following acquisition parameters in 12 bipolar diffusion-encoding gradient directions: thickness = 1.7 mm, FOV = 21.7 cm $\times$ 21.7 cm, TR = 7200 ms, TE = 156 ms, b = 1000, and NEX = 3. For each DWI sequence, the corresponding DTI volume was then obtained by fitting six independent parameters of a single second-order tensor at each voxel to the twelve measurements from the DWI sequence [12]. We

generate fiber-tract models of the whole brain by integrating (second-order Runge-Kutta integration) the major eigenvector field of the diffusion tensor field bidirectionally starting at seed points. We integrate with a constant step size of 0.5 mm and stop integration when a gray-matter area is reached or the linear anisotropy or signal-to-noise (SNR) ratio at the current point passes predetermined thresholds.

4.3.2 Measuring Similarities Between Fiber Tracts

We quantify the similarity between two tracts using a distance measure. This measure tries to capture how much any given two tracts follow a similar path, while giving more weight to the points closer to tract ends. Given two integral curves $C_i = \{C_i^1, \ldots, C_i^m\}$ and $C_j = \{C_j^1, \ldots, C_j^n\}$ that are represented as polylines with m and n vertices respectively, we first find mean weighted distances d_{ij} and d_{ji}, and then determine the maximum of these two distances as the distance D_{ij} between the two curves:

$$d_{ij} = \frac{1}{m}\sum_{k=1}^{m} \alpha_i{}^k \text{dist}(C_i^k, C_j) \tag{4.1}$$

$$d_{ji} = \frac{1}{n}\sum_{k=1}^{n} \alpha_j{}^k \text{dist}(C_j^k, C_i) \tag{4.2}$$

$$D_{ij} = D_{ji} = \max(d_{ij}, d_{ji}) \tag{4.3}$$

The function $\text{dist}(p, C)$ returns the shortest Euclidean distance between the point p and curve C. Also, $\alpha_k = \frac{1}{Z} e^{|k-(m+1)/2)|^2/\sigma^2}$, where the normalizing factor $Z = \sum_{k=1}^{m} e^{|k-(m+1)/2|^2/\sigma^2}$. We set the parameter σ automatically, proportional to L_C the length of the fiber tract, such that $\sigma = \lambda L_C$, where $\lambda \in (0, 1]$. We set $\lambda = 0.5$ for the datasets used for this chapter.

We compute distance between each pair of integral curves as we denoted and assemble the measures to create a distance matrix. Note that our measure is symmetric and non-negative but does not necessarily satisfy the triangle inequality, therefore, it is not a metric. While our approach is independent of a particular similarity measure, good results in practice require a good similarity measure – one that reflects users' understanding of the similarity between data points (i.e., tracts) and works well for the task at hand.

4.3.3 Clustering

For a given tractography dataset we compute a clustering tree using an average-linkage hierarchical clustering algorithm on the tract distance matrix (e.g., [58]). We choose the average-linkage criterion because it is less sensitive than the minimum-linkage to broken tracts due to tracking errors. The output of the clustering algorithm is a hierarchical tree called dendrogram. The height of the tree can be thought as the radius of the bounding ball of the dataset–in the units of the similarity measure used. And any horizontal cut on this tree provides a clustering of the dataset. Therefore, for example, the root node represents a clustering with a single cluster containing all the data points. Conversely, the leaf nodes correspond to a clustering where every data point is a cluster.

We obtain a clustering of tracts by manually setting a cut threshold on the dendrogram. This threshold can be also interactively changed by user to control the coarseness of the clustering. A constant cut at 60% of the clustering tree's height gave consistent results across the six datasets we experimented with.

4.3.4 Planar Projections of Fiber Tracts

Fiber Tracts as Embedded Points

We use a simple iterative force-directed method for embedding tracts in the plane [59]. Embedding refers to a one-to-one smooth mapping from fiber tracts to points in the plane that preserves the "structure" of the fiber tracts. In this context, the distance matrix computed can be considered as a manifestation of the structure in the fiber tract space.

For a given dataset with M tracts, we start with M corresponding points, all placed at the origin initially. We then iteratively adjust the positions of these points by moving the pairs of points closer to or further from each other to match with the corresponding D_{ij} entries in the similarity matrix. To achieve an interactive performance, we use a stochastic sampling technique described in [36] for updating the "forces" between data points. Briefly, instead of computing forces on a point x_i from every other point in the dataset, we limit the points acting on x_i to $x_j \in \mathscr{F} = \{N_i \cup S_i\}$, where N_i and S_i are disjoint sets with a constant size. We iterate over data points and resample S_i each time by uniformly randomly selecting points from the whole dataset. For a randomly selected point x_k, if D_{ik} is smaller than $\max_{x_l \in N_i} D_{il}$, the maximum distance from x_i to any member of N_i, then x_k is assigned to N_i, otherwise it is assigned to S_i. N_i loosely represents the neighborhood of the point iterated.

How Embedding Can Be Used to Color Fiber Tracts

Given a similarity measure, a good coloring of fiber tracts should reflect the similarities between the tracts such that similar tracts are assigned to similar colors and different tracts are assigned to different colors. Embedding fiber tracts in perceptually-uniform color spaces, which are subsets of $\mathbb{R}^3$, provides a practical way to approximately achieve this goal. A perceptually uniform color space is an empirically-constructed color space where the Euclidean distances between color triplets are approximately proportional to perceptual differences between them. L*a*b* and Luv are two common examples of such color spaces [65].

Embedding fiber tracts in the L*a*b* space is the general coloring scheme used in the interactive tools presented here. For this, we compute an approximation of the L*a*b* color gamut, as visible on the right panels of Figure 4.6, and use it as a container for force directed embedding. To avoid having to adjust a repulsive container force, which would likely need a hard-to-control, steep gradient, we perform a physically accurate simulation with container contact detection. The embedding begins in the center of the gamut and is gradually expanded until most of the space is filled. During implementation we observed that the largest distances are often embedded along the luminance axis, the vertical (in paper coordinate frame) axis, of the color gamut. This is problematic because luminance offers little resolution and can be interpreted as a lighting effect. We therefore apply a "flattening" force at the beginning of a simulation cycle to force large distances to lie in the horizontal plane (a*b*-plane). These force components, acting on the luminance axis towards the center of the gamut, wear off as the embedding moves towards a steady state. The force computation used is the same as for the 2D embedding, with straightforward 3D modifications. In terms of interaction, the color embedder only supports collapsing and color grabbing.

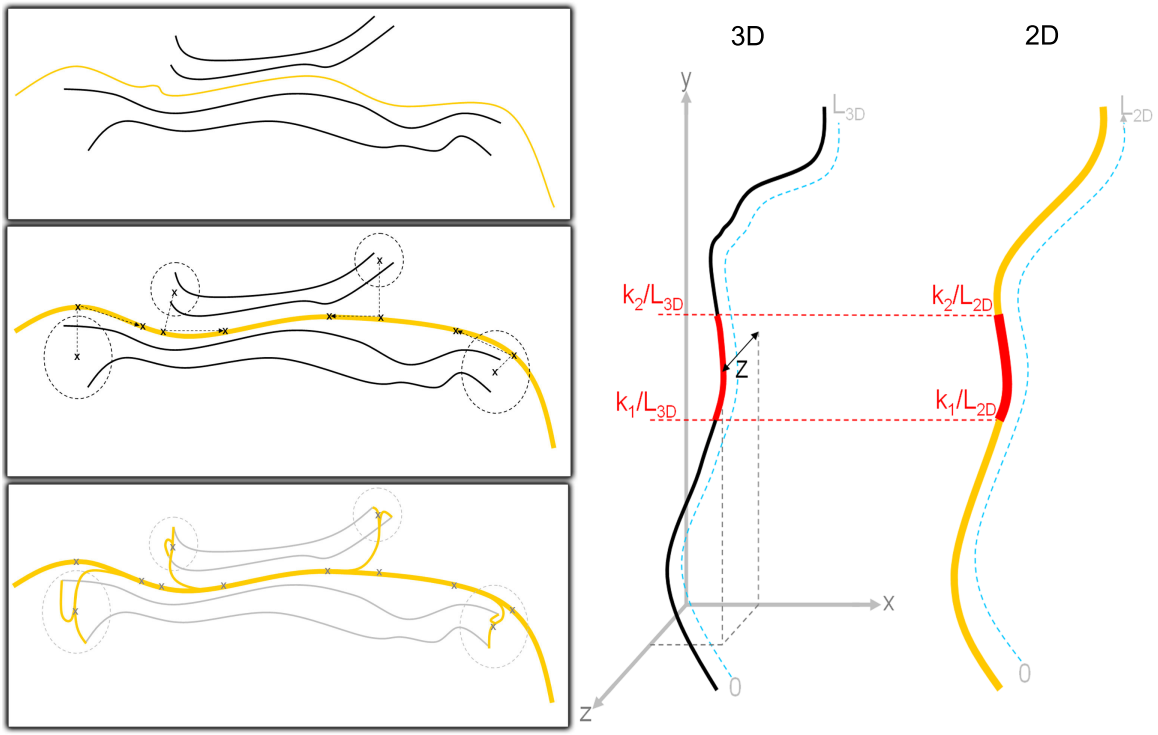

(a) (b)

Figure 4.3: (a) Schematic tract-cluster representation. (Top) 2D projections of a tract-bundle, with an associated centroid curve (orange), are determined from a hierarchical clustering of initial 3D tracts. (Middle) The centroid curve is smoothed by a spline and the endpoints of non-centroid curves are clustered using their initial 3D coordinates (four clusters); for each cluster, three control points linking the center of the cluster to the centroid spline are computed. (Bottom) Splines are run from each curve endpoint through the control points of its corresponding cluster. (b) Depth ordering of 2D paths. For each segment of a 2D spline, we locate a corresponding segment on the 3D curve from which the spline was derived by traveling the same fractional distance along both curves. The depth of the 2D segment is the same as the depth of the middle of its corresponding 3D segment.

Fiber Tracts as Planar Curves

For a given tractography dataset, we create schematic views of fiber tracts projected on the sagittal, coronal, and transverse planes.

We start with obtaining a clustering of the dataset by picking a cut on the hierarchical clustering tree computed earlier. We then create simple orthogonal projections of tracts on each plane. Suppose the sagittal plane is aligned with the xy-plane. And let $v = (x, y, z) \in \mathbb{R}^{n\times 3}$ be a vertex of a tract. Then the projection onto the sagittal is given by this simple equation $\pi(v) = (x, y)$. We cull out tracts that do not contribute significantly to the projection. If the ratio of projected tract length to true tract length is under a threshold value, we remove the tract from the corresponding cluster. We set the culling threshold to 0.65 for the projections used in our experiments. Finally, we compute a centroid for each cluster by choosing the tract with the smallest maximum distance to any other tract in the cluster. We found that for illustration purposes it is desirable to avoid broken

tracts. We therefore weigh the centroid selection to favor longer tracts by dividing the maximum distance from each tract to any other tract by the tract's length.

We opted for a non-photorealistic rendering of brain projections to avoid their interpretation as 3D views and to harness user's intuition about 2D maps. The rendering assumes a given clustering with assigned centroid tracts, which can be computed as described in the previous section. The centroid tracts will define a schematic neural skeleton on top of which the non-centroid tracts are scaffolded. Projections of centroid curves are smoothed prior to rendering to achieve a schematic representation and to reduce clutter. This is done by sampling a number of evenly distributed control points (5 in our implementation) along the tract projection and using them as control points for a piecewise cubic spline with 30 segments. The thickness of a centroid curve is proportional to the square root of the number of tracts in the bundle. Once centroid tracts are represented as 2D splines, endpoints of non-centroid curves are linked to their cluster's centroid spline following the procedure illustrated in Figure 4.3a. First, the end-points of non-centroid curves in a bundle are clustered based on the end-points' initial 3D coordinates. Two endpoints are placed in the same cluster if the distance between them is less than 2 mm. Then, for each such endpoint cluster we compute three control points that link the geometrical center of the endpoint cluster to the centroid spline: the first point is the center itself, the second is a point on the centroid spline closest to the center point, and the third is determined by traveling from the second point down the centroid spline, towards each curve's other endpoint, for a predefined distance (e.g., half of the distance between the first two points). Ultimately, splines are run from each tract endpoint through its cluster's three control points, thus linking each endpoint to the centroid path. The thickness of these endpoint linkage splines gradually increases from unit thickness (i.e., single-tract thickness) at the tract endpoint to a thickness proportional to the square root of the endpoint cluster size, where it merges with the centroid spline.

We depth-order spline segments so that 2D centroid splines crossings can indicate the depth ordering of their corresponding 3D shapes. The depth ordering is done differently for centroid splines and non-centroid splines, since while centroid curves are close representations of actual 3D tracts, non-centroid curves are abstract representations obtained through the process described above. Furthermore, the depth ordering is approximate (as discussed in the following paragraph) and may produce artifacts. For centroid splines, the depth of a spline segment is computed by finding a matching segment on the 3D tract from which the spline was derived, and taking the depth of that segment's center (see Figure 4.3b). The matching segment on the 3D tract has its endpoints at the same fractional distance from the start of the 3D tract as the 2D segment's distance from the start of the 2D spline. The depth of any non-centroid spline is determined by averaging the depth of the corresponding 3D tract.

In the following two sections, we give details on how we use 2D neural path representations as part of an interactive application and as standalone digital maps.

4.3.5 Linked Multiview Interaction

We expect a typical use of low-dimensional representations to be as part of interactive applications where views and interactions of conventional representations are linked with that of low-dimensional representations. We have developed two interactive visualization systems using the 2D point and path representations,

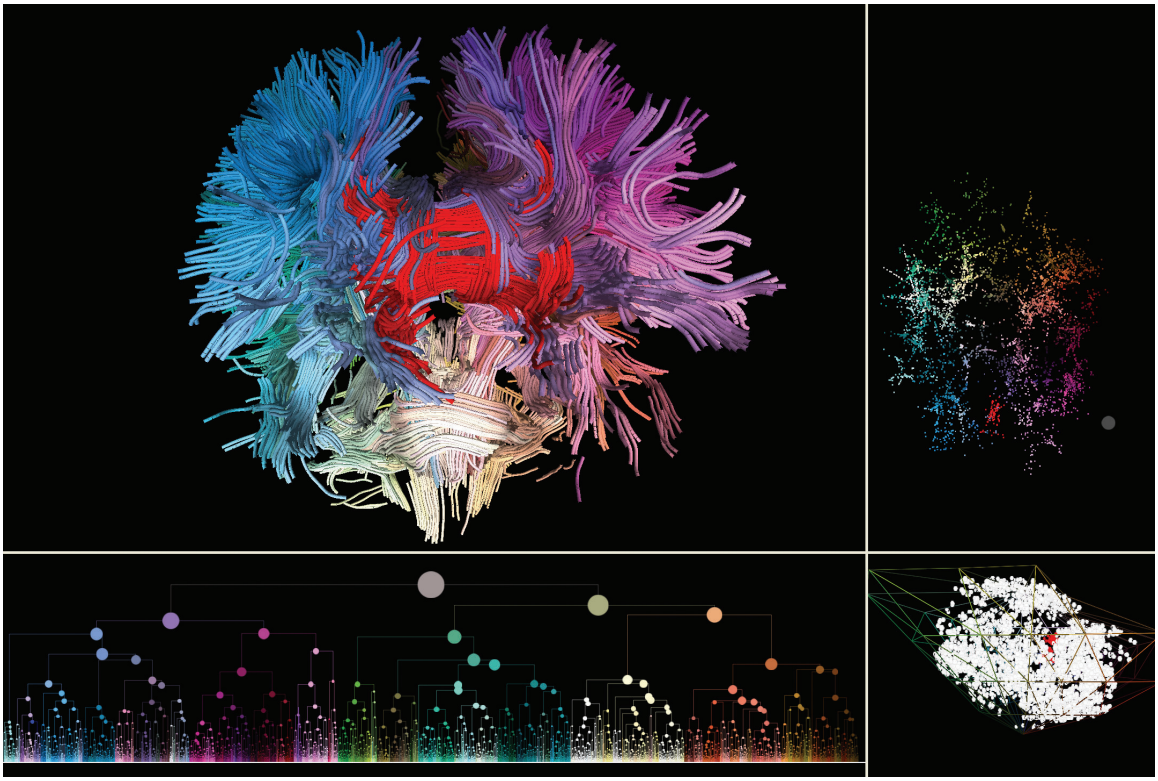

Figure 4.4: Coordinated DTI tractogram model exploration in lower dimensional visualizations: 2D embedding (upper right), hierarchical clustering (lower left), and L*a*b* color embedder (lower right). A selection of a fiber bundle (red) in the hierarchical clustering is mirrored in the other views.

Figure 4.5: An interactive analysis system using linked views and planar tract-bundle projections. Three planar representations, along the coronal, transverse and sagittal planes (bottom panels), are linked to a 3D streamtube model (upper left) and a 2D point embedding of tract similarities (upper right). Selections in the projection views can be performed by clicking or cutting across cluster curves and are mirrored in the 3D view. Points corresponding to the selected tracts are interactively embedded into the plane and used to refine selections at tract level.

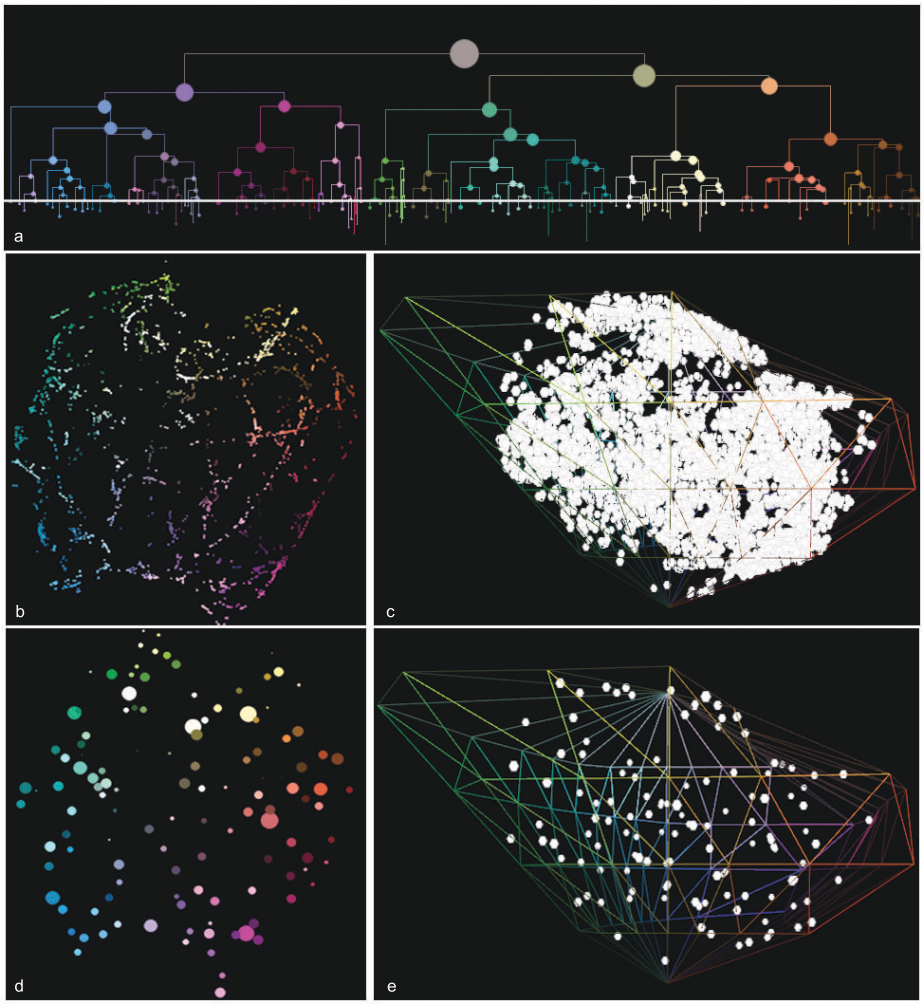

Figure 4.6: A clustering cut in the dendrogram view (top row) is applied to the linked 2D embedding and 3D colorer (middle row). Points belonging to the same cluster are collapsed to their centroids (bottom row).

respectively, to demonstrate this mode of use (see Figures 4.4 and 4.5).

Both applications have a view of tractography data visualized using 3D streamtubes. Coloring is generated through the embedding of tract similarity into the L*a*b* color space. In addition to the standard 3D viewing interactions, we have two basic 3D selection/deselection interactions on streamtube models: sphere-selection and brushing. Sphere-selection, like box-selection, enables user to select the intersecting tracts by moving a sphere of desired radius. Brushing lets user draw 2D curve on the viewing plane and select the intersecting tracts.

Both sphere-selection and brushing can be used to further prune the current selection but they cannot be used to grow it. For that purpose, we provide a selection growing interaction that gradually adds tracts closest to the current selection. Proximity is determined again by the distance measure discussed above.

On the 2D point representation, we provide point selection, point collapsing. Selection is performed by

clicking and dragging; multiple selection can be performed to select points from non-adjacent regions. Collapsing groups a set of points into a single clustered representation. This can be used either for easier tract bundle selection or as a mechanism for manually refining embeddings–points belonging to the same tract bundle can be grouped together if the embedding algorithm places them apart. The centroid of the grouping will be used in subsequent embedding iterations. The hierarchical clustering tree provides a filtration of the dataset via cuts. Figure 4.6 shows how a cut on the dendrogram, which results in a clustering, provides a coarser representation. Both the 2D embedding and the 3D colorer clusterings can be further refined interactively after a cluster cut, by expanding or collapsing individual clusters.

In the 2D path tool, we link projective views on the sagittal, coronal, and transverse planes to a standard 3D streamtube model. The clustering cut threshold that defines the specificity of the projected bundles can be altered interactively during visualization. Tract clusters in the planar projections can be selected by drawing line segments that select intersecting bundles. A selection in any of the planar views is mirrored in the 3D model view as well as all other 2D projections.

4.3.6 Digital Map Interface

Brain mapping is one of the quintessential problems in neurosciences. We believe that a geographical map metaphor is well-suited to the visualization and analysis of results obtained in that direction. Therefore, having a representation of the brain that is viewed, interacted, queried, and enriched like an online geographical map was one of the motivations behind our creation of 2D path representation.

For this, we use the Google Maps API, an Ajax framework used to render large maps, to interactively display our tractogram maps on the web. The Google Maps API receives input image data in the form of a set of small images, called tiles, that when assembled together form the different zoom levels of the map. Each zoom level, z, consists of a rectangular grid of tiles of size $2^z \times 2^z$. The API decodes the zoom level and coordinates of the currently viewed map region to retrieve and display the visible tiles. The developer can load a custom set of tiles in the API by implementing a callback function that translates numerical tile coordinates and zoom level into unique paths to the custom tiles.

The API provides basic functionality such as zooming and panning and allows programmatic extension or customization with markers and polyline overlays, information pop-ups and event management. The API can easily be integrated into any webpage supporting Javascripts.

Our visualization system can render our 2D projections into a set of image tiles instead of the screen. For each cluster, including both tract-bundled and endpoint clusters, we export information required for interaction and browsing. Selection information consisting of evenly spaced points along splines and thickness radii for splines contained in a cluster is exported. In line with the tile paradigm, instead of exporting this information to a single large file, we divide it geometrically across corresponding tiles and write it as multiple tile-content text files. Upon user selection, the content file of a clicked tile is fetched from the server and its data analyzed for an intersection. This approach avoids loading and searching through large files. A valid cluster selection is marked on the map with polyline overlays running over tract splines contained in the selected cluster (see Figure 4.7). For this purpose, spline coordinates for each cluster are exported to files indexed by a unique cluster identifier.

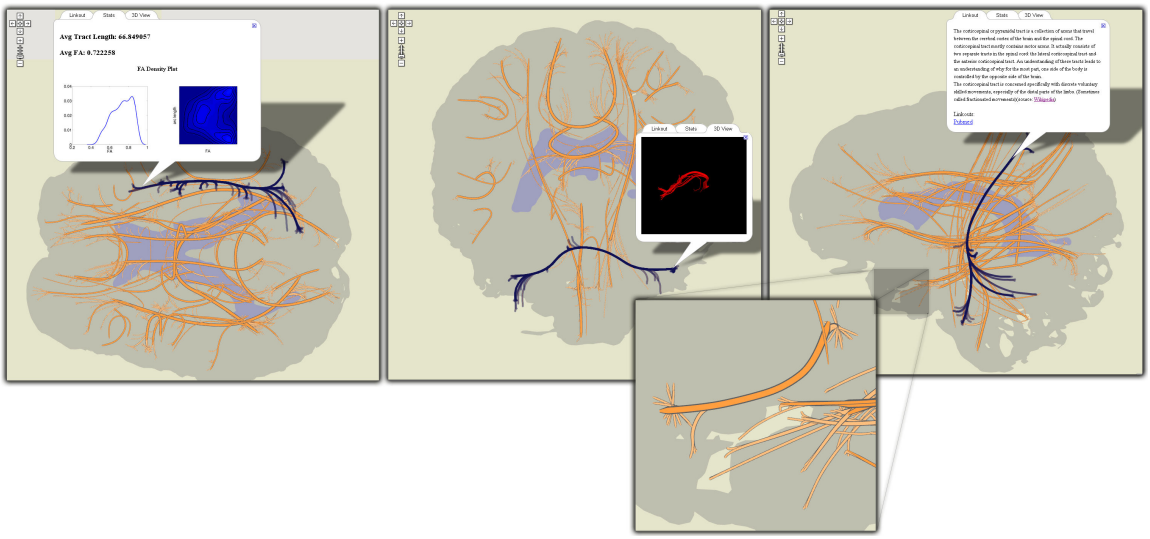

Figure 4.7: DTI tractography data projected onto the sagittal, coronal and transverse planes. Major tract bundles are represented schematically by their centroid tract; individual tracts in bundles are linked from the centroid bundle to their projected end points. Zooming in allows access to smaller clusters of tracts. Bundles can be selected and pre-computed statistical data along with 3D views of the tract bundle ("brain view") can be displayed.

Finally, for each tract cluster we export a variety of metadata accessible during map browsing in information boxes, as shown in Figure 4.7. A short description and links to the most relevant publications or research can be manually added for major tracts. A few 3D poses of each tract bundle are prerendered and exported as animated GIF images, indexed by the cluster identifier. Statistical data, in both textual and graphical form, are computed for each cluster and written as HTML content to cluster indexed files. This information is loaded and displayed in tabbed information boxes at the user's request.

4.3.7 Implementation

We implemented both interactive systems in C++ using G3D and Qt libraries [70, 132]. We created the web interface for the neural path representation using the Google Maps API [75]. It can be accessed via the url link [169].

4.4 User Evaluation

We compared the two representations both anecdotally and quantitatively.

4.4.1 Anecdotal Study: Methods and Results

In the anecdotal study we showed a prototype that implements both 2D point and neural path representations to 3 neuropsychologists. They were all interested in the relationship between fiber tracts and cognitive and

behavioral function in the brain. Similarly, all have used computational tools for analyzing DTI data, though only one of them has used fiber-tract visualization tools in his clinical research. The participants had research interests in vascular cognitive impairment, early Alzheimer's disease, and HIV, focusing on specific tracts and regions such as the corpus callosum (CC), frontal lobe, basal ganglia, cingulate bundle, superior and inferior longitudinal fasciculi, anterior internal capsule, and the uncinate fasciculus.

Our anecdotal evaluation protocol was straightforward: we demonstrated the prototype while asking questions and collecting participant's feedback. Two of the experts also tried both interfaces themselves by selecting a set of major TOIs, the CC, cingulate bundle, uncinate anterior internal capsule, and the corticospinal tract. There was agreement that the 2D neural path representation was more intuitive and easier to use and learn than the 2D point representation.

Our experts also found the web interface with the digital map interaction useful. Although they believed that the standalone application with linked representations would remain necessary for quantitative analyses that require interactive fine selection, they thought the web accessibility opened up interesting possibilities. They were particularly excited about browsing through datasets while commuting or at home, of quickly inspecting unfamiliar datasets, and of sharing such visualizations with collaborators.

4.4.2 Quantitative Study

In the quantitative evaluation, we compared the point and path representations by measuring user performance on a bundle selection task.

Applications

We used our interactive path representation system (Section 4.3.5) and the 2D point representation tool presented in [37] for running the comparative study. The reason that we used this tool instead of ours is that it was already compared to other known tractography tools such as CINCH, MedINRIA, and BrainApp, and reported to be preferable over them. Using it would give us some idea about, if any, relative merits of the 2D path representation tool over not only the 2D point representation but also these other tractography applications compared before. Chen *et al.*'s application offers a brush tool that works similarly to ours in 3D and as a lasso tool on the 2D point representation. Users can select tracts or points and then remove them or, conversely, remove everything else from a current selection.

Participant Pool

We had 4 subjects all familiar with neuroanatomy and tractography. They also had experience with one or more tractography visualization tools. Our first subject was a neuroscience graduate student, working on tracing white-matter tracts from frontal subregions to basal ganglia and the medial temporal lobe. Our second user was a postdoc in neuropsychology and had five years of experience with DWI in clinical research. This user, who participated in the anecdotal study as well, studied white matter variation with neurodegenerative diseases as specified above. Our third subject was a biomedical engineering graduate student and had

	time (secs)				confidence			
	cb	cst	slf	mean	cb	cst	slf	mean
2D point	227	361	234	274	4.1	3.3	3.1	3.5
2D path	136	165	215	172	4.1	3.8	3.7	3.9

Table 4.1: User performance on bundle selection task.

significant tract-selection experience working as a rater for a neuroscientist. Our last subject was a computer-science graduate student doing research on computational DWI algorithms. Two of the users were male and two female.

Task

We measured user performance on bundle selection, a typical real world task in tractography tools. Users were asked to select three major bundles, the cingulate bundle (cb), corticospinal tract (cst), and right superior longitudinal fasciculus (slf), in two different brain datasets. We choose these bundles because they represent the easy-to-hard selection-difficulty range well and were used for evaluation in [37].

For each system, we explained to users the underlying visualization concepts and demonstrated the basic interactions, mainly involving brushing on 2D and streamtube representations. After this introduction, users were asked to select the bundles (cb, cst, and slf) on two different training datasets. Following training, the users performed the task on two different test datasets while we collected their task-completion times. After each selection they provided subjective confidence estimate in the range 1-5 (1: not confident, 5: very confident) for their selection. They could give fractional estimates. After completing the task on both systems, users were asked to complete a post-questionnaire giving qualitative feedback on their experience. Half the users performed the task first on the 2D point-representation tool and the other half on the 2D path tool.

Factors and Measures

The sole factor considered in our quantitative experiment was the type of low-dimensional representation: 2D point and 2D path. All subjects used both types of representation. We recorded the users' bundle-selection times and subjective confidence values as measures of performance.

Results

In order to understand if the differences between user performances on the two tools were significant, we ran the paired t-test on our measurements. Results show that users were significantly faster on the 2D path tool than the 2D point tool ($p = 0.02$). Users were also significantly more confident with using the 2D path representation than the 2D point representation ($p = 0.01$). Table 1 summarizes users' overall and per-bundle mean performances on each tool. Figure 4.8 shows the difference between the means of performance measures per user (2D-path-performance values are subtracted from 2D-point-performance values) and the mean over users. Errors bars indicate the standard error of per-user differences.

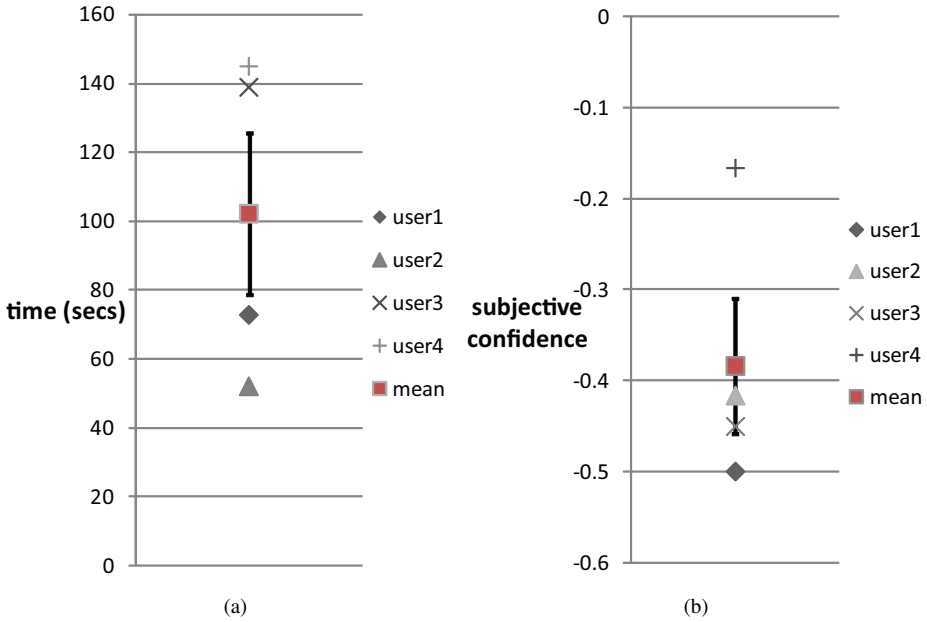

Figure 4.8: Per-user differences between (a) time and (b) confidence measurements with the two tools. Differences are obtained by subtracting 2D-point-tool performance values from 2D-path-tool performance values. Red squares show the mean performance difference between the tools. Errors bar around the red squares indicate the standard error of per-user differences.

We observed some interaction patterns worth reporting. We noticed two distinct selection strategies used with the 2D path tool. Two of the users consistently brushed over large areas of the projection to ensure that the targeted bundle was selected and then relied on the 3D view to clean up the selection. The other two users aimed for fine selections in the 2D projections and then inspected the 3D view to determine whether any fibers the selection. They added the missing tracts using short, targeted brush strokes and then removed tubes that were erroneously added during this operation. These users seemed to have a better understanding of the mapping between the 3D view and the 2D projections, perhaps explaining the difference in strategies.

All subjects used the 2D point representation relatively rarely. The most common operation was to remove points they were completely confident were not part of the selection (e.g., half of the brain, or peripheral U-shaped bundles). However, in the absence of a clear contextual mapping between the 2D point and streamtube views, subjects were hesitant to perform bold operations in 2D, at least in the short run.

4.5 Discussion

It is important to note that our representations rely on the anatomical fidelity of the intermediate results at each step. For example, broken trajectories due to fiber tracking errors can reduce the effectiveness of the representation. Similarly, our method expects the clustering algorithm and similarity measure to provide anatomically plausible results. However, it is difficult for a single distance measure to capture the anatomical

similarity completely. Furthermore, on the same data, a good similarity measure for one purpose can be entirely irrelevant for another. While the choice of similarity measure makes clustering a subjective task, clustering algorithms themselves also have intrinsic limitations [98].

One potential limitation of the planar path representation is that bundles surrounded by the other bundles similar in orientation and shape may not be clearly visible. While we have not found this to be an issue in practice, moving projection planes along major axes while restricting the projecting tracts to a volumetric window moving with the projection plane can help solve potential problems. Also, results of user studies should be taken with a grain of salt. In general, it is difficult to run experiments that vary one factor while keeping all the other factors constant. For example, an earlier study [37] compared the 2D point representation tool with other tractography applications and reported that users were faster with the former. However, in our user evaluation, we observed that users rarely used the 2D point representation and the brush tool dominated their interaction. This brings up an open question of whether the performance difference in the reported evaluation was due mainly to the brush tool or to the 2D point representation. An experiment that replaced the brush tool with a more standard box-selection tool, say, might resolve this question. In either case, we believe that abstract representations, including the 2D point representation, are useful in the long run, as users gain more experience with the mapping between brain tractograms and low-dimensional representation primitives. In general, however, it is not realistic to expect practitioners to learn the correspondence between the new representation and the actual fiber-tract collection quickly, unless the tools are easily interpretable using a conventional anatomical framework. Furthermore, in order any tool using a new representation of tractograms to have a clinical relevance, it should provide anatomical context and intuitive functionality for region of interest (ROI) analysis on both conventional and new representations.

While we have focused on planar spatial representations here, it is possible to create abstract representations of tractograms. For example, the hierarchical clustering tree itself can be considered as a representation of the tractogram. Or consider the circular map of connectivities (or dependencies) shown in Figure 4.9. We obtain this dependency graph representation by first clustering tract end points using hierarchical clustering and setting an implicit dependency between the end points of each tract. We then visualize the resulting hierarchical tree with pair-wise connectivities using hierarchical edge bundling [82]. Feedback from a neuropsychologist suggest that it might be useful for understanding connectivity densities and profiles.

Although DWI is the only imaging protocol to estimate the brain neural architecture *in vivo*, there are *in vitro* imaging techniques, such as the three-dimensional electron microscopy particularly used in the emerging field of connectomics, with which it is possible to extract neural structures on much smaller scales (e.g., individual axon bodies) [85]. We believe that the general ideas as well as specific techniques presented in this chapter can extend to the visualization and analysis of visually complex axonal structures originating from these high-throughput imaging techniques.

4.6 Conclusions

Combining traditional 3D model viewing with intuitive low-dimensional representations with anatomical context can ease navigation through the complex fiber tract models, improving exploration of the connectivity

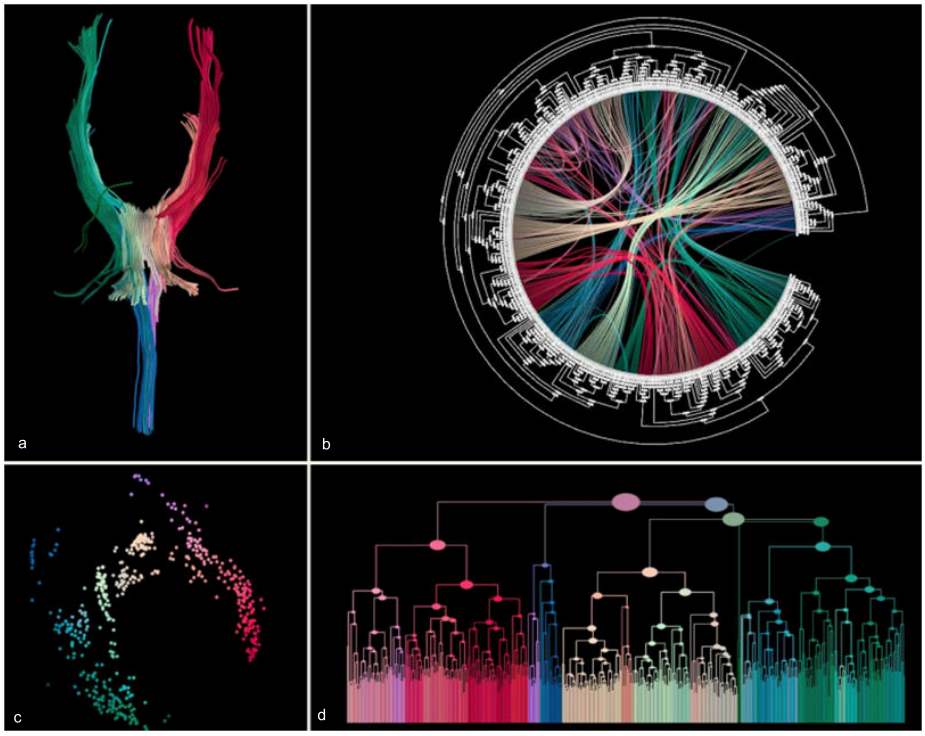

Figure 4.9: Low-dimensional representations can have varying degree of abstractions. The internal capsule represented with (a) streamtubes, (b) a circular connectivity map, (c) a set of embedded points in the plane, and (d) a hierarchical clustering tree.

in the brain. We have presented two planar maps, point and path representations of tractograms, that facilitate exploration and analysis of brain connectivity. Both representations are essentially applications of abstraction and filtration concepts to tractograms. We achieve abstraction by simplifying and generalizing fiber tracts, both geometrically and topologically, with points and schematic curves in the plane. We create filtrations of tractograms by computing hierarchical clustering trees. These help create better abstractions as well as provide a multiscale view of the data, which is important in reducing visual complexity and noise. We compare the two representations both qualitatively and quantitatively with help of experts. Results suggest that the 2D path representation is more intuitive and easier to use and learn than the 2D point representation.

We also introduced a novel way of making tractography data accessible by publishing neural maps online through a digital map framework. Our representation is conducive by construction to such a geographic map interface. This interface leads to new possibilities for enriching tractography datasets using the mass knowledge base available on the web. User feedback indicates that our web interface can be particularly useful for browsing unfamiliar datasets quickly, for analysis tasks that do not require fine selection and for sharing data in collaborative settings.

Part II

VISUALIZATION

Chapter 5

Coloring 3D Line Fields Using Boy's Real Projective Plane Immersion

5.1 Introduction

It is often useful to visualize a *line field*, a function that sends each point P of the plane or of space to a line through P (see Figure 5.2a-b); such fields arise, for example, in the study of tensor fields, where the principal eigendirection at each point determines a line. The direction is a line rather than a vector because diffusion is spatially symmetric; if $\mathbf{v}$ is an eigenvector, so is $-\mathbf{v}$. Second-order tensor field data is common in physics, engineering, medical imaging, and structural imaging. Line fields also appear in computer vision, material science, and chemistry as edge, orientation, or rotation fields.

To visualize a line field, we often assign a color to each line, typically according to its direction. This can be thought of as a mapping from the set of all line directions to color space. When we color a line field, there are some natural desiderata:

- Smoothness: the color should change smoothly as the line changes smoothly; i.e., the color-triple

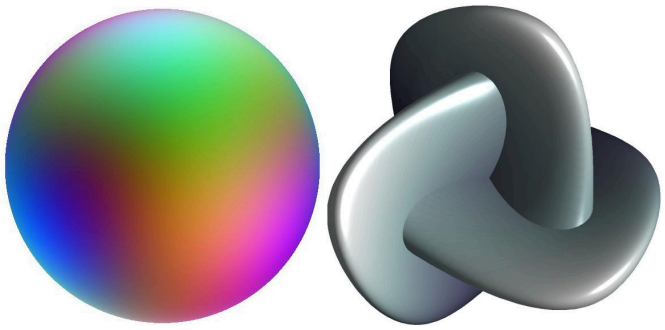

(a) (b)

Figure 5.1: (a) Sphere colored by immersing RP^2 in RGB color space. (b) Boy's surface

should be a smooth function of the coordinates of the direction vector of the line;

- Locality: lines with similar directions should map to similar colors;
- Uniqueness: lines with very different directions should map to very different colors.

Previous methods for line field coloring suffer from either discontinuities or uniqueness problems (i.e., multiple lines corresponding to a single color). One reason is that the problem, as formulated, has no solution! Finding a mapping from RP^2 to the space of color triples, which is topologically a subset of $\mathbb{R}^3$, while satisfying the above criteria is equivalent to embedding RP^2 in 3-space. But it is possible to show that no such mapping exists [107]: RP^2 is a nonorientable surface, it admits no embedding in 3-space.

Our primary contribution is introducing Boy's RP^2 immersion as a color model for line fields. The resulting coloring method is smooth and one-to-one except on a set of measure zero. We also present an unpublished parametrization of Boy's surface. Note that, all or parts of this chapter were originally published in substantially similar form at [46, 45].

We apply the proposed model to color the orientation in slices of diffusion tensor magnetic resonance imaging (DTI) brain volumes (a second-order tensor field) and integral curves, which represent the fiber tracts, extracted from these volumes. DTI provides a non-invasive means to measure microarchitecture of fibrous tissues induced by the Brownian motion of water molecules [12]. Compare to conventional MRI, DTI makes possible the acquisition of much richer anatomical information about the white matter structures (myelinated axons) in the brain. During the last decade this has led to a number of important applications for clinical and neuroscience research into the health of the brain [116].

5.1.1 Related Work

Color mapping of line fields has been studied in the context of diffusion tensor magnetic resonance imaging (DTI) visualization widely. There have been several methods based on rotational and mirror-symmetry assumptions. An extensive study of different coloring methods for visualizing DTI can be found in [124]. One standard method for visualizing DTI slices is to map the absolute values of the *xyz*-coordinates of principal eigenvectors directly to RGB color-triples [128]. We refer to this method as the *absolute value method* throughout this chapter. Clearly, this does not have uniqueness: there are 4 different orientations associated to the color (r,g,b) whenever r, g, and b are all nonzero (Figure 5.2-c).

However, the absolute value method provides an easy-to-read color model, as it aligns the major colors with the major axes. As the data also tend to be aligned with these major axes, colors can often be compared across datasets. There have been other methods proposed in different contexts such as visualization of material microstructures [160] and rotation fields; since every nonzero 3D rotation has a unique axis of rotation a rotation field induces a line field [6]. All these methods, however, suffer either from uniqueness or continuity problems. The coloring method proposed here is continuous and one-to-one almost always. Furthermore, in the few cases where uniqueness is not satisfied, context usually provides disambiguation.

There have been different parametrizations proposed for Boy's surface. The first parametrization was found by Morin to be used as a half-way model for sphere eversion [117, 67]. Based on Morin's work, his

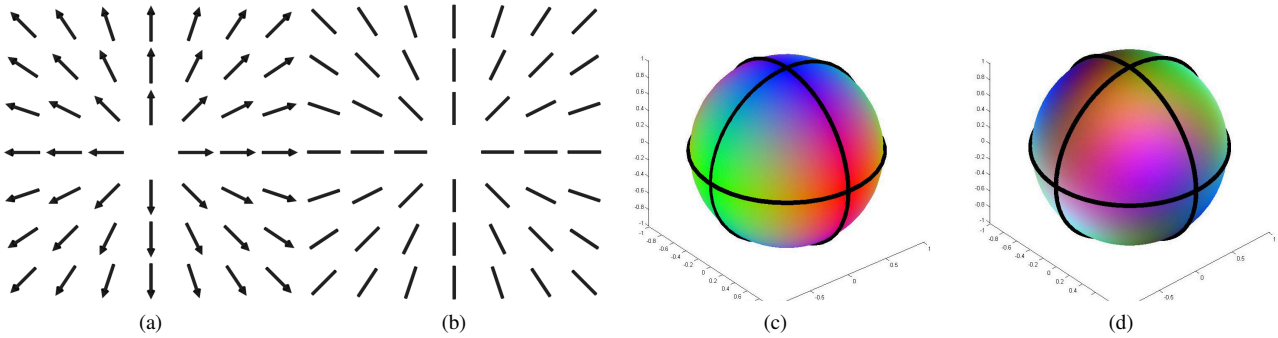

Figure 5.2: (a) A vector field. (b) A line field. Sphere colored using (c) absolute value method and (d) Boy's immersion. Three great circles aligned with the major axes are also shown. Note that in the case of the absolute value method colors are repeated eight times, once within each spherical triangle bounded by the great circles.

student Apery found another parametrization [7, 67]. Bryant and Kusner also proposed a smooth immersion which minimizes the Willmore energy of the surface (the least bent and most spherical) using complex arithmetic [29, 99]. Our approach is to represent a Boy-like immersion using low-degree polynomials restricted to the unit sphere, determined by treating Boy's immersion as an approximation to a polyhedral shape.

5.2 The Real Projective Plane

The real projective plane (RP^2) is the set of lines through the origin in $\mathbb{R}^3$; this set corresponds naturally to a surface. There are several ways to describe this surface; we show the most common ones in Figure 5.3.

The fundamental idea in all our descriptions is that each element of RP^2 — each line through the origin — intersects the upper hemisphere somewhere. For most lines, the intersection point is unique. But for

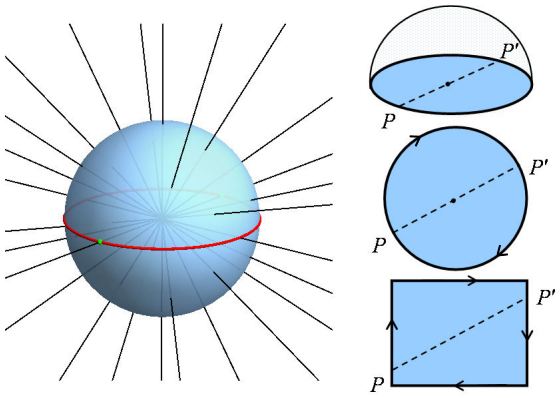

Figure 5.3: Models of RP^2. Top: The upper hemisphere, with antipodal points on the equator identified. Middle: a unit disk, with antipodal boundary points identified. Bottom: A rectangle, with opposite edges identified, with a reversal in orientation.

horizontal lines, there are two intersections, at antipodal points on the equator. This establishes a correspondence between RP^2 and a set whose elements are (a) points in the interior of the upper hemisphere and (b) pairs of antipodal points on the equator. If we imagine "gluing together" antipodal points (just as we "glue together" the two copies of the international dateline in wrapping a mercator projection onto a sphere), we get a representation of RP^2; unfortunately, it's hard to imagine doing this in three dimensions.

The upper hemisphere is topologically the same as a disk, so we can instead consider RP^2 as a disk with antipodal boundary points identified. And a disk is topologically equivalent to a rectangle, so we can consider RP^2 as a rectangle with opposite sides identified. All three of these approaches are shown schematically in Figure 5.3.

5.2.1 Realization in 3D and Boy's Surface

Because the real projective plane is a non-orientable surface (i.e., it contains a Möbius stript) and is closed (has no boundary), it admits no embedding in $\mathbb{R}^3$. This implies that there can be no one-to-one and smooth coloring of 3D line fields.

Nonetheless, it is still possible to *immerse* RP^2 in 3-space in a relatively simple way. A smooth mapping $f : M \to \mathbb{R}^n$ is an *immersion* of the surface M if $Df(p)$ the derivative at the point p, has rank two at each point $p \in M$, and if for every point $p \in M$, there is a neighborhood U_P of P such that f, restricted to U_P, is one-to-one; intuitively speaking, an immersion is a mapping that looks like an embedding locally, but which may have self-intersections. The first immersion of RP^2 in 3-space was discovered by the German mathematician Werner Boy using hand-drawn figures in his 1901 thesis [24]; this immersion is known as *Boy's surface*. It is this three-fold symmetric immersion of RP^2 that we use for coloring line fields.

5.2.2 Parametrization

Boy's immersion was described visually; we generated a smooth mapping from RP^2 to $\mathbb{R}^3$ whose image has the general shape indicated by Boy. We first generated a polyhedral shape (the *boxy approximation*) that approximated Boy's immersion, in much the way that a cube could be said to approximate a sphere (see Figure 5.4). We then triangulated the unit sphere with an antipodally-symmetric triangulation (i.e., if P is a vertex, so is $-P$; if there is an edge from P to Q, there is an edge from $-P$ to $-Q$, etc.) so that each antipodal pair of vertices in the triangulation corresponded to a single point in the boxy approximation, and each antipodal pair of triangles on the sphere corresponded to a single triangle in the boxy approximation of Boy's surface. This defines a mapping f from the sphere to the boxy approximation: each point $\pm(x,y,z)$ in a spherical triangle was represented by barycentric coordinates with respect to its spherical triangle; it was mapped to a point in the corresponding triangle in the boxy approximation with the same barycentric coordinates. We write $f(x,y,z) = (f_1(x,y,z), f_2(x,y,z), f_3(x,y,z))$.

The continuous function $f : S^2 \to \mathbb{R}^3$ could be described as a "piecewise linear version of Boy's immersion." Each of the three component functions f_i of f is a continuous real-valued function on the sphere. Every such function can be written as an infinite weighted sum of spherical harmonics, using a method completely analogous to Fourier analysis, which expresses a periodic function on the real line as a weighted sum of sines

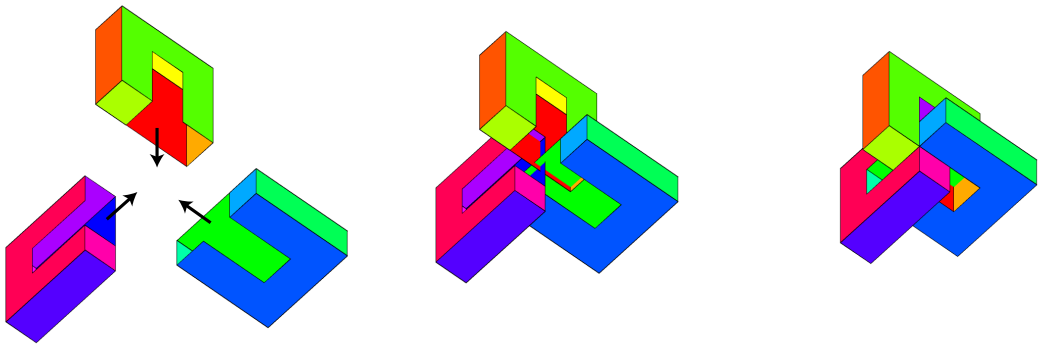

(a) (b) (c)

Figure 5.4: Boxy approximation of Boy's surface, shown here as a union of three congruent pieces in (a) partially exploded, (b) intermediate, and (c) assembled views.

and cosines [30].

The first component function, f_1, therefore can be written in the form

$$f_1(x,y,z) = \sum_{i=0}^{\infty} c_i h_i(x,y,z)$$

where the h_i are the spherical harmonics, and the c_i are real-valued weights. Because of the antipodal symmetry ($f_1(x,y,z) = f_1(-x,-y,-z)$ for all x,y,z), only the spherical harmonics of even total degree (when expressed as polynomials in $x,y,$ and z) can have nonzero coefficients in this expansion.

As defined, f_1 is continuous but not smooth. To smooth it, we can take a *partial* sum:

$$f_1(x,y,z) \approx \sum_{i=0}^{n} c_i h_i(x,y,z)$$

of just n terms of its "Fourier series". Since each spherical harmonic is smooth, the resulting function is a smooth approximation of f_1. We can do the same with f_2 and f_3, of course.

This is similar to Fourier-analyzing a square-wave s (which is discontinuous) as

$$s(x) = sin(x) + \frac{1}{3}sin(3x) + \frac{1}{5}sin(5x) + \ldots$$

and then noting that the first term of the series gives a smooth approximation of the discontinuous square wave. If we sum the first several terms, to improve the approximation, we get "ringing artifacts" in which the approximation "overshoots" the square wave near discontinuities. If we instead sum the first n terms, but multiply the ith coefficient by $\frac{n-i}{n}$, this artifact is substantially reduced; this is called "smoothing with the Fejer kernel."

We did the same thing with Boy's surface: we ordered the spherical harmonics by degree (each spherical harmonic can be written as a polynomial in the coordinates $x,y,$ and z on S^2), sum up the low-degree terms in the Fourier series, attenuating by degree; the result is a smooth approximation of the boxy approximation; the smooth approximation is expressed as three fourth-degree polynomials in the x,y,z coordinates on the unit

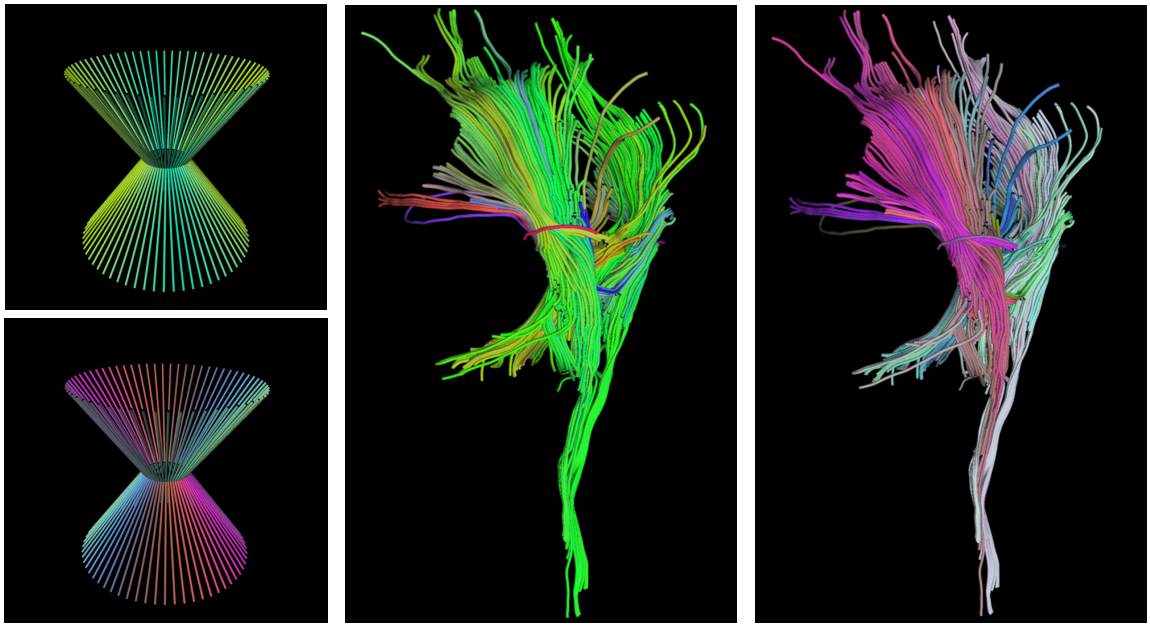

(a) (b) (c)

Figure 5.5: Tracts colored by encoding the lines determined by the end-points of the tracts (also known as the *end-point vector*, the direction of which does not matter); a common way coloring fiber tracts in DTI. (a) Synthetic tracts colored using (top) the absolute value method and (bottom) Boy's surface immersion in RGB; tracts designed to demonstrate the symmetries of the absolute value method. The internal capsule, extracted from a normal person's DTI brain data set, colored by encoding the end-point vector with (b) the absolute value method, and (c) Boy's surface. While Boy's RP^2 immersion assigns unique colors to unique lines, the absolute value method assigns the same color to the mirror-symmetric lines.

sphere. As a final step, we rounded the coefficients to two decimal places, and then adjusted some coefficients by hand to achieve an aesthetically pleasing shape, because our original goal (25 years ago) in creating this polynomial immersion was aesthetically motivated. The entire process of finding this polynomial immersion was ad hoc: the boxy approximation was easy to create; the triangulation of the boxy faces was done to maximize minimum angles; the spherical triangulation was hand-adjusted to have "nicely shaped" triangles whose areas were similar to those of the corresponding triangles in the boxy approximation.

5.3 Coloring

Given the parametrization above (or any other parametrization of Boy's surface), coloring is straightforward. To color an element of a line field (which is determined, up to sign, by a unit vector $\mathbf{v} = (x,y,z)$), we consider the element $[x,y,z] \in RP^2$; to this, we apply Boy's immersion f, producing an element $f([x,y,z]) \in \mathbb{R}^3$, which we in turn treat as a color-triple. We implemented our coloring based on Boy's immersion in Matlab and make it available online at [136].

5.3.1 Color Space of Immersion

We immerse Boy's surface both in the RGB and L*a*b* color spaces. L*a*b* is a perceptually uniform color space, where the perceptual difference between any two colors is approximately equal to the Euclidean distance between them in the color space. For example, if there is an isometric map from data points to a Euclidean subspace, the L*a*b* would provide a coloring reflecting "distances" between the data points as perceptual differences. Since Boy's immersion is smooth and one-to-one locally, the L*a*b* is a natural choice. However, the gamut of the L*a*b* color space has an irregular shape, and saturated colors lie close to the gamut boundary. This consequently limits the utilization of the saturated colors. On the other hand, the RGB has a rectilinear shape allowing simpler use of the color volume. Note also that, in both color spaces, the level sets of luminance are two-dimensional planes; Therefore, it is not possible to create an iso-luminant color mapping using a non-planar immersion like Boy's.

5.4 Results

We applied our model to sectional visualizations of DTI brain volumes and to fiber tracts extracted from these volumes by tracking principal eigenvectors of the underlying tensor-field. For comparison, we also obtained coloring results for the standard absolute value method. The DTI datasets were gathered by scanning two healthy subjects. Results, shown in Figures 5.5-5.9, demonstrate that coloring based on Boy's surface assigns unique colors to unique tissue orientations in a smooth, continuous manner, while the absolute value method does not.

5.5 Discussion

Note that Boy's surface has a curve of self-intersections that winds around and ultimately intersects itself at a triple point. This means that certain pairs of distinct directions will be coded by the same color. But, unlike the absolute value coloring method, directions nearby these will be coded differently, so surrounding context should often help distinguish the directions. Also, it should be possible to obtain a solution of the double curve numerically and ε-perturb the points on it.

The power of our coloring scheme (or any other smooth, nearly-unique continuous mapping) shows itself not only in visualization but also in data inspection for verification (e.g., see Figure 5.9). RP^2 is a natural model for studying line fields and has been used in physics, particularly for describing singularities in nematic crystals topologically [111]. We expect it to find similar applications in the analysis and visualization of tensor fields in the future.

5.5.1 Comparing Boy's Surface with the Absolute Value Method

While uniqueness and smoothness are important criteria, they are not the only ones. For example, the absolute value method is an easy-to-read color model because it aligns the additive color primaries with the coordinate directions of the DTI datasets. Our parameterization of Boy's surface happens to send six points fairly near

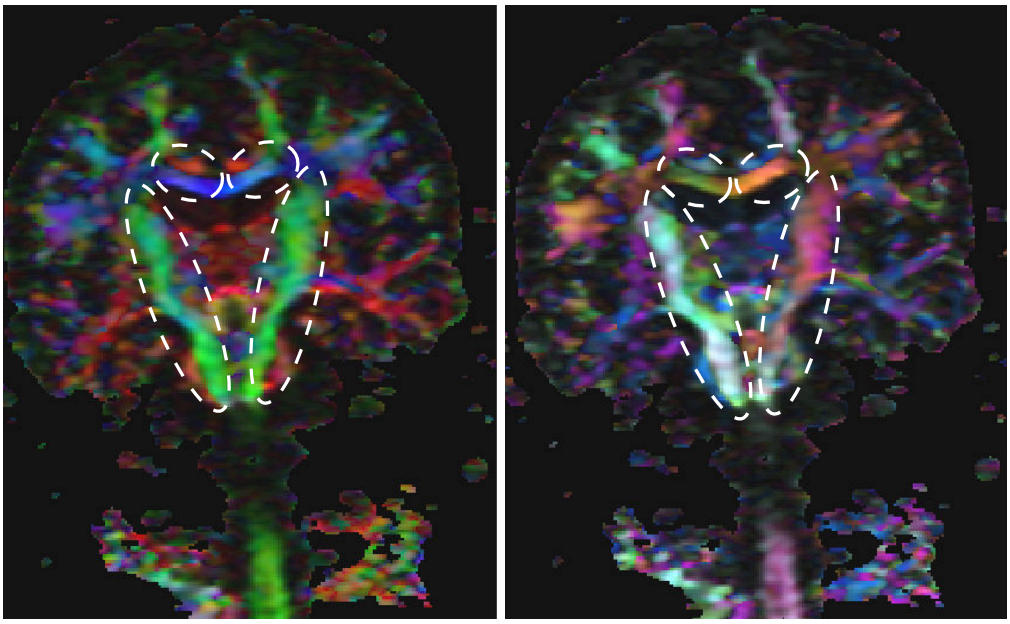

(a) (b)

Figure 5.6: Mid-coronal slice sagittal slice from a DTI brain data set of a normal subject; colored by mapping principal tensor axis at each point to an RGB color using (a) the absolute value method and (b) Boy's RP^2 immersion. Major anatomical structures selected, including the corpus callosum and internal capsule, show that while Boy's immersion assigns different colors to different tissue orientations, the absolute value method does not.

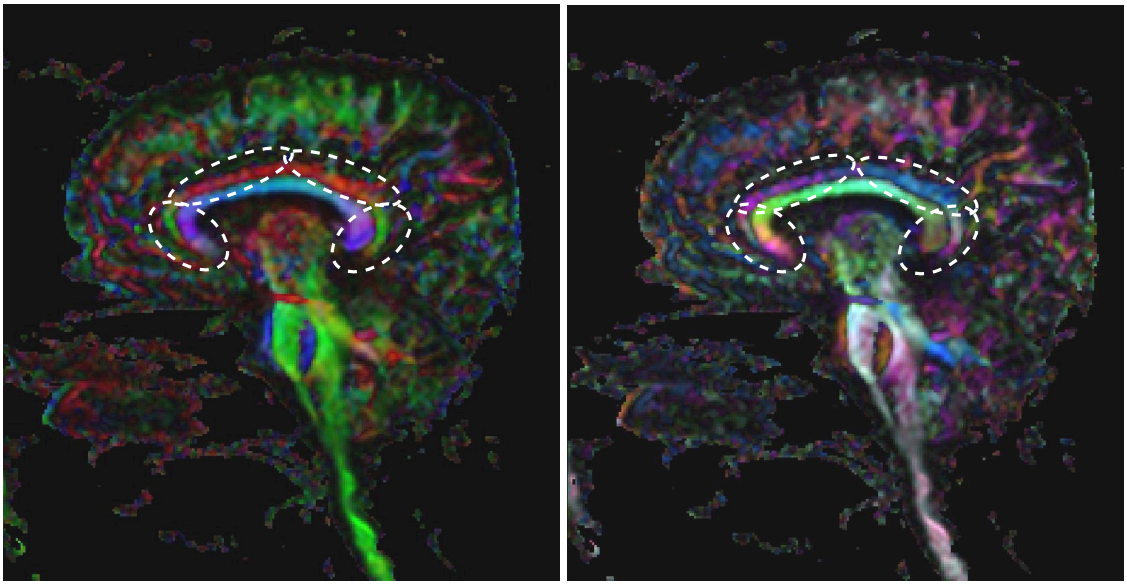

(a) (b)

Figure 5.7: Left-sagittal slice from a DTI brain data set of a normal subject; colored by mapping principal tensor axis at each point to an RGB color using (a) the absolute value method and (b) Boy's RP^2 immersion. Major anatomical structures selected include the cingulum bundle and corpus callosum.

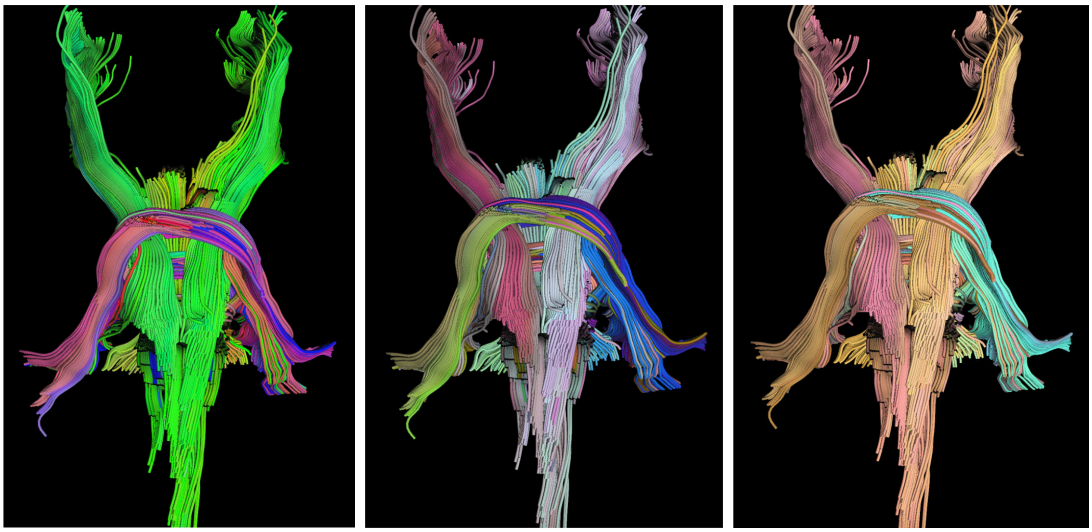

(a) (b) (c)

Figure 5.8: The internal capsule and the middle cerebellar peduncle (the tilted U-shape structure). Tracts colored by encoding the end-point vector (the direction of which does not matter) with (a) the absolute value method, (b) Boy's surface immersed in RGB, and (c) Boy's surface immersed in L*a*b*, an approximation for a perceptually uniform color space.

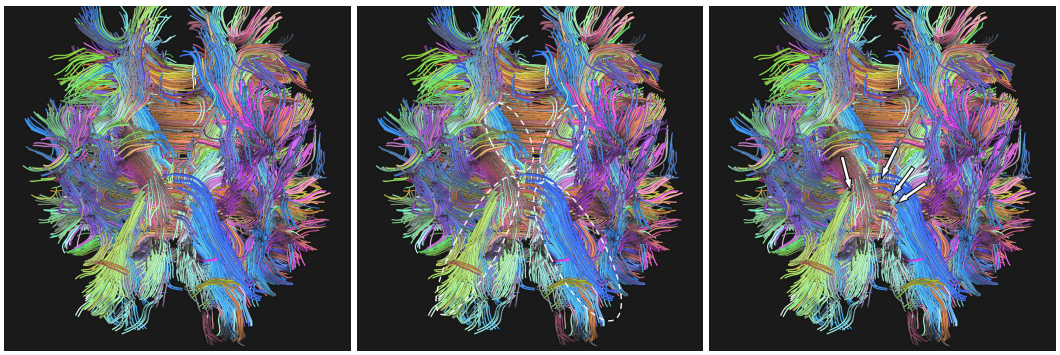

Figure 5.9: (a) Fiber tracts obtained from the whole DTI brain dataset of a healthy subject. Tracts colored encoding the end-point vector via immersing Boy's surface in RGB. (b) Unique coloring helps finding (c) broken curves (indicated with arrows).

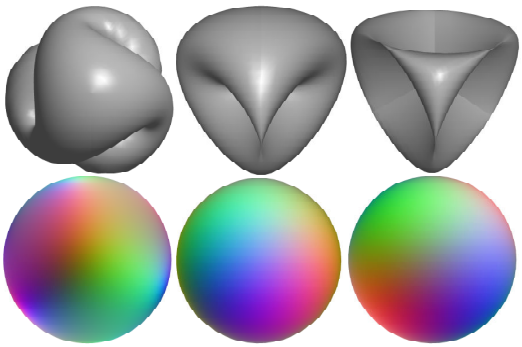

Figure 5.10: Sphere colored by immersing different realizations of RP^2 in RGB color space. (a) Bryant-Kusner parametrization of Boy's surface, (b) cross-cap, (c) Roman surface.

the equator (i.e., three line-pairs that pierce the sphere near the equator) to the tips of the three lobes of the immersion, while the north/south-pole pair is sent to the "center of the cap"; it is possible, by applying a rotation to the sphere, and adjusting the triangulation slightly, to make a very similar parameterization that sends the directions of the three major axes to the tips of the three lobes, which then can be placed on the three axes in color space, so that the Boy's-surface-based coloring shares the easy-to-read property with the absolute value method, at least for directions near the primary directions; we have not implemented this, however. We instead show coloring results for an immersion of the Roman surface transformed to maximize the alignment with major colors in the following section.

The absolute value method has a four-fold symmetry and, hence, can be useful for displaying structures with four-fold symmetry. One of the reasons that the absolute value method has been popular in visualizing DTI brain data sets is the fact that the brain has a two-fold (left-right) symmetry across the midsagittal plane at a certain scale. However, such symmetry does not exist in other directions, or even in the left-right direction within the hemispheres of the brain, not to mention other organs of interest such as the heart. Furthermore, it is always possible to turn a one-to-one color mapping into a symmetric one, given a reference of symmetry. Therefore it is a color model with the uniqueness property would typically be preferable to one without.

The absolute value method "squeezes" the color space into a spherical triangle with area one eighth of the sphere. The eight identically colored copies of the triangle are shown in Figure 5.2-c. Because the vertices of this triangle are fully saturated, all of the colors tend to be quite saturated, which can be desirable in certain application settings. This also hints of a trade-off between the color saturation and "bandwidth." Both of the methods can be improved, in the latter sense, by increasing the surface area within the color space by superimposing smooth wrinkles.

5.5.2 Other Realizations of RP^2 in $\mathbb{R}^3$

Boy's surface has other well-known parametrizations as we discussed above. A particularly elegant one is the Bryant-Kusner parametrization. This parametrization is the optimum mapping from RP^2 to $\mathbb{R}^3$ in the sense

that it minimizes the bending (or Willmore) energy, which is equivalent to locally maximizing sphericity. This feature of the Bryant-Kusner parametrization can help conveying dissimilarities more faithfully, particularly when a perceptually uniform color space is used. There are also other realizations of RP^2 that were known before Boy's surface, including the cross-cap and Roman surface, both of which are Steiner surfaces with singularities. We show the surface models of these realizations and coloring results obtained by immersing them in the RGB color space in Figure 5.10.

Which coloring is perceptually optimum entails further investigation, possibly including a user study. While both parametrizations of Boy's surface have nice properties such as non-singularity, smoothness, minimal distortion, and better coverage area, the Roman surface or cross-cap may be easier to align with the major axes. For example, by applying simple transformations (rigid and non-rigid) to the Roman surface, we obtained colorings approximately aligned with the major color axes (see Figure 5.11). As pointed in the previous section, it is possible to construct a similar alignment for Boy's surface as well, combining its desirable properties with ease of use.

5.6 Summary and Conclusions

Any coloring of a 3D line field is a mapping from RP^2 to a Euclidean subspace. We have introduced Boy's immersion of RP^2 as a method for such a coloring. Given that a smooth, one-to-one coloring for line fields is impossible in general, the coloring based on Boy's surface is a good approximation, because it is one-to-one and smooth except along the self-intersection curve of Boy's surface. While we have applied our model to visualization of DTI-derived data, it is likely to be useful in any domain where data instances are represented as lines, including edge, orientation, and rotation fields.

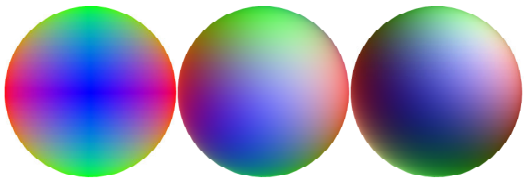

Figure 5.11: Sphere colored by (a) using the absolute value method, (b) immersing the Roman surface rotated to align with the major color axes in RGB, and (c) immersing the Roman surface transformed (non-rigid) to align with the major color axes in RGB.

Chapter 6

Similarity Coloring of Fiber Tracts

6.1 Introduction

Diffusion tensor magnetic resonance imaging (DTI) makes possible the exploration of fibrous tissues such as brain white matter and muscles non-invasively *in-vivo* [12]. It exploits the fact that water in these tissues diffuses at faster rates along fibers than orthogonal to them. Integral curves that represent fiber tracts by showing paths of fastest diffusion are among the most common data derived from DTI volumes [115]. The ability to estimate fiber tracts *in vivo* is in fact one of the key advantages of DTI over conventional imaging techniques. Integral curves are generated from DTI data by bidirectionally following the principal eigenvector of the underlying diffusion tensor field and are often visualized with streamlines or variations of streamlines in 3D. Reflecting the intricacy of the connectivity in the brain, these 3D models obtained from DTI brain data sets are generally visually dense, and it is often difficult to ascertain tract projections as well as anatomical and functional structures clearly. Therefore, the ability to see similarities and differences is fundamental to exploring tractography data. In this context, we present a coloring method that conveys spatial relations between tracts effectively: similar tracts are assigned to similar colors and different tracts are assigned to different colors in a smooth and continuous manner. In other words, variation in similarities (and dissimilarities) is represented with proportional variation of colors. To this end, we first compute pairwise similarities between the tracts then pose the problem of coloring tracts as an optimization problem and solve it.

Our contributions are twofold. First, we present a coloring method in which variation in similarities among tracts is reflected by variation in perceptual differences among their colors. Second, we introduce a geometric coloring model, the flat torus, that allows fiber-tract color differences to be adjusted arbitrarily. We also introduce a new fiber-tract similarity measure that uses the complete geometry of fiber pathways while giving higher "importance" to the endpoints. Note that, all or parts of this chapter have appeared in substantially similar form at [48, 52].

We show example visualizations of the corpus callosum and internal capsule (see Figure 6.1). The corpus callosum is the largest white-matter structure connecting the left and right hemispheres in the brain.

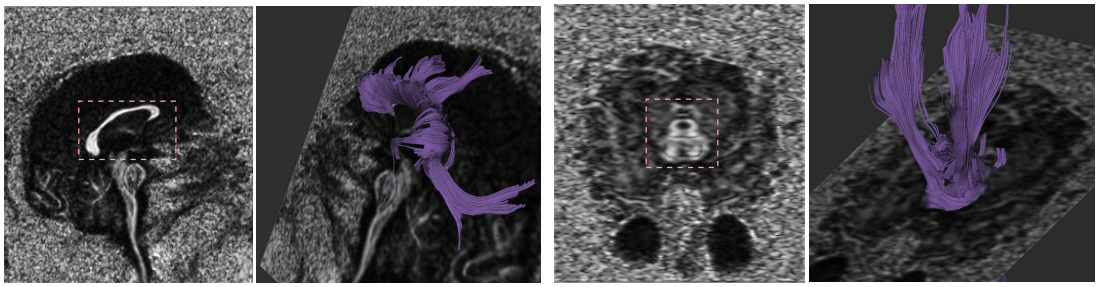

(a) (b)

Figure 6.1: Two brain fiber bundles superimposed with slices of fractional anisotropy (FA) map. (a) Corpus callosum with the mid-sagittal slice and (b) internal capsule with the mid-coronal slice.

The internal capsule runs between the cerebral cortex and the medulla, containing major longitudinal (both ascending and descending) pathway systems, including the corticospinal tract, medial lemniscus, and corticobulbar tract. Both the corpus callosum and internal capsule are targets for clinical and neuroscience research into normal developmental vs. pathological changes in white-matter integrity across the lifespan and the functional correlates of those changes.

In the following section we discuss related work. Then, we give details of our coloring method in Section 9.4 and present and discuss the results in Section 9.5. We finish the chapter by summarizing our work and contributions in Section 9.6.

6.2 Related Work

Mapping data values to colors is a fundamental operation in scientific visualization. Previous work based on empirical studies addressed the problem of generating perceptually effective colormaps [156, 80, 103, 96].

Several different geometric models, including lines, planes, cones, cylinders, and B-spline surfaces have been proposed for univariate, bivariate or trivariate color mapping [127, 134]. We extend the earlier models by introducing the flat torus model to give a continuous 2D color mapping that is approximately perceptually uniform and that can be repeated an arbitrary number of times in both directions to increase sensitivity.

Pajevic *et al.* proposed methods to colormap DTI cross-sections according to principal eigenvectors of tensor-valued voxels using different color spaces, including the perceptually uniform CIE L*u*v* color space [124]. The authors point at the potential limitations due to the irregularity of the L*u*v* space. Our flat torus model addresses the limitations due to the shape of the color space by allowing cyclic continuous bivariate mapping.

Integral curves generated from DTI volumes have been visualized generally with streamlines in 3D with different geometric (i.e., hyperstreamlines, streamtubes, etc.) and coloring combinations [97, 167].

In a work that is the closest to ours, Brun *et al.* colored DTI fiber tracts by embedding them in the RGB color space using a non-linear dimensionality-reduction technique [28]. The authors use end-point distances of tracts to define the similarity. Our work differs from this work in 1) the similarity measure that we use, 2) the embedding approach that we take, 3) the perceptual color space that we embed the curves in, and 4) the

bivariate coloring model that we introduce.

6.3 Methods

The goal of our coloring method is to represent the spatial variation among the fiber tracts as perceptual variation among their colors. To this end, we first compute a "distance" matrix quantifying similarities (or dissimilarities) between pairs of tracts. Then, we express the coloring as an optimization problem such that perceptual distances among colors of fiber tracts are proportional to similarities among fiber tracts quantified in the distance matrix. The problem as posed is a distance embedding problem and we approximate the solution using a spectral approach. We refine this global result using a simple mass-spring-based heuristic restricted to local neighborhoods. In the context of our work, embedding fiber tracts means finding a set of corresponding points in a Euclidean subspace where distances among the points optimally preserve the corresponding distances among the fiber tracts. We use our embedding procedure in two ways. In one, we embed tracts in the three-dimensional color space directly. In the other, we embed them into a plane and color via the flat-torus model by covering the flat-torus surface with the planar embedding.

6.3.1 Fiber Tracts as Integral Curves

Integral curves are models extracted from DTI volumes to approximate neural fiber tracts in the brain [1]. These integral curves are solutions to the following first-order differential equation: $\frac{dC}{ds} = \vec{v_1}(C(s))$, where s parametrizes the curve and v_1 is the principal eigenvector of the underlying diffusion tensor field at the point $C(s) = (x(s), y(s), z(s))$. We compute the integral curve $C(s)$ passing through a given seed point $C(0)$ by integrating the above equation for $s > 0$ and $s < 0$ (i.e., both directions from the seed point). We represent integral curves as lists of connected line segments (i.e., polygonal chains).

6.3.2 Similarities Between Fiber Tracts

We quantify how fiber tracts relate to each other by computing the anatomically motivated pairwise distance measure introduced in 4.3.2. Note that our measure does not necessarily satisfy the triangle inequality; therefore, it is not a metric. We compute distance between each pair of integral and assemble the measures to create a distance matrix. While we believe our distance measure is a good approximation of the notion of similarity in the domain, our coloring method is independent of a particular distance measure.

6.3.3 Embedding Fiber Tracts

In our coloring approach, we aim to reflect the boundaries in distance changes between integral curves as perceptual boundaries of colors. Intuitively, we are looking for a set of color coordinates x in a given color space that optimally preserves relations (as defined in D). Therefore, given a similarity matrix D, we define a

[1] We use the terms integral curve and fiber tract interchangeably throughout this chapter

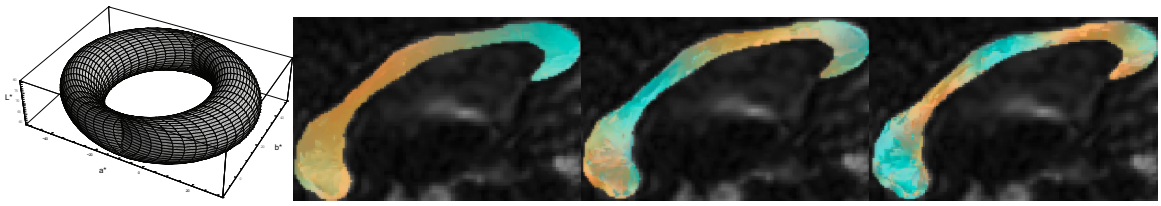

(a) (b) (c) (d)

Figure 6.2: (a) Bohemian Dome. (b-d) Flat torus model adjusting the sensitivity of color mapping by rescaling the planar representation (embedding) of the integral curves and "wrapping" around the flat torus as many times as needed. (b-d) Visualization of the mid-sagittal plane of the corpus callosum with increasing sensitivity of coloring.

cost function $E(x)$ and minimize it.

$$E(x) = \sum_{i<j} (||x_i - x_j|| - D_{ij})$$

Searching for x that minimizes $E(x)$ is indeed a distance embedding problem and we take a spectral approach to approximate it as suggested in [38]. The reason for choosing this particular spectral embedding method is its simplicity and speed. We refine the spectral embedding result using a mass-spring-based heuristic restricted to an ε neighborhood [59].

Perceptually uniform color spaces such as the L*a*b* are natural choices for our method as perceptual differences between colors are encoded as Euclidean distances in these color models. A color space is said to be perceptually uniform if the perceptual difference between any two colors in just-noticeable-difference (JND) units is equal to the Euclidean distance in that color space [65].

We use the embedding procedure above in two different ways to color fiber tracts. In the first, we embed the fiber tracts into the three dimensional L*a*b* color space [2] directly. In the second approach, we use the flat torus as a bivariate coloring model. For this, we first create an embedding in the plane and then map this planar embedding to cover the flat-torus surface arbitrarily many times to adjust the sensitivity of the coloring. Then, we submerse the flat torus in the color space to color the curves.

Flat torus

We propose the flat torus as a new geometric model for bivariate color mapping. A flat torus in 4-space is a Cartesian product of two circles in R^2. It can be obtained by a mapping $W : R^2 \rightarrow R^4$ such that

$$W(x,y) = (u,v,s,t) = (r_1 \cos x, r_1 \sin x, r_2 \cos y, r_2 \sin y) \tag{6.1}$$

where r_1 and r_2 are the radii of the circles. The flat torus has zero Gaussian curvature everywhere; therefore, a plane can be wrapped around it without distortion [56]. This can be particularly useful if all four coordinate values are used in a perceptually meaningful way, because wrapping the embedding plane onto the flat torus does not distort the embedded distances. One of the primary advantages of using flat torus is that we can

[2] L*, a*, and b* are the three orthogonal axes of the color space, where L* determines the lightness of the color (ranging from 0 to 100) while a* (ranging green to red) and b* (ranging from blue to yellow) together determine the chroma and hue of the color. See [65] for further details.

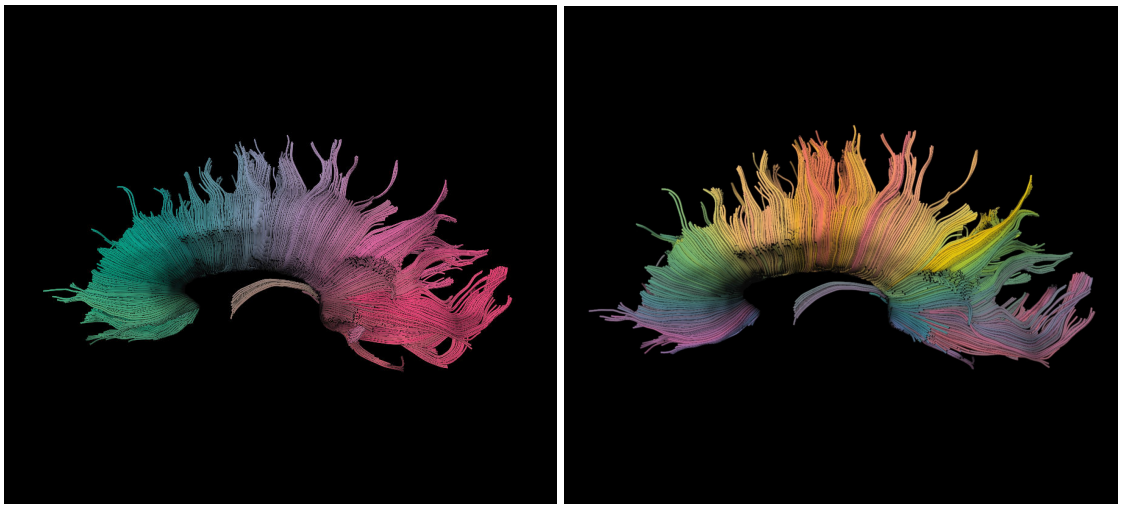

(a) (b)

Figure 6.3: Corpus callosum colored by (a) embedding the tracts in the L*a*b* color space directly and (b) submersing the flat torus, covered with a planar embedding of the fiber tracts, in the L*a*b* color space. Notice non-local color repetitions in the latter case.

adjust the sensitivity of the color mapping by rescaling the data values (i.e., points on the plane) uniformly and wrapping around the two circles, determined by r_1 and r_2, continuously an arbitrary number of times. This in turn provides a more flexible use of the color bandwidth. Our second coloring approach wraps the planar representation generated by the embedding algorithm onto a flat torus creating 4D (u,v,s,t) coordinates for the plane points. Then we project the flat torus onto a quartic surface called the Bohemian Dome (see Figure 6.2) centered at $(L_o, a_o + r_1, b_o)$ in a visible portion of the L*a*b* color space as follows:

$$(L*, a*, b*) = (L_o + t, a_o + r_1 + u + s, b_o + v)$$

The resulting image shows larger changes in color where neighboring integral curves differ more. Note that this projection is not isometric. It has two lines of self-intersection (where different (x,y) points map to the same colors); it also distorts the angles between the coordinate directions. For the examples shown in this chapter, we locate the projected Bohemian Dome in the L*a*b* interval $I = (I_L, I_b, I_c)$, where $I_L = [45,95], I_a = [-50,70], I_b = [-20,70]$, and use $r_1 = 45$, $r_2 = 25$. Note that $L_0 = 70, a_0 = 10$, and $b_0 = 25$.

6.4 Results

We implemented our distance computation routines in C++ and the embedding procedure in Matlab. Source codes are available online at [142]. We set $\varepsilon = 4.0$ in the spring embedding refinement such that only distances within ε are iteratively refined. We apply our coloring procedure to the corpus callosum and internal capsule. Figure 6.3 shows the coloring of the corpus callosum via the flat torus model. Similarly, Figure 6.4 shows the coloring of the internal capsule using direct embedding into the L*a*b* space.

A common way of coloring fiber tracts in the DTI community is to color-encode the lines determined

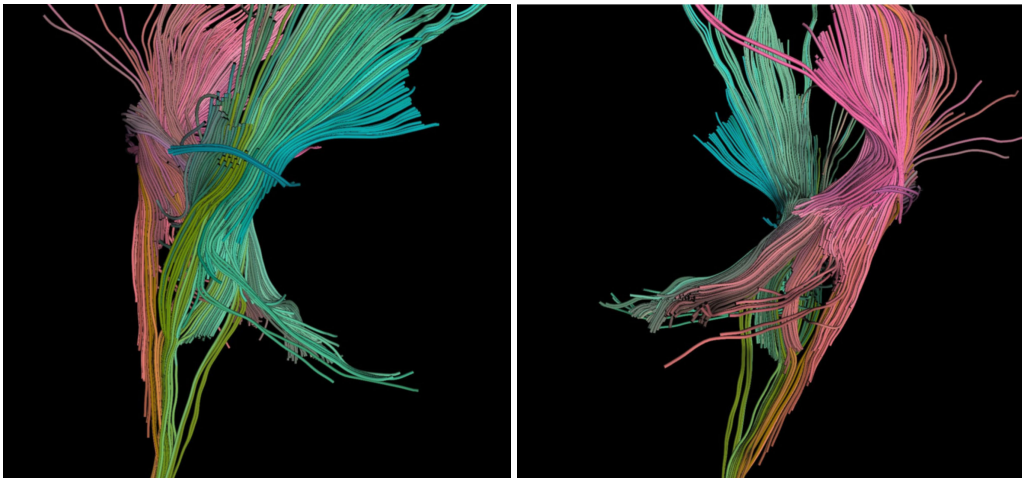

Figure 6.4: Two views of the internal capsule colored by embedding the tracts in the L*a*b* space directly.

by the end-points of the tracts (also known as the *end-point vector*, the direction of which does not matter) by mapping the absolute value of the normalized end-point vector to RGB color-triples [128]. Figure 6.5 demonstrates the advantage of our similarity coloring method over this de-facto standard. The end-point vector coloring gives almost uniform colors (due to mirror symmetries resulted by taking the absolute value of the vector) for both the internal capsule and corpus callosum models, while our method clearly represents the spatial variation among tracts as variation among their colors.

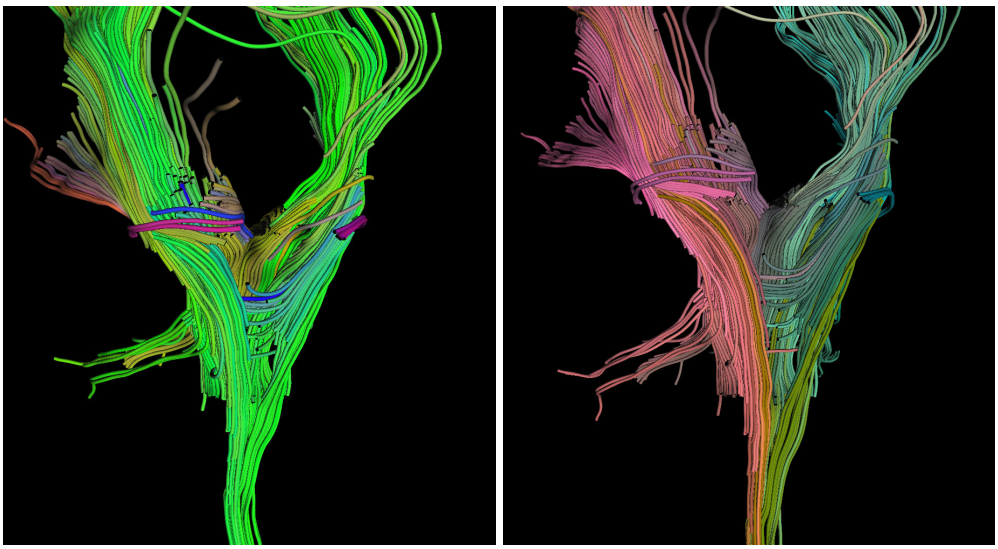

(a) (b)

Figure 6.5: Internal capsule colored (a) using the end-point vector coloring and (b) embedding the tracts in the L*a*b* space directly.

6.5 Discussion

It is important to note that the perceptual uniformity when the flat torus is used is an approximation, because the flat torus cannot be mapped to three space isometrically. Our projection can deemphasize changes in certain regions of the flat torus. There are other projections that may be closer to isometric, although our experimentation with linear and non-linear metric embedding methods such as metric multidimensional scaling (MDS) and isomap [149] using geodesic distances on the flat torus, which are projections of the lines in the lattice space, did not work as well as the Bohemian Dome projection. It also may be possible to add a fourth perceptual dimension like texture to the three color dimensions, removing the need for a projection and preserving the properties of the flat torus. It should be also noted that the perceptually uniformity of the L*a*b* color space is an empirical approximation obtained under certain display settings. Deviation from these settings (e.g., when monitors used are not calibrated) further distorts the approximation of perceptual uniformity [65].

Unless there is an isometry between the "manifold" of fiber tracts and the embedding space (e.g., R^3, R^2), it is not possible to preserve similarities among all pairs of fiber tracts in the embedding space. Our two-step embedding process aims to preserve similarities locally without distorting the global structure too much. However, as the number of fiber tracts increases while ε is kept fixed, it should be expected that both the computational cost and embedding distortion will increase, in general.

The two coloring approaches discussed here have their own merits. While coloring via the flat torus provides a cyclic use of the color space, direct embedding makes better use of the volume of the color space resulting in more saturated colors. Still, for both approaches, initial results indicate that our coloring method can help to quickly ascertain tract projections and find subtler anatomical divisions in fiber tracts.

6.6 Conclusions

We presented a coloring method that reflects spatial relations between neural fiber tracts effectively, assigning similar tracts to similar colors and different tracts to different colors, with differences preserved proportionately. We also introduced the flat-torus model that gives the ability to adjust the sensitivity of the color mapping by rescaling the data values (i.e., points on the plane) uniformly and wrapping around the two circles, making possible the continuous use of the color "bandwidth."

We applied our method to the corpus callosum and internal capsule in the brain. Feedback from neuroscientist collaborators suggests that our visualization methods can be useful in identification of smaller caliber anatomically or functionally related white-matter structures, particularly those that are contained within large bundles or fasciculi that project to multiple areas.

While our application interest has been in DTI, the coloring method described here can be applied to curves representing flow, surveillance trajectories, paths, etc. In fact, given a measure of similarity, our method can easily be used for coloring data entities arising from any field.

Part III

ANALYSIS

Chapter 7

A Slicing-Based Coherence Measure for Clusters of DTI Integral Curves

7.1 Introduction

In this chapter, we introduce a measure of coherence for a "hypothesized" cluster of DTI integral curves. As we discussed earlier, these curves correspond to neural tracts in the brain. The cluster coherence measure we propose relies on evaluating the stability of further subdivision of the cluster. To this end, we use the configurations at what we call "slices"– cross-sections of the cluster. Each slice is effectively an embedding of the curves into points in two-dimensional space. These points can be clustered using any off-the-shelf clustering algorithm. Each slice therefore provides a "vote" for each pair of curves being together or separate in the overall clustering. Furthermore, assuming reasonable smoothness of the curves, we can assess the temporal coherence of these votes: two adjacent slices carry more weight if they vote the same way than if their votes are opposite.

We demonstrate our measure's use to improve an agglomerative hierarchical clustering algorithm that has been shown to work relatively well in clustering integral curves corresponding to neural tracts [112]. When our slice-based method detects that a stable split exists in the cluster, it provides a specific partitioning that can be used as part of a clustering algorithm. Expert evaluation shows that this mechanism may be superior to the standard hierarchical clustering approach. While our primary motivation in designing the slicing-based coherence measure is to refine an initial clustering assignment, it can also be used for validating clusterings, quantifying connectivity or parametrizing clusters.

In the next section we discuss related work briefly, focusing on clustering methods. We then describe the details of the slice-based coherence measure in Section 7.3. Its application in improving hierarchical clustering results is described in Section 7.4, followed by discussion and conclusions in Section 7.5. Note that, this chapter was originally published in substantially similar form at [34].

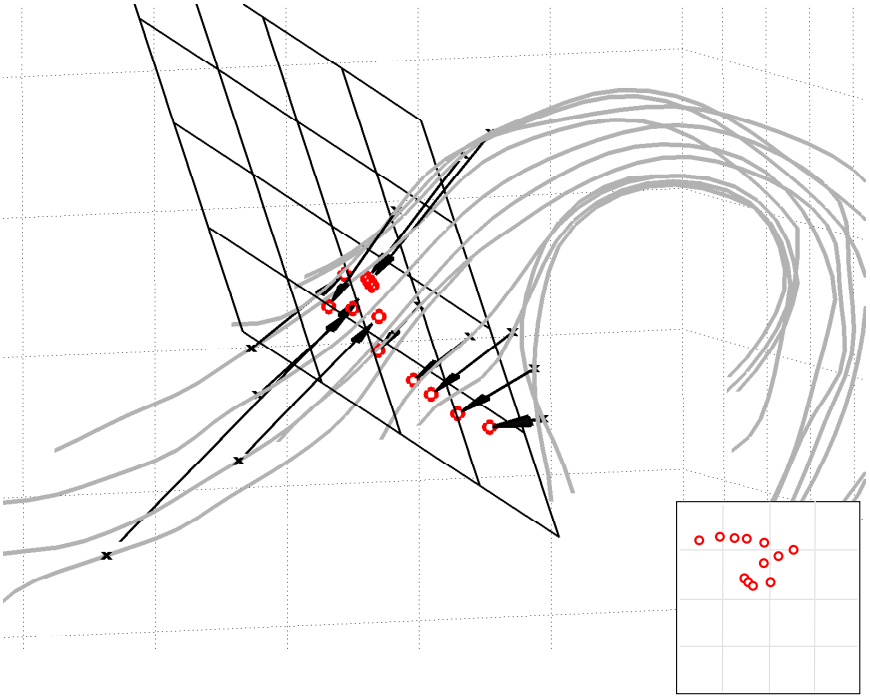

Figure 7.1: Illustration: the cluster is sliced using the arc length ratio $\alpha = 0.2$. Crosses (black): points on the curve corresponding to the arc length parameter αS_i. Circles (red): projections on the slicing plane. The bottom-right legend shows the embedding of the curves in the axis-aligned plane of the current slice.

7.2 Related Work

There have been several clustering methods proposed for DTI curves. All of them are adaptations of some of the well-known clustering methods including fuzzy c-means (a variation of k-means) [140], agglomerative hierarchical clustering [55, 40], and spectral clustering [27, 122]. An evaluation of the most popular fiber clustering of algorithms can be found in [112]. While we are not aware of any coherence measure specific to DTI integral curves (or 3D curves of any origin, for that matter), we found stability (or confidence) argument for cluster analysis in statistics to be common [25, 69, 100] and often used to compare and validate clustering methods. In this context, our notion of stability can be seen as an internal geometric index primarily aiming to quantify the quality of hypothesized individual clusters.

7.3 Slicing and Cluster Stability

The basic idea of our method is to assess the behavior of curves, that comprise a candidate cluster, relative to each other in a number of cross-sections (slices). A clustering pattern with more than one clusters, consistent over neighboring slices, is found, the cluster is considered incoherent, and split, with the splitting details deduced from the slices. Otherwise, the cluster is considered coherent, and preserved intact. Note that a clustering pattern in our case is a particular number of clusters and cluster membership. Therefore, consistency of a clustering pattern suggests consistency in number of clusters as well as consistency in cluster membership.

7.3.1 Slicing A Candidate Cluster

Given a cluster of undirected curves $X_1,\ldots,X_N$ in $\mathbb{R}^3$, we first orient the curves (in the sense of assigning start- and end-point designations), in a way that makes orientations consistent within the cluster. This is done by computing the start-to-end vector for each curve and iteratively re-orienting the curves until all the vectors are in the same half-space.

Next we sample each curve uniformly along its path in the following manner. Let S_i be the arc length of the curve X_i; i.e., the curve can be parametrized as $\mathbf{X}(s)$ where s goes from 0 to S_i (arc length parametrization). We take M samples $x_i^{(1)},\ldots,x_i^{(M)}$ from each curve such that $x_i^{(m)} = X_i(\alpha S_i)$ where $\alpha = \frac{m}{M}$. Note that while the arc length ratio α is the same at the mth sample for all curves, the arc length parameter αS_i will be different for each curve unless the curves have the same arc length. In other words, the arc length distance from the beginning of the curve to $x_i(m)$ will be different for every curve in general.

Now for each m, we have a point set $X^{(m)} = \{x_1^{(m)},\ldots x_N^{(m)}\}$. Each point set is then projected onto a two-dimensional slicing plane. Intuitively, we would like this plane to be normal to the "cluster tangent" at a given arc length. We estimate this by computing the tangent $\tau_i^{(m)}$ to the i-th curve at $x_i^{(m)}$ and averaging this direction over the curves, to yield $\tau^{(m)}$. The slicing plane is spanned by the columns of the 3×2 matrix $\mathbf{P}_m$ that are set to be orthonormal and orthogonal to $\tau^{(m)}$. Furthermore, we require the plane to pass through the mean location $\mu^{(m)} = \frac{1}{N}\sum_{i=1}^{N} x_i^{(m)}$. The geometry of this construction is illustrated in Figure 7.1.

This slicing mechanism is motivated by the following intuition. Suppose that the cluster is coherent, i.e. the curves comprising it follow similar paths. In that case, the slicing plane will be close to normal to the tangent of each individual curve, and moreover the points corresponding to the m-th sample will be close in space. On the other hand, if the cluster contains a number of distinct subgroups, there will be two sets of samples, which will be clustered around distinct sub-means. Moreover, in such a case we can expect that projecting those samples onto the slicing plane, orthogonal to the cluster, will emphasize the separation between the subgroups.

Given the m-th slicing plane, we represent the cluster by a set of two-dimensional projections, which we call the "slice", $Y^{(m)} = \{\mathbf{y}_1^{(m)},\ldots,\mathbf{y}_N^{(m)}\}$:

$$\mathbf{y}_i^{(m)} \;=\; \mathbf{P}_m^T(\mathbf{x}_i^{(m)} - \mu^{(m)}) \tag{7.1}$$

Before we proceed, we note that an alternative slicing technique could be based on fitting a plane to $X^{(m)}$ using Principal Component Analysis [92]. A straightforward application of this in our experiments has proved inferior to the tangent-based technique described above, due to the effect of outliers. This could be alleviated by using robust PCA; we do not, however, pursue it further here.

7.3.2 Cluster Co-membership Within Slices

We now treat each slice separately. Effectively, each slice is an embedding of the set of curves into two-dimensional space. The set of points in this space can in principle be clustered using any off-the-shelf clustering method. However, we are specifically interested in determining whether there is an "interesting" partition in the slice. Therefore, of most interest to us are methods that allow automatic determination of

the number of clusters. We use the Gaussian mixture clustering, accompanied by the Bayesian Information Criterion (BIC) for setting the number of components [138, 19].

7.3.3 Spatial and Temporal Coherence

Once a Gaussian mixture model has been fit to the slice $Y^{(m)}$ for a range of values of k, we select the optimal model based on BIC. With the mixture model, we cluster the data by assigning each point to the component with the highest responsibility (i.e., the highest posterior probability of the point drawn by the associated Gaussian distribution). We will denote the label of $\mathbf{y}_i^{(m)}$ by $c_i^{(m)}$.

There is of course no direct relationship between the cluster labels across slices, since those are assigned arbitrarily. Even if the same partition of the curves to two clusters is reached in two slices, a given set of curves could be labeled 1 in one slice and 2 in the other slice. The information of interest to us is conveyed by the co-membership of a given pair of curves in a particular slice. Specifically, we define $J_{ij}^{(m)}$ such that

$$J_{ij}^{(m)} = \begin{cases} 1 & \text{if } c_i^{(m)} = c_j^{(m)}, \\ 0 & \text{otherwise.} \end{cases} \tag{7.2}$$

Intuitively, the m-th slice votes for X_i and X_j being in the same cluster if $J_{ij}^{(m)} = 1$. One could now simply combine the values of this vote across all slices. However, it is possible for a particular pair of curves to be separated by the clustering in a given slice m simply due to the randomized nature of the procedure (placement of slicing plane, random initialization of the EM, etc.) despite genuinely belonging to the same subgroup. If indeed such an accidental result occurs, we expect that it will not persist in the neighboring slices, $m-1$ and $m+1$. This suggests the notion of temporal coherence. We formalize it with the value $T_{ij}^{(m)}$ defined as

$$T_{ij}^{(m)} = \begin{cases} 1 & \text{if } J_{ij}^{(m)} = J_{ij}^{(m-1)} \text{ and } J_{ij}^{(m)} = J_{ij}^{(m+1)}, \\ 0 & \text{otherwise.} \end{cases} \tag{7.3}$$

Finally, we can now represent the vote of the m-th slice regarding the similarity of curves i and j, weighted by the temporal coherence of that slice (with respect to that pair!):

$$W_{ij}^{(m)} = (1 - J_{ij}^{(m)}) \cdot T_{ij}^{(m)}. \tag{7.4}$$

$W_{ij}^{(m)}$ takes the value of 0 or 1. When it is 1, it indicates that the m-th slice supports a split where X_i and X_j are separated. $W_{ij}^{(m)}$ with 0 value, on the other hand, does not mean that the slice supports keeping X_i and X_j in the same cluster. It simply means that there is no evidence to the contrary. This may be due to the two curves being separated in the slice (zero co-membership), or the lack of temporal coherence, i.e., the co-membership of i and j being unstable in this slice–or both. This sort of asymmetric reasoning is similar in spirit to the statistical hypothesis testing formalism, in which the null hypothesis is either rejected or not, but never "accepted". In our case, this reflects the notion that not splitting the cluster is the default action.

We are now ready to describe the algorithm for evaluating the coherence of the cluster. Given the set of curves, we calculate a set of M slices, and cluster each of them using the Gaussian mixture clustering, with BIC model selection. We then compute the cluster co-membership values $J_{ij}^{(m)}$ and the temporal coherence

$T_{ij}^{(m)}$ for each $m = 1,\ldots,M$ and $i, j = 1,\ldots,N$. This yields for each pair i, j a set of M votes $W_{ij}^{(m)}$. We combine the evidence regarding X_i and X_j across all slices in a single value:

$$W_{ij} = M - \sum_{m=1}^{M} W_{ij}^{(m)}. \tag{7.5}$$

This provides us with a measure of similarity for each pair of curves, organized in the form of an $N \times N$ matrix $\mathbf{W}$. By default, its diagonal is set to zero. For example, a value of 0 indicates that the two curves are consistently separated in all slices; the value of M indicates that no evidence for separating the two curves is provided by any of the slices. It is important to note that our similarity values are computed in the context of the given cluster, specifically with the objective to evaluate potential splits. This is in contrast to "global" distance measures, that operate on the same scale throughout the data set.

7.4 Refining Hierarchical Clustering Results

Given initial hypothesized (candidate) clusters as a result of hierarchical clustering, we evaluate the decision of splitting the cluster by applying the spectral clustering algorithm described in [120] to $\mathbf{W}$. Briefly, the algorithm is based on eigendecomposition of the symmetric matrix $\mathbf{D}^{-1/2}\mathbf{W}\mathbf{D}^{-1/2}$, where $\mathbf{D}$ contains the sums of corresponding rows of $\mathbf{W}$ on its main diagonal and zeros everywhere else. When two clusters are requested, the algorithm divides the data according to the sign of the corresponding entries of the $N \times 1$ second largest eigenvector of the matrix above. Note that this can lead to either two clusters, or a single one (if all the entries in the eigenvector are of the same sign).

7.4.1 Expert Evaluation

In order to assess the utility of our method for a practitioner, we conducted a comparative evaluation study with a domain expert. The fiber tracking data used in our experiments were obtained from DTI brain data sets scanned from four volunteers. We first obtained initial candidate clusters by applying the single linkage hierarchical clustering algorithm, with a distance measure adapted from [167]. Note that the adapted distance measure does not prevents curves with radically different lengths to be in the same cluster. Let $\mathscr{D}$ represent the similarity matrix obtained using this distance measure. We used the cut-off threshold of 3 (set heuristically to produce reasonable cluster sizes). The expert has significant experience with DTI and uses fiber-track models (integral curves) generated from DTI data sets in clinical research regularly. We compared three methods. The first one is our slicing-based method described above; we used $M = 50$ slices for each candidate cluster. The second method applies the single-linkage hierarchical clustering algorithm on $\mathscr{D}$ for the candidate cluster, with the objective of obtaining two sub-clusters. In other words, this method finds the optimal split of the appropriate subtree in the original dendrogram. The third method applies the spectral clustering algorithm as does our method, but it uses the similarity matrix $\mathscr{D}$. We displayed 93 cases where at least one of the methods produced a different split to the expert. The expert was shown the results from the three methods side-by-side, and asked to rank them. The evaluation was blind (i.e., the expert was not told which of the methods produced each result). We used a streamtube representation for the curves and juxtaposed clusters with the

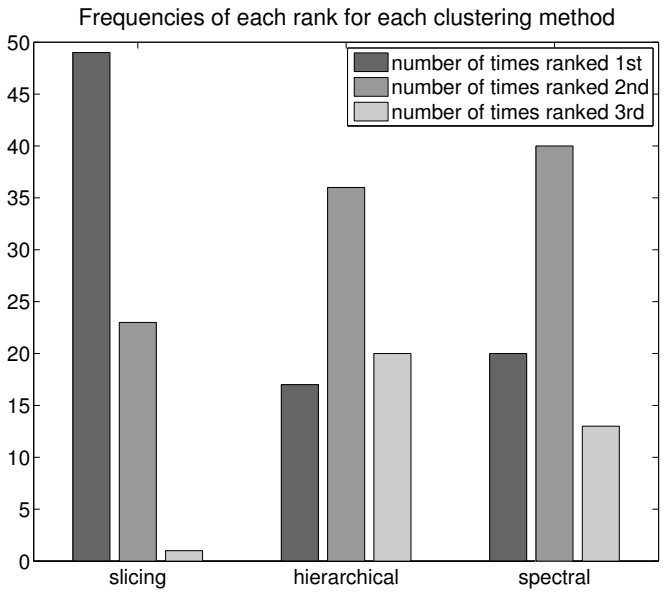

Figure 7.2: Ranking frequencies of each clustering method for 93 cases as evaluated by an expert.

surface of lateral ventricles (areas of cerebrospinal fluid in the brain) extracted to provide an anatomical landmark. The expert were able to interactively manipulate the viewpoint, zoom-in and out, and rotate the models. The ranking decisions were based on the following criteria: anatomical correctness (whenever the expert recognized a candidate cluster), anatomical and physiological plausibility, and amount of information conveyed by the resulting clustering. The goal in this scenario is to evaluate whether the resulting clustering decision helps to identify biologically distinct structures in the DTI data.

7.4.2 Results

Out of 93 cases considered, 20 were ranked as three-way ties (i.e., undecided, equally good, or equally arbitrary). In the remaining 73, the expert ranked our slicing-based method the best method in 49 cases while only in 9 of these cases there was another method ranked the same. Furthermore, our method ranked worse than the other two methods only in one case. We summarize these results in Figure 7.2. In Figure 7.3, we show examples where our method produced better results (according to the expert feedback) than the other two methods.

7.5 Discussion and Conclusions

The main contribution of the present work is the novel coherence measure obtained by combining an intuitive geometric idea, slicing, with known statistical machine learning techniques. An important property of our method is its reliance on the context of the cluster in evaluating similarity between curves. The quantities

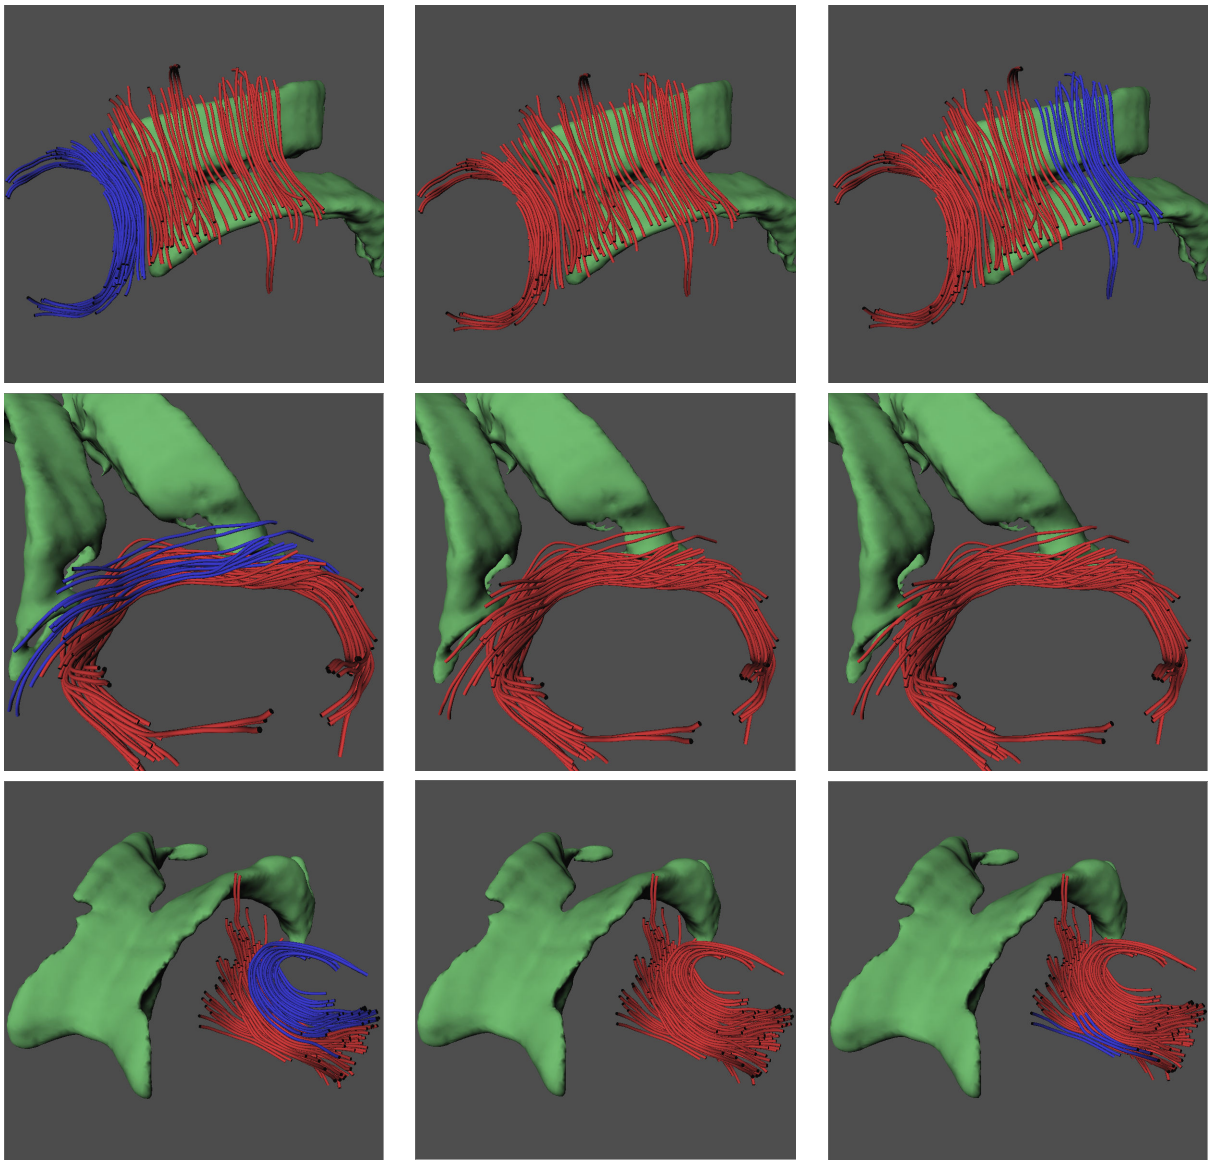

Figure 7.3: Examples of clusters where our method was ranked to be the best by the expert. Split, if any, is visualized by two colors (red and blue). The surface of lateral ventricles (green) provides an anatomical landmark. Left: the slicing-based split. It works well on clusters with curves having high curvatures as well as varying arc lengths. Middle: split based on the dendrogram used in hierarchical clustering. Right: split based on spectral clustering using the distance measure used to obtain the initial candidate clusters.

computed in each slice, and across slices, are tied directly to the clustering objective. This is in contrast to the more standard setup in which distances and thresholds are defined in a global context. Our method works reasonably well on clusters with curves having high curvatures and varying arc lengths. While our motivation and the experimental evaluation were on 3D DTI integral curves, our method generalizes to higher and lower dimensions easily and may apply to curve data in other domains. There are a number of technical aspects

that we believe could be improved. Specifically, a more robust slicing and projection method that explicitly down-weights outliers could help reduce uncertainty in the per-slice quantities. Also, while we are not aware of any "gold-standard" DTI fiber-track clustering data set, it is still possible to validate our method more quantitatively. The approach taken by Mobert *et al.* [112] can be a good starting point. A more challenging extension of the idea of slicing that we would like to pursue is to build a semi-parametric generative model for a cluster. Although our primary interest is in the analysis of DTI data, we believe our method is general and can be applied to any domain where data instances are represented by curves or trajectories.

Chapter 8

Cycles of Neural Tracts

8.1 Introduction

There are many local descriptors that use diffusion-derived measurements for quantifying structural variation in brain white matter (e.g., [144, 123, 74, 53]). Geometry augmented with statistics is a natural tool for creating local descriptors. On the other hand, one can propose a complementary global analysis. Topology offers a set of invariants that can help characterize a space globally, among them the number of independent cycles in the space (i.e., the Betti number). For example, the pth Betti number indicates the number of p-dimensional independent cycles. In essence, the Betti sequence, the sequence of Betti numbers with increasing dimension, is a proxy measure to capture the number of holes in the space. The computation of this invariant in a discrete setting amounts to a set of matrix reduction operations.

Therefore, we propose to compute the Betti numbers of neural tracts obtained using diffusion MRI so as to provide a global characterization of structural brain connectivity not available through existing, local biomarkers. This analysis has two main advantages: First, it captures the effects of axonal dropout, which creates "holes" in the white-matter space, caused by diffuse neurological disorders (Figure 8.1). Second, it provides a topological signature of the structural brain connectivity.

The branch of topology that studies cycles and invariants based on equivalence classes of cycles in topological spaces is called homology. Edelsbrunner et al. introduced the persistent homology framework to compute the Betti numbers robustly at all scales [60]. And the work proposed in this chapter is essentially an application of persistent homology to brain fiber pathways. In the following section, we briefly discuss recent applications of persistent homology. Section 8.3 provides a background on simplicial homology and related concepts. Readers familiar with simplicial homology can skip to Section 8.3.4, where we detail the persistent homology computation for neural tracts. Section 8.5 reports results from computation of the proposed descriptor on brain datasets obtained from patients with vascular cognitive disorder (VCI) and tumor. We discuss our assumptions and their potential ramifications in Section 8.6. Section 8.7 recaps our contributions and provides conclusions.

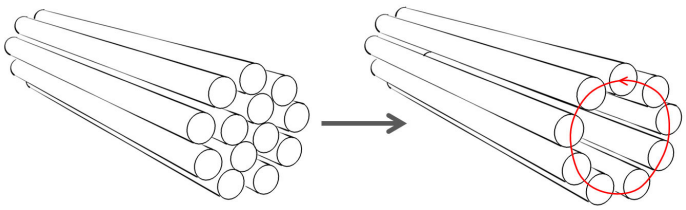

Figure 8.1: Tract dropout can create a hole in the tract space.

8.2 Related Work

Most of the recent applications of computational topology have been based on persistent homology. Persistent homology provides a framework for performing these computations in a multiscale fashion that is robust to noise [60]. Carlsson et al. use persistent homology to analyze the topology of the space of natural images, a collection of 3-by-3 image patches preprocessed to be restricted to a seven-dimensional sphere (S^7) [31, 101]. The authors show that the underlying "persistent" space of natural images has the topology of the Klein bottle. In a similar spirit, Singh et al. apply persistent homology to the study of population activity in the primary visual cortex (V1) [143] and show that the space of the activity patterns has the topology of a two-sphere (S^2). Persistent homology has been applied in analysis of biomedical data (e.g., [95, 121]) and recently to functional brain networks [102].

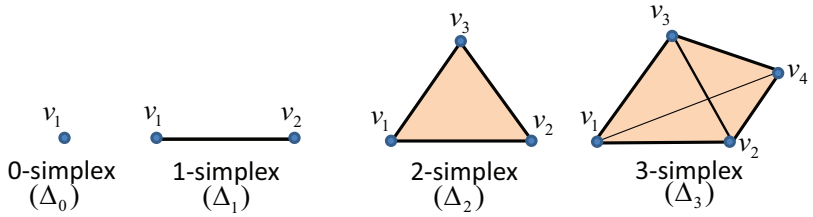

Figure 8.2: Examples of simplexes.

8.3 Background

Topology provides a formalism for qualitatively characterizing the structure in data [79]. Like geometry, topology is also concerned with the classification (or congruence) of the objects it studies, which requires a notion of invariance. In this context, a topological invariant is a property of the space that does not change under continuous stretching and bending of the space. As part of algebraic topology, homology theory studies cycles and invariants based on equivalence classes of cycles (homology classes) in topological spaces [72]. The basic invariant in homology theory is the Betti number. Specifically, the pth Betti number is the rank of the pth homology group formed by equivalence classes of p-dimensional cycles. Simplicial homology, homology over simplicial approximations of continuous spaces, provides a practical way to compute Betti

numbers by means of linear algebra. Note that while Betti numbers are topological invariants, they are not unique: two spaces that are not topologically equivalent can have the same Betti sequence.

8.3.1 Combinatorial Representation of Topological Spaces: Simplicial Complexes

There are different ways of representing topological spaces. We are interested in one that would let us use linear algebra in studying the topology of the space. One way is to use simple discrete convex building blocks that are convenient for computation (manual or otherwise). Typical examples of this approach would include the use of edges (line segments) to represent graphs, triangles to represent surfaces, or tetrahedra to represent solid structures. A *simplex* is a generalization of triangle-like convex building cells to higher dimensions. Specifically,

Simplex: A *p-simplex* is the convex hull of $p+1$ affinely independent points (i.e., p+1 points are not co-planar in the $p-1$ dimensional plane in $\mathbb{R}^n$). Formally, $\Delta_p = \text{conv}\{v_1, v_2, \ldots, v_{p+1}\}$ where $v_1, v_2, \ldots, v_{p+1}$ are affinely-independent points in $\mathbb{R}^n$. The dimension of the simplex is p.

To refer an unoriented simplex with its vertices, we use the notation $(v_1 v_2 \ldots v_{p+1})$, a concatenation of the vertices enclosed by parentheses. This notation indicates that the order of vertices listed does not matter. Conversely, we use $[v_1 v_2 \ldots v_{p+1}]$ to indicate an oriented simplex. Henceforth all simplexes referred to in this chapter are unoriented.

Face: The convex hull of any $k+1$ non-empty subset of $\{v_1, v_2, \ldots, v_{p+1}\}$ is a face of Δ_p, where $0 \leq k \leq p$. If the subset is *proper* (i.e., not the whole set), then the face is called a *proper* face. If Δ_k is a face of Δ_p we denote it by $\Delta_k < \Delta_p$. Note that a p-simplex has $\begin{pmatrix} p+1 \\ k+1 \end{pmatrix}$ k-faces (faces of dimension k) and a total of $2^{p+1}-1$ faces. See Figure 8.2 for examples.

Coface: If Δ_k is a face of Δ_p, then Δ_p is a coface of Δ_k denoted by $\Delta_p > \Delta_k$.

Boundary of a Simplex: The boundary of a geometric p-simplex is the union of all of its $(p-1)$-dimensional faces. This is the same as saying that the boundary of a p-simplex is the union of all its proper faces.

Simplicial Complex: A simplicial complex is a finite collection K of simplexes in $\mathbb{R}^n$, where

1. If $\Delta_p \in K$ and $\Delta_k < \Delta_p$, then $\Delta_k \in K$ (i.e., every face of a member simplex is also in the set K).
2. If $\Delta_p \in K$ and $\Delta_r \in K$, then $\Delta_p \cap \Delta_r = \emptyset$ or Δ_k with $\Delta_k < \Delta_p$ and $\Delta_k < \Delta_r$ (i.e., two simplices in K either do not intersect or intersect at a face of both).

See Figure 8.3 for examples.

The definition of a simplicial complex above is based on a realization in $\mathbb{R}^n$. However, sometimes we would like to work with simplicial complexes without worrying about their realizations.

Abstract Simplicial Complex: An abstract simplicial complex is a pair $X = (V, S)$, where V is a finite set of vertices and S is a set of non-empty subsets of V. Each element of S is a simplex of X. $s \in S$ is a p-simplex of the complex if s has $p+1$ elements. S should have the following properties:

1. If $v \in V$ then, $\{v\} \in S$
2. If $s \in S$, $t \subset s$, $t \neq \emptyset$ then $t \in S$

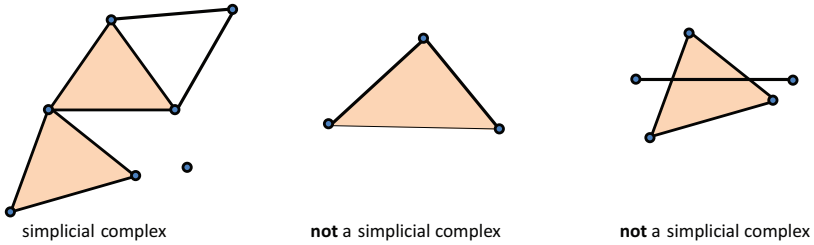

Figure 8.3: Example of a simplicial complex and two simplex collections that are not simplicial complexes.

Given a simplicial complex K, we can create the corresponding abstract simplicial complex X by simply setting V to be all the vertices of K and adding the vertex set $\{v_1, v_2, \ldots, v_{p+1}\}$ to S if and only if $\Delta_p = (v_1 v_2 \ldots v_{p+1}) \in K$. In fact, two simplicial complexes are considered the same (or isomorphic) if the abstract simplicial complexes they determine are isomorphic.

Chain Complexes: Let K be a simplicial complex and p be a dimension. A *p-chain* is a formal sum of p-simplexes in K: $c = \sum a_i \Delta_{\mathrm{p}}^{\mathrm{i}}$. The coefficients a_i, sometimes called multiplicities, can be reals, rationals, and integers, elements of a ring or field. Chains are added componentwise, like polynomials. Given two p-chains $c = \sum a_i \Delta_p^i$ and $c' = \sum b_i \Delta_p^i$ in K, we add them $c + c' = \sum (a_i + b_i)\, \Delta_p^i$ by adding their coefficients componentwise. As a convention, we omit simplexes with 0 coefficients and use Δ_p^i for Δ_p^i in writing p-chains.

The set of all p-chains in K is denoted by $C_p(K,G) = \left\{\sum_{i=1}^{n_p} a_i \Delta_p^i\right\}$, where n_p is the number of p-chains in K and $a_i \in G$. If G is a field such as $\mathbb{Z}_2$, $C_p(K, \mathbb{Z}_2)$ becomes a *vector space*. This is useful because it makes the existing machinery for vector spaces available to us. Due to their desirable computational properties, we often use mod-2 ($\mathbb{Z}_2$) coefficients. Unless otherwise stated, chain coefficients are in mod-2 throughout this chapter.

Boundary of a Simplex as a Formal Sum: We can express the boundary of a simplex as a formal sum:

$$\partial_p \Delta = \sum_{j=1}^{p+1} (v_1, \ldots, \hat{v}_j \ldots, v_{p+1})$$

where $\hat{}$ indicates that the vertex underneath is omitted from the sequence[1].

An important property of the boundary operator is that the boundary of the boundary of a simplex is always 0 as a formal sum:

$$\partial_{p-1} \circ \partial_p \Delta_p = 0$$

[1] The use of ∂ notation as the boundary operator is not arbitrary. If we think of the simplex representation $\delta = v_1 v_1 v_2 \cdots v_{p+1}$ as a symbolic multiplication of $p+1$ terms, then each term in the boundary chain of the simplex can be seen as the sum of the partial derivatives with respective to each vertex: $\partial \Delta = \sum_{i=1}^{p+1} \frac{\partial \Delta}{\partial v_i}$[5].

Proof:

$$\partial_{p-1} \circ \partial_p \Delta_p = \sum_{j=1}^{p+1} \partial_{p-1}(v_1, \ldots, \hat{v}_j, \ldots, v_{p+1}) \tag{8.1}$$

$$= \sum_{i \neq j} (v_1, \ldots, \hat{v}_i, \ldots, \hat{v}_j, \ldots, v_{p+1}) \tag{8.2}$$

$$= \sum_{i<j} (v_1, \ldots, \hat{v}_i, \ldots, \hat{v}_j, \ldots, v_{p+1}) + \sum_{i>j} (v_1, \ldots, \hat{v}_j, \ldots, \hat{v}_i, \ldots, v_{p+1}) \tag{8.3}$$

$$= 0 \tag{8.4}$$

In other words, every $(p-2)$ face of the simplex is shared by exactly two $(p-1)$ faces of the simplex and therefore the summation cancels in pairs to 0.

Cycle Chains: A *p-cycle* is a p-chain with an empty boundary ($\partial c = 0$). Note that 0 here means all the coefficients of the resulting chain are zero. Since the boundary operation commutes with addition, the set of all p-cycles in a simplicial complex K is a subspace (or subgroup) of $C_p(K,G)$ and denoted by

$$Z_p(K,G) \subset C_p(K,G)$$

Boundary Chains: A *p-boundary* is a p-cycle that is a boundary of a $(p+1)$-chain. The boundary of a chain is the formal sum of the boundaries of its simplexes. First we express the boundary of a simplex as a formal sum:

$$\partial_p \Delta = \sum_{j=1}^{p+1} (v_1, \ldots, \hat{v}_j \ldots, v_{p+1})$$

Therefore the boundary of a simplex is a formal sum of its $(p-1)$-faces. The boundary of a $(p+1)$-chain is then

$$\partial_{p+1}(c) = \sum a_i \partial \Delta^i_{p+1}$$

a p-chain that is a cycle Therefore, the set of p-boundaries in a simplicial complex K constitutes a subgroup (or subspace) of the cycle group $Z_p(K,\ G)$, denoted by

$$B_p(K,G) = \partial(C_{p+1}(K,G)) = \left\{ \sum_{i=1}^{n_{p+1}} a_i \partial \Delta^i_{p+1} \right\} \subset Z_p(K,G)$$

Notice that we have the nested group structure $B_p(K,G) \subset Z_p(K,G) \subset C_p(K,G)$.

Boundary Homomorphisms: Coefficient-wise addition (formal sum) of chain complexes can be seen as a first step in a linear algebraization of topology. Now, we introduce a type of map called the *boundary homomorphism* between chain complexes that is analogous to a linear map between vector spaces. Homomorphism is a general name for maps between algebraic objects that preserve the objects' algebraic structure. For example, if A and B are two groups with a binary operation $+$, like the chain groups discussed above, then a homomorphism $f : A \to B$ is a map such that $\forall a,b \in A$, $f(a+b) = f(a) + f(b)$. Hence "structure-preserving" here simply refers to preservation of the addition operation. Therefore, the boundary homomorphism $\partial = \partial_p : C_p \to C_{p-1}$ between chain groups is defined by

$$\partial_p \left(\sum a_i \Delta^i_p \right) = \sum a_i \partial_p (\Delta^i_p)$$

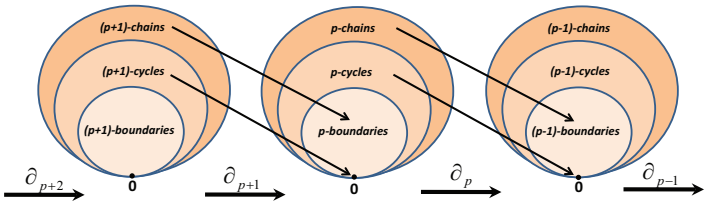

Figure 8.4: Boundary homomorphisms among chain groups (adapted from [60]).

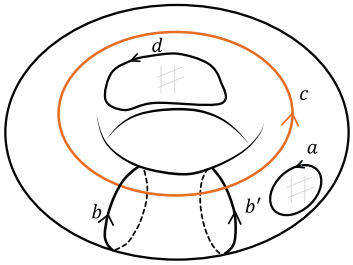

Figure 8.5: Examples of 1-dimensional cycles in a torus. The cycle a bounds a region, separating the torus in two parts. b does not bound a region, indicating the presence of a hole. b' does not bound a region either but goes around the same hole as b. Therefore b is "homologous" to b'. c does not bound a region and is not homologous to b, indicating another hole in the space.

The fundamental property of structure-preserving boundary maps (boundary homomorphisms) is that the boundary of the boundary of any chain is 0 (trivial):

$$\partial \circ \partial = \partial_{p-1} \circ \partial_p : C_p \to C_{p-2}$$

Proof:

We have shown that $\partial_{p-1} \circ \partial_p \Delta_p = \partial_{p-1}(\partial_p \Delta_p) = 0$.

$$\partial_{p-1} \circ \partial_p C_p = \partial_{p-1} \circ \partial_p \underbrace{\sum_{i=1}^{n_p} a_i \Delta_p^i}_{C_p} \tag{8.5}$$

$$= \sum_{i=1}^{n_p} a_i \underbrace{\partial_{p-1}(\partial_p \Delta_p^i)}_{0} \tag{8.6}$$

$$= 0 \tag{8.7}$$

The most important corollary of the above proposition is that every boundary is a cycle!

8.3.2 Homology Groups

The motivation for homology is to understand and quantify topological spaces with holes or voids (of any dimension) in them because the number of holes is a topological invariant. Since by definition voids are not there, we study theem indirectly by looking at cycles because, under certain conditions, the existence of a cycle indicates the existence of a hole. Two basic facts in this context should be understood early on. First, not all cycles indicate the presence of a hole. Specifically, only non-bounding cycles (i.e., cycles that are not boundaries) indicate the existence of a hole. Second, not all non-bounding cycles point to the existence of unique hole. In fact, most of them "go around" the same hole (Figure 8.5). What these two intuitive facts hint at is that in order to quantify the unique holes in the subject space, we need to find independent cycles that are not boundaries. In other words, since $B \subseteq Z$ (i.e., every boundary is a cycle), we are interested in the rank of the quotient space Z/B (the space of cycles that are not boundaries); Z_p/B_p is called the *pth homology group*.

As this chapter attests, preparing the technical machinery for computing homology groups on simplicial complexes is tedious but this effort pays off with the ability to directly use basic linear algebra techniques in the computation.

8.3.3 Boundary Matrices

The very first step in using linear algebra to compute homology groups is to use matrices for representing incidence relations between every pair of p and $p-1$-simplices in a given complex K. Formally, the *pth boundary matrix*, ∂_p, of a simplicial complex is a matrix in which rows correspond to the $(p-1)$-simplices and columns correspond to the p-simplices in K. Let n_p and n_{p-1} be the number of p- and $(p-1)$-simplices, respectively, in K. Then, ∂_p is an n_{p-1} by n_p binary matrix with entries $a_i^j \in \{0,1\}$, where the entry $a_i^j = 1$ if the ith $(p-1)$-simplex is a face of the jth p-simplex and is 0 otherwise.

$$\partial_p = \begin{bmatrix} a_1^1 & a_1^2 & \dots & a_1^{n_p} \\ a_2^1 & a_2^2 & \dots & a_2^{n_p} \\ \vdots & \vdots & \ddots & \vdots \\ a_{n_{p-1}}^1 & a_{n_{p-1}}^2 & \dots & a_{n_{p-1}}^{n_p} \end{bmatrix} \tag{8.8}$$

Understanding ∂_p is key to understanding homology computation. First, ∂_p uses an arbitrary but a *fixed* ordering of simplices; this ordering must stay fixed across boundary matrices in all dimensions. Second, we can express boundary homomorphisms as matrix multiplications using boundary matrices. For example, let $c_p = [b_1 b_2 \dots b_{n_p}]^T$ be a p-chain represented by a column vector of its coefficient entries. Note that the ordering of the p-simplices in this vector is the same as their ordering in the boundary matrix ∂_p. b_i is 0 if the ith simplex is absent from the chain. Now we can compute the boundary of the chain by a matrix multiplication: $\partial_p(c_p) = \partial_p c_p$. Third, Z_p, the pth cycle group of the complex, is by definition the null space of ∂_p. Recall that the null space of ∂_p consists of all solutions to $\partial_p x_p = \mathbf{0}$. We remarked above that most of these cycles are not independent in general. Therefore, a natural question to ask is what the rank of the null space is. Again, from linear algebra we know that if the linear system $\partial_p x_p = \mathbf{0}$ has r independent rows (or

equations) then it has $n_p - r$ independent solutions. In other words, the rank of the null space is the rank of the matrix subtracted from the column size of the matrix:

$$\text{rank}Z_p = n_p - \text{rank}\partial_p$$

The next question is what the rank of the boundary matrix means. By construction, the columns of a boundary matrix are boundaries of the corresponding p-simplices.[2] And the sum of any subset of these columns gives the boundary of the p-chain that would be obtained by adding the p-simplices corresponding to the columns. Therefore, the rank of the boundary matrix is, not surprisingly, the rank of $(p-1)$-boundaries:

$$\text{rank}\partial_p = \text{rank}B_{p-1}$$

If we rewrite the rank of the cycle group Z_p, we get

$$\begin{aligned}\text{rank}Z_p &= n_p - \text{rank}\partial_p \\ &= n_p - \text{rank}B_{p-1}\end{aligned}$$

8.3.4 Computing Homology

As observed above, homology studies non-bounding independent cycles in a topological space because they indicate the presence of unique holes in the space. And the number of holes as a topological invariant in turn helps determine whether two given spaces are topologically equivalent or not. We defined the pth homology group to be the pth cycle group modulo the pth boundary group: $H_p = Z_p/B_p$. The rank of the pth homology is what we have been looking for, the number of non-bounding independent p-dimensional cycles.

$$\begin{aligned}H_p &= Z_p/B_p \\ \text{rank}H_p &= \text{rank}Z_p - \text{rank}B_p \\ &= n_p - \text{rank}B_{p-1} - \text{rank}B_p \\ &= n_p - (\text{rank}B_{p-1} + \text{rank}B_p)\end{aligned}$$

$\text{rank}H_p$ is also called the pth *Betti number*: $\beta_p = \text{rank}H_p$. So, in order to find the pth Betti number, we need to gather the ranks from the pth and $(p+1)$th boundary matrices—recall that $\text{rank}\partial_p = \text{rank}B_{p-1}$. Similar to Gaussian elimination (while staying in the domain of the coefficients, which means mod-2 arithmetic in our case), using column and row operations, we can reduce a boundary matrix to a normal form called the *Smith normal form*. The reduced form of the matrix, $\bar{\partial}_p$, has an initial segment of nonzero diagonal elements

[2] For expository reasons we use column representations here. Recall that the row rank and column ranks are the same for matrices.

$d_1 d_2 \ldots d_r$, where d_{i+1} is an integer multiple of a_i, and 0s everywhere else:

$$\bar{\partial}_p = \begin{bmatrix} \overbrace{\begin{matrix} d_1 & 0 & 0 & \ldots \\ 0 & d_2 & 0 & \ldots \\ 0 & 0 & \ddots & \\ \vdots & & & d_r \\ & & & & 0 \\ & & & & & \ddots \\ 0 & & & \ldots & & & 0 \end{matrix}}^{\text{rank}B_{p-1}} & \overbrace{\begin{matrix} 0 \\ 0 \\ 0 \\ \vdots \\ \\ \\ \\ 0 \end{matrix}}^{\text{rank}Z_p} \end{bmatrix}$$

Note that in the mod-2 coefficient domain, the nonzero diagonal elements $d_1 d_2 \ldots d_r$ of the reduced matrix are all 1s. There are efficient algorithms for computing the Smith normal form of boundary matrices [60].

8.3.5 Persistence Homology

Persistence homology is an attempt to make simplicial homology computation robust to noise [60, 168]. The idea is to create multiscale nested representations of the space and perform the homology computation at each scale while keeping track of the independent cycles that appear and disappear (birth and death) from one scale to the next (for details). The hope is that recording the topological evolution of the space from fine to coarse scale will provide a robust and complete picture of the topology. In this context, the concept of *filtration* is important. A filtration of a simplical complex K is a nested subsequence of complexes $\emptyset = K^0 \subseteq K^1 \subseteq \ldots \subseteq K^m = K$. From K_i to K_{i+1} no simplex can disappear; only new simplices can be added. Basically, this is a construction of K by possibly adding new collections of simplices at each step i. It is a convenient and efficient way of representing as a single object nested sequences of simplicial complexes, each of which represents a different scale of the same space, as a single object Filtration as a mathematical object by itself, however, does not guarantee a multiscale representation; we want the multiscale representation constructed with help of a metric (or monotone function) to be a filtration (see Section 8.4.1).

8.4 Computing Homology Groups of White Matter

Diffusion models quantifying anisotropic diffusion locally make it possible to estimate neural pathways in the brain as a collection of space-curves, tractograms. Given a whole-brain tractogram, we compute the ranks of persistence homology groups of a given tractogram in three steps. First, we compute pairwise "distances" between the tracts in the tractogram, obtaining an affinity matrix. For this, we use the distance measure introduce in 4.3.2. Second, using the classical multidimensional scaling (MDS) on the affinity matrix, we generate a point-set representation of the tractogram in $\mathbb{R}^3$. Third, we construct alpha-complexes from the point-set while varying α between 0 and ∞, recording the births and deaths of cycles in all three dimensions (Figure 8.6).

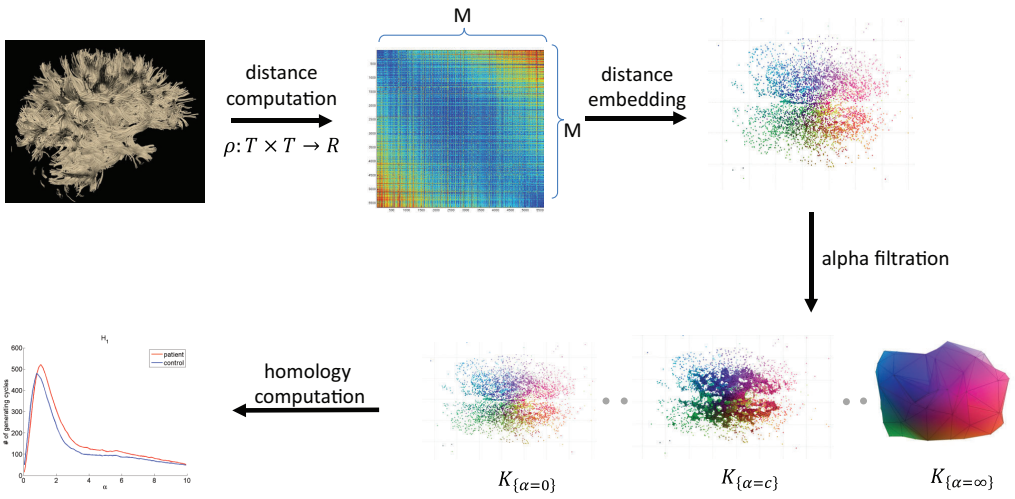

Figure 8.6: Overview of our method.

8.4.1 Alpha-Complex Filtration

We construct an alpha-complex filtration of the point-set P, while recording births and deaths of generating (independent) cycles in all three dimensions. The *alpha-shape* is a generalization of the convex hull of a point set P sampled from an underlying, often unknown space. Unlike the convex hull, alpha shapes can recover disconnected non-convex spaces with holes, of course, given the "correct" α parameter [61]. Specifically, for a given real parameter $\alpha \in [0,\infty)$, alpha balls are balls of radius α centered at the points in P. The alpha shape of P is then the union of the convex hulls of the points whose alpha balls intersect. Similarly, the *alpha complex* is the union of the Delaunay triangulations restricted to the same points [61, 33]. So, the alpha complex is a simplicial approximation of the alpha shape, as intended. When α is 0, the complex is just made of the point set P (i.e., all 0-simplices) and for sufficiently large α it is the Delaunay triangulation of P. In this context, the *alpha complex filtration* of this final Delaunay complex is the sequence of alpha complexes of P obtained by increasing α from 0 to ∞.

$$\emptyset = K^0 \subseteq K^1 \subseteq \ldots \subseteq K^m = \text{Delaunay complex}$$

Note that this is a filtration induced by the alpha parameter. Let α_i and α_j be the alpha parameters used to construct the alpha complexes K_i and K_j, respectively,k and let the simplexes $\sigma_i \in K_i$ and $\sigma_j \in K_j$. If σ_i is a face of σ_j, then $\alpha_i \leq \alpha_j$.

8.5 Results

We present results from two different group comparisons. For this, we compute and plot histograms of persistent homology intervals (i.e., birth-death intervals) averaged over subject groups. Our first result is from a comparison of control subjects with patients suffering from vascular cognitive disorder (VCI). VCI is a general term for vascular cognitive deficits caused by injuries to the white matter from cerebrovascular diseases.

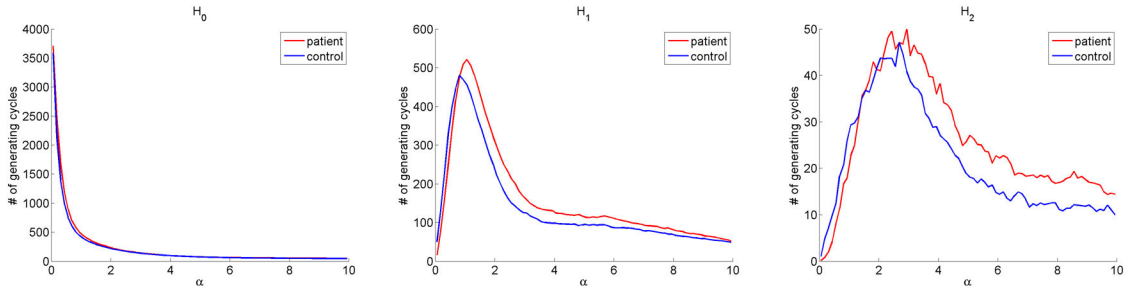

Figure 8.7: Evolution of the ranks of H_0, H_1, and H_2 for whole-brain tractograms of VCI patients and control subjects over varying α values. There is a persistent trend that the ranks of H_1 and H_2 for patients are greater than that of controls over a large α interval, confirming VCI's well-known global atrophic effects.

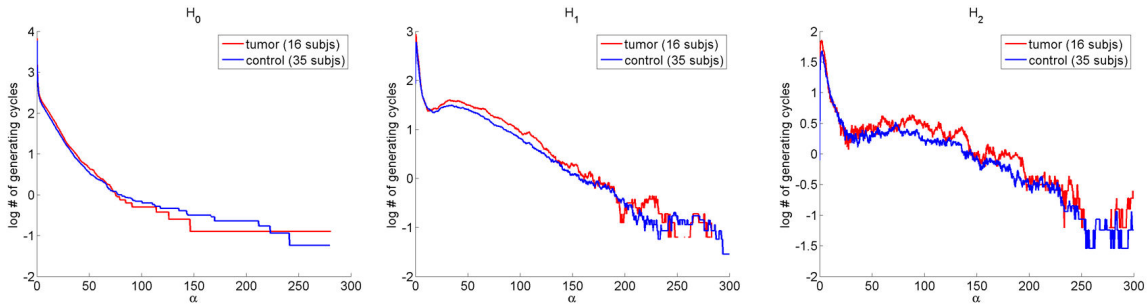

Figure 8.8: Evolution of the ranks of H_0, H_1, and H_2 for whole-brain tractograms of brain tumor patients and control subjects over varying α values. The ranks of H_1 and H_2 for patients are greater than those of controls over large α intervals. The differences, however, are not as strong as in the VCI-control comparison. The y-axis is in the log scale.

VCI typically manifests itself with problems in speech, language, and the ability to organize thoughts and track actions. Its effects on memory, however, are considered mild in comparison to Alzheimer's disease. Deficits in cognitive domains caused by VCI are often similar to those common in individuals of advanced age. Our subjects were 19 patients with VCI and 20 healthy individuals who served as a control group. The subjects were age-matched. Results indicate that the ranks of H_1 and H_2 for patients are greater than those of controls over large α intervals (Figure 8.7). Our second result is from a comparison of control subjects with glioma patients. Glioma is an aggressive, recurring, malignant type of brain tumor. Gliomas originate from glial cells, a type of support cells that fill the space between axons in the white matter; they grow by destroying nearby tissues, white and gray matter alike. Our subjects were 16 patients with tumor and 35 healthy volunteers used as a control group. Results show a trend similar to, but not as strong as, the results from the first case: The ranks of H_1 and H_2 for patients are greater than controls over large α intervals (Figure 8.8).

8.6 Discussion

We have presented preliminary results from our effort to quantify the global space of brain white matter using simplicial homology. We have throughout made a few implicit assumptions, each of which is likely to add a layer of approximation to the final characterization. First, we assumed that there is a *natural* metric that captures the anatomical space of white-matter tracts and that our similarity measure is a good approximation to this natural metric. However, to begin with, our measure is a semi-metric, not a real metric. The effects of using a semi-metric, instead of a full metric, on the final characterization need to be investigated further.

Second, for computational reasons, we created a point-set representation of white matter in $\mathbb{R}^3$ and then constructed a geometric simplicial complex from it. But this projection of course preserves the topology only approximately. We could, on the other hand, choose to create an abstract simplicial complex representation, which obviates the projection; however, this is not computationally feasible and prone to noise effects.

Finally, the distance measure used determines the manifold space of the white matter and its topology. Different metrics induce different topological spaces in general and one can thus talk about not only one but a family of topologies determined by a family of metrics used.

8.7 Conclusions

We have presented a global analysis of brain white matter based on persistent homology. This analysis has two main advantages: First, it captures the effects of axonal dropout, which creates "holes" in the white matter space caused by diffuse neurological disorders. Initial results indicate that the proposed analysis can capture global atrophic effects of neuropathologies in the brain. Second, the proposed analysis provides a topological signature of the structural brain connectivity. The latter can be useful in understanding and classification of the topological structure of the neural pathways in the brain.

Chapter 9

Tract-based Probability Densities of Diffusivity Measures in DT-MRI

9.1 Introduction

Tract-based approaches to analysis of diffusion brain data sets are attracting interest because they can geometrically localize quantities while reducing sensitivity of associated statistics to computational errors and noise (e.g., [144, 123, 74]). In this context, we propose using probability densities in characterizing diffusion along diffusion-tensor MRI (DTI) fiber tracts generated from diffusion-weighted imaging data sets. We experiment with well known diffusivity measures such as fractional anisotropy (FA), mean diffusivity (MD), radial diffusivity (RD), and axial diffusivity (AD) [145]. We estimate univariate and bivariate diffusion densities of these four measures nonparametrically, using histogram and kernel density estimation methods. We use the estimated densities for hypothesis testing of differences between a group of patients with vascular cognitive impairment (VCI) and a control group. Results indicate that probability density estimates can detect group differences with high statistical significance.

Our main contribution is to introduce tract-based probability density functions, including the joint density of tract arc length and scalar diffusivity measures, as potential metrics for characterizing diffusivity in DT-MRI brain data sets. Our approach is a simple and effective addition to existing methods. We also show that probability density entropy itself is a useful biomarker. Note that, this chapter was originally published in significantly similar form at [49].

9.2 Related Work

Using voxel or integral curve (as in tractography) representations of fiber tracts, previous research has proposed tract-based analysis of diffusion-derived measures in order to increase the specificity and robustness of related statistics [144, 123, 74]. A typical approach in these earlier studies is to find a representative skeleton (or curve) first and then project the diffusivity measures of individual subjects on this representative.

For example, tract-based spatial statistics (TBSS) has been proposed to overcome some of the disadvantages of approaches using voxel-based morphometry (VBM) [144]. TBSS reduces the sensitivity of statistics to alignment problems by projecting subject data to a voxel skeleton of the white matter as characterized by FA. Using tractography makes the incorporation of fiber tract arc length into the analysis easier. One of the joint densities evaluated here is the probability of a diffusivity value given an arc-length distance. Parametrization of tracts based on arc-length distance has already been shown to be useful [39], and tract arc length by itself has been proposed as a metric [41]. As in TBSS, previous studies using tractography have also been able to detect differences between groups of subjects by analyzing diffusivity measures projected onto an arc-length parametrized representative tract [123, 74].

Recently, probability density functions (PDFs) have been used as shape descriptors in cortical folding analysis and in quantifying FA change with respect to the thickness of the surface sheet covering the corticospinal tract [10, 162].

While earlier tract-based analysis methods essentially collapse bundles into a single representative subset (skeleton or curve) and run the subsequent analysis through this subset, our approach considers each bundle directly and separately, modeling the full spectrum of individual bundle's data as a sample from a probability distribution characterizing the corresponding subject. Also, our use of PDFs is simple, general (can be applied to any tracts and combine any number of variables), and truly spatial, since we orient tracts.

9.3 Data Collection

The underlying premise of using tract-based probability densities in clinical research is that they have different values in patients with and without known white-matter injury. In order to verify this, we generated tractograms of a group of patients with vascular cognitive impairment (VCI) and a group of healthy control subjects and compared the group differences using probability densities. VCI is a general term for vascular cognitive deficits caused by injuries to the white-matter from cerebrovascular diseases. VCI typically manifests itself with problems in speech, language, and in the ability to organize thoughts and track actions. Its effects on memory, however, are considered mild in comparison to Alzheimer's disease. Often, deficits in cognitive domains caused by VCI are similar to those common in individuals of advanced age. Below we give only a summary of our data collection process; details can be found in [41].

Subjects

Our subjects were 19 patients with VCI and 20 healthy individuals who served as a control group. The subjects were age-matched. Diffusion-weighted MRI (DWI) sequences for each subject's brain were acquired on a 1.5T Siemens Symphony scanner with the following parameters in 12 bipolar diffusion encoding gradient directions: thickness = 1.7 mm, FOV = 21.7 cm $\times$ 21.7 cm, TR = 7200 ms, TE = 156 ms, b = 1000, and NEX = 3.

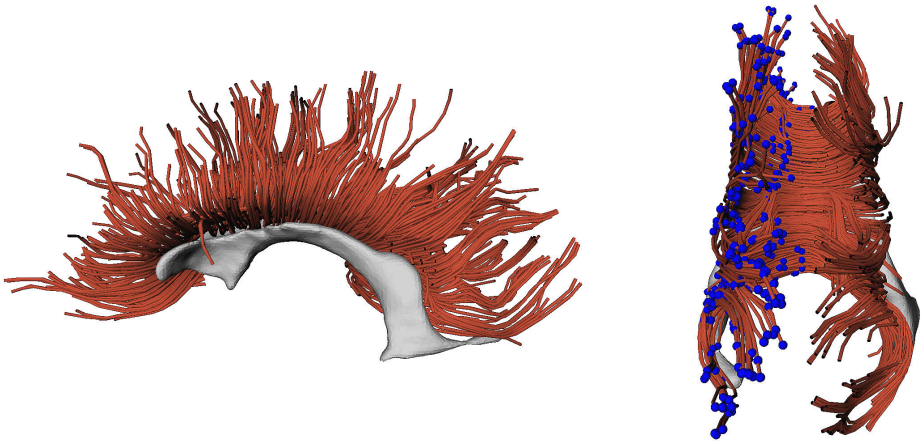

Figure 9.1: Transcallosal fiber tracts obtained from a healthy subject. The gray surface representing the ventricles is displayed for anatomical correspondence. Left: Sagittal view. Right: Tilted coronal view, where spheres (blue in color) indicate the "beginning" of each tract.

Fiber-tract Generation

From acquired sequences, we calculate tensors and then derive the three principal eigenvalues and eigenvectors for each image voxel after interleaving the three sets of DWIs. We then generate fiber-tract models of the whole brain for each subject by integrating (second-order Runge-Kutta integration used) the major eigenvector field of the diffusion tensor field bidirectionally at seed points (see Figure 9.1). Transcollosal fibers, defined as all trajectories passing through the corpus callosum, for each participant were selected manually by a rater using a custom interactive visualization program. Selected models were checked by two experts for anatomical correctness. Since tract of interest (TOI) selection was based on each participant's own anatomy, registration was not necessary.

Orienting Tracts

Since we compute the joint probabilities of arc length and the diffusivity measures, we need to assign start- and end-point designations consistently within and across fiber tracts and across subjects. We achieve this by simply computing the start-to-end vector for each curve and iteratively reorienting the curves until all the vectors are in the same half-space; we repeat this interactively between data sets.

9.4 Methods

Density Estimation

Since we do not assume a particular distribution of diffusivity measures along tracts, we take a nonparametric approach to estimation. We use two widely studied estimators, histogram and kernel estimators [141, 139]. Let $X_1^N = \{X_1, \ldots, X_N\}$ be independent and identically distributed (i.i.d.) random variables drawn from a

density $f(x)$. The histogram estimator of f is obtained simply by dividing the range of the samples into bins of width h and counting the number of samples falling into each bin.

Kernel density estimation (KDE) addresses some of the obvious disadvantages of histograms, including non-differentiability and reference-point dependency. The kernel estimator of f using X_1^N is given by $\hat{f}(x) = \frac{1}{Nh}\sum_{i=1}^{N} K(\frac{x-X_i}{h})$, where $K(x)$ is some kernel satisfying $\int_{-\infty}^{\infty} K(x)dx = 1$ and h is the kernel *bandwidth*, analogous to the histogram bin width. We omit here multivariate versions of the equations; they can be found in any good multivariate statistical analysis book (e.g., [141, 139]).

Both bin width and kernel size can affect the results significantly. Therefore we use a data-driven rule originated from the L_2 theory of histograms, minimizing the mean integrated square error (MISE) [139, 68]. We set the histogram bin width $h = 3.5\hat{\sigma}N^{-\frac{1}{d+2}}$, where d is the dimensionality of the data (e.g., $d = 2$ for a bivariate distribution) and N is the number of samples. We select $\hat{\sigma} = min\{s, IQR/1.349\}$ as discussed in [154], where s is the standard deviation of samples and IQR is the inter-quantile range. We use a Gaussian kernel with both univariate and bivariate estimators. As for the kernel bandwidth h, we apply the "normal scale" rule introduced in [141], which is also asymptotically optimal.

Non-parametric Multivariate Hypothesis Testing

Both histogram and kernel estimates (kernel estimators in their discretized forms) essentially provide a vector of values identifying the underlying distribution. These "density vectors" can be used for testing differences between groups, which requires a multivariate test statistic. We use the permutation Hotelling's T^2 test [73]. Hottelling's T^2 statistic is a multivariate extension of Student's t statistic and permutation is a common method for running non-parametric tests using scores obtained with parametric statistics.

Entropy

One of the usual quantities computed over probability densities is entropy (or Shannon entropy) [42]. While entropy has different interpretations in different domains, it can be viewed here as a measure of randomness in data (i.e. uniformity of its distribution). We compute the entropy of the density f using the estimated density as $H(\hat{f}) = -\sum_x x\hat{f}(x)$ and evaluate it in group comparison of our subjects.

Classification

Classification is now a standard approach to exploring medical data sets. We use a support vector machine (SVM) classifier to assess the usefulness of quantities derived from tract-based diffusivity probability densities. SVM is a maximum-margin classifier minimizing the classification error while maximizing the geometric margin between the classes [20]. It is particularly suitable for two-class problems.

9.5 Results

We estimate univariate densities of FA, MD, AD, and RD as well as joint densities of these measures with fiber tract arc length, where arc length is normalized to 1, in each subject's transcallosal fiber tract. Overall,

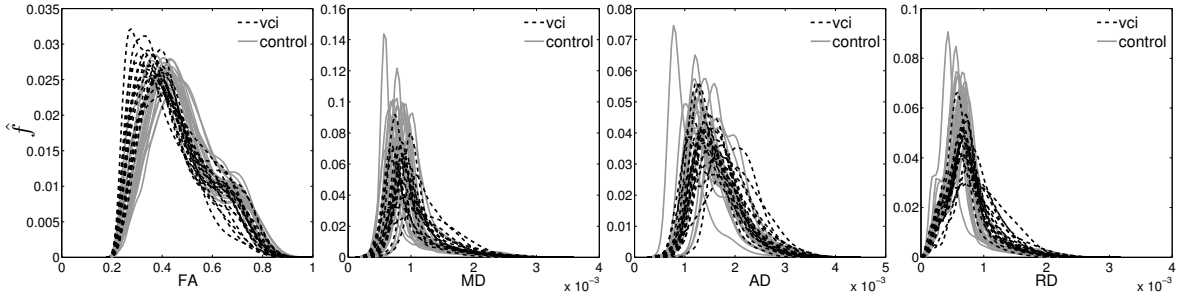

Figure 9.2: Univariate density estimates of FA, MD, AD, and RD measures for control subjects and subjects with VCI. They suggest an increase in variances of MD, AD, RD densities with VCI. Also, the shifts in the densities of the two groups are consistent with previous reports: while FA decreases, MD, RD and, to a lesser extent, AD increase with atrophy in the brain.

results show that the tract-based probability densities help detect differences between groups and classify individual subjects with high accuracy. Figure 9.2 shows kernel estimates for univariate FA, MD, AD, and RD probability densities. Even in this simple univariate form, estimated PDFs are rich in information. They suggest an increase in variances of MD, AD, RD densities with VCI. Also, the shifts in the estimated densities of the two groups support previous reports that, while FA decreases, MD, RD and, to a lesser extent, AD increase with atrophy in the brain. Figure 9.3 shows kernel estimates for the joint probability of FA with the arc length. These 2D density results are visually informative and can help generate hypotheses. For example, it is clear from the figures that the third local maximum (on the mid-right) starts disappearing with VCI and age. Similarly, the decrease in MD with VCI is clear in MD and the arc length joint density, shown in Figure 9.4.

9.5.1 Group Comparisons

Using Entropies

We expect neurological diseases to affect the entropy of probability distributions of diffusivity-related measures. We test this hypothesis for all the measures using their univariate and bivariate densities. The results show that MD, RD, and AD entropies increase significantly while FA decreases (see Figure 9.5 and Table 9.1a). This outcome is in line with previous findings that AD is more sensitive to atrophy than FA [109].

Density Vectors

Using vectorized "raw" density values, we compare the group of subjects with VCI with our control group, employing the multivariate permutation Hotelling test [73, 63]. We reduce the dimensionality of the density vectors using principal component analysis (PCA) before running the test, while preserving 99% of the variance. The results show that the density vectors of univariate estimates are able to detect group differences (see Table 9.1b).

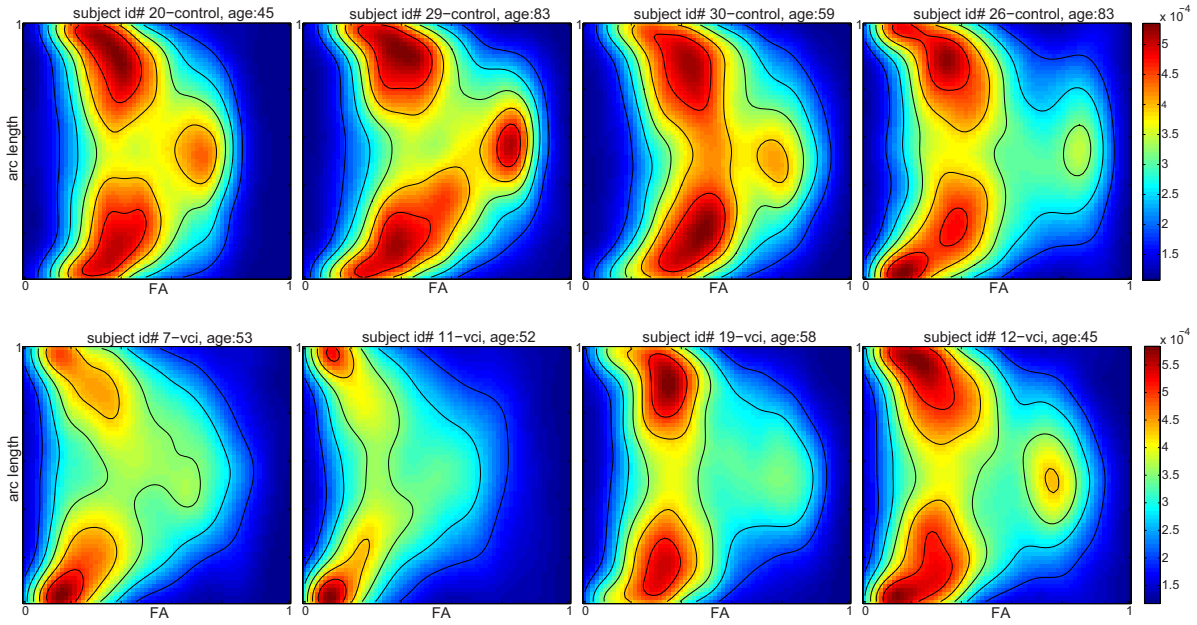

Figure 9.3: Kernel density estimates of FA and arc length joint PDF for 8 subjects–4 controls (first row) and 4 patients (second row). The vertical axis corresponds to the normalized arc-length, horizontal axis corresponds to the range of FA values and the color intensities represent the estimated probability (density) values. The shapes of the density functions are of particular interest in this case. For example, the third local maximum (on the mid-right) starts disappearing with VCI. Also, observe, in the last column on the right, the similarity between a healthy subject of advanced age and a young patient.

	univariate		bivariate	
	histogram	kernel	histogram	kernel
FA	0.36797	0.01239	0.0297	0.8379
MD	0.00000	0.00000	0.00000	0.00000
AD	0.00106	0.00878	0.0042	0.00016
RD	0.00000	0.00000	0.00000	0.00000

(a)

	univariate
FA	0.00078
MD	0.00000
AD	0.00005
RD	0.00002

(b)

Table 9.1: (a) p-values for the group comparison (VCI vs. control) using entropies of the estimated probability densities. Except for the FA, the entropy measure was able to detect the group difference for all densities. (b) p-values for the group comparison (VCI vs. control) using probability density vectors. Histogram estimates of univariate densities were able to detect the group differences.

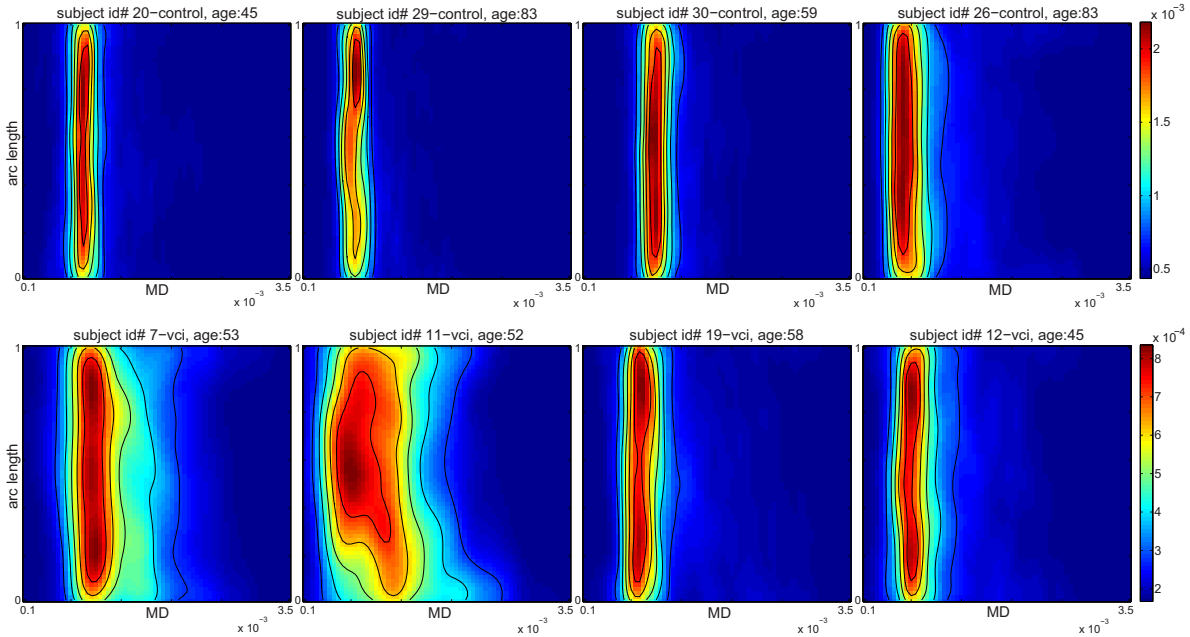

Figure 9.4: Kernel density estimates of MD and arc length joint PDF for 8 subjects–4 controls (first row) and 4 patients (second row). The vertical axis corresponds to the normalized arc-length, horizontal axis corresponds to the range of MD values and the color intensities represent the estimated probability (density) values. The joint PDFs tend to spread out (i.e., increase in variance) with VCI. As seen in the last column, aging can have an effect similar to VCI.

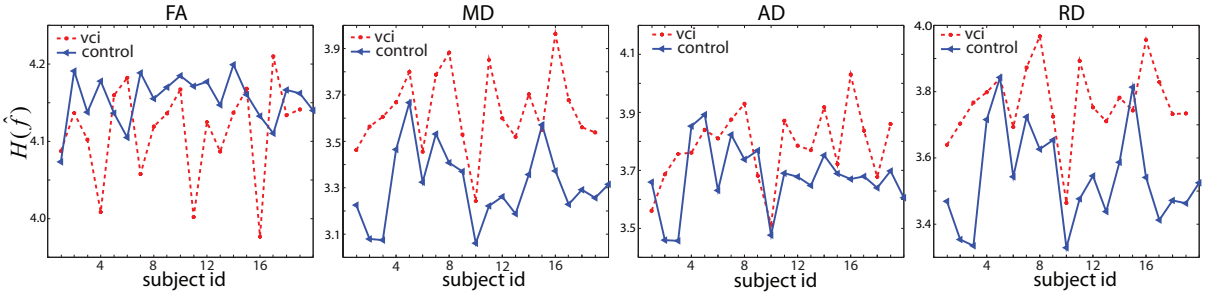

Figure 9.5: Entropies for univariate probability density estimates of FA, MD, AD, and RD for subjects. MD, AD, and RD density entropies significantly increase with VCI, and the FA entropy decreases, though not significantly (see also Table 9.1b).

density	sensitivity	specificity	error
FA	83.3%	81%	17.9%
MD	90%	94.7%	8%
AD	77.2%	88.2%	17.9%
RD	89.5%	90%	10.2%

Table 9.2: Classification results for our 39 subjects, obtained using an SVM classifier with 10-fold cross-validation. Data points used for classification are univariate histogram density estimates.

9.5.2 Classification of Normal and Pathological Data Sets

We ran an SVM with a polynomial kernel on 39 data sets using 10-fold cross-validation. Using density estimates of MD yields the best classification. Table 9.2 summarizes the results.

9.6 Discussion and Conclusions

We have proposed a simple and practical addition to existing tract-based multi-subject diffusion data analysis techniques, which are essentially based on collapsing bundles into a representative subset and sampling of diffusion measures along this subset. Our approach considers all the tracts in a given bundle and treats the associated diffusivity measures as samples from a PDF characterizing the subject. We estimate this PDF using the histogram and kernel estimators. While using the histogram density estimator is fast and simple, the kernel estimator provides smoother results with useful differential properties.

Given a reasonable choice of variables, estimated PDFs are likely to be descriptive and rich with information, and therefore well suited for analysis of individual and group differences in diffusion data sets. In the case examined here, both PDFs themselves and their entropy measures appear to be good indicators of group differences, although additional work is required to understand how this generalizes to other cases. Similarly, our initial results on the classification of subjects based on densities suggest the potential utility of PDFs in detecting individual differences. Note that using natural distance measures on PDFs such as the Bhattacharyya distance and Kullback-Leibler (KL) divergence could further improve the analysis discussed above. Finally, although we have not dealt with alignment issues, the probabilistic approach discussed here can also increase the robustness of the analysis to registration errors in fully automated settings.

Chapter 10

Summary and Conclusions

Computational brain connectivity is the study of efficient algorithms, models, and tools for exploring, understanding, and characterizing brain connectivity. This dissertation presents our contributions to modeling, visualization, and analysis of structural brain connectivity data obtained using diffusion MRI.

The diffusion tensor model of diffusion at each voxel is the most common model of diffusion used in diffusion MRI. It is well known that a single tensor is inadequate for modeling a mixture of fiber orientations, since it prevents reconstruction of configurations such as crossing, branching, and tangential fibers. In Chapter 3, we presented a proof of concept for efficiently modeling tensor configuration distributions with MRFs and inferring the most likely tensor configurations with MAP estimations. We demonstrated the plausibility of our method by resolving fiber crossings in a synthetic dataset. The MRF model proposed here needs to be validated on real datasets with different noise levels. In general, fitting multiple tensors to diffusion MRI measurements using a limited number of diffusion gradient directions is an underconstrained problem. Our MRF model attempts to bring additional constraints from first-order neighbors (local context). Further constraints can be added by using anatomical priors learned from training datasets, where the ground truth fiber configuration at each voxel is known, using long-range features.

A typical clinical research study that uses tractography datasets involves selection of a bundle of interest or two across dozens of different datasets. Tractography datasets obtained using diffusion MRI have, however, the complexity proportional to the axonal connectivity of the brain, and any interactive tool that aims to make these datasets useful for practitioners needs to address this complexity at levels of representation and interaction. With this in mind, in Chapter 4 we described and compared two low-dimensional representations, 2D point and 2D path, for neural tracts. The goal is to facilitate exploration of dense tractograms (collection of curves representing neural tracts) by reducing visual complexity both in static representations and interaction. The proposed planar maps have desirable properties such as visual clarity, easy of tract-of-interest selection, and multiscale hierarchy. The 2D path representation also conveys the anatomical familiarity of 3D brain models and cross-sectional views. We demonstrated the utility of the representations in two interactive systems in which the views and interactions of the standard 3D streamtube representation are linked to those of the planar representations. We also demonstrated a web interface that integrates precomputed

neural-path representations into a geographical digital-map framework with associated labels, metrics, statistics, and linkouts. We compared the two representations both anecdotally and quantitatively via expert input. Results indicate that the planar path representation is more intuitive and easier to use and learn. Similarly, users are faster and more accurate in selecting bundles using the path representation than the 2D point representation. Finally, expert feedback on the web interface suggests that it can be useful for collaboration as well as quick exploration of data. This work can be extended in several directions. For example, we can parse anatomical and functional cortical and subcortical gray-matter structures onto map representations, generating the labels and linkouts (e.g., PubMED, WebMD, Wikipedia, etc.) at several scales automatically. There are exciting possibilities here: extracting images of the labeled regions from web documents, annotating with crowdsourcing, querying paths between cortical regions, etc. Similarly, we can pursue an extension to electron microscopy (EM). Extraction of axonal connections from EM datasets results in extremely dense forest structures. Generating stylized 2D map representations, like Cajal's drawings, that emphasize the synaptic connections can improve understanding of these structures.

Color mapping is a fundamental visualization technique and can facilitate exploration of brain connectivity by characterizing structures qualitatively. Chapters 5 and 6 presented coloring techniques formalizing the intuition that a good coloring of data points should map similar data points to similar colors and dissimilar data points to dissimilar colors and that if the data points change smoothly, the color mapping should change smoothly as well. In Chapter 5, we reduced the problem of smooth one-to-one coloring of 3D line fields to the problem of embedding RP^2 in R^3 and hence showed that the problem has no solution. We introduced Boy's immersion of RP^2 as a method for such a coloring. Given that a smooth, one-to-one coloring for line fields is impossible in general, the coloring based on Boy's surface is a good approximation, because it is one-to-one and smooth except along the self-intersection curve of Boy's surface. While we applied our model to visualization of data derived from diffusion MRI, it is likely to be useful in any domain where data instances are represented as lines, including edge, orientation, and rotation fields.

While the set of 3D lines through the origin corresponds to a surface, this is not the case for most data types. Therefore, in Chapter 6, we introduced a general coloring method that conveys spatial relations among neural fiber tracts effectively; similar tracts are assigned similar colors and different tracts are assigned different colors in a smooth and continuous manner. To this end, we combined a standard spectral approach with a mass-spring heuristic to embed fiber tracts into a perceptually uniform color space, L*a*b*. We also introduced a new geometric bivariate coloring model, the flat torus, that allows finer adjustment of coloring arbitrarily. Results indicate that our method allows quick evaluation of tract projections (endpoints).

In Chapter 7, we introduced a novel neural tract cluster coherence measure obtained by combining an intuitive geometric idea, slicing, with known statistical machine-learning techniques. An important property of our method is its reliance on the context of the cluster in evaluating similarity among curves. The quantities computed in each slice, and across slices, are tied directly to the clustering objective. This is in contrast to the more standard setup in which distances and thresholds are defined in the global context. Our method works reasonably well on clusters with curves having high curvatures and varying arc lengths. While our motivation was 3D DTI curves and the experimental evaluation was on such curves, the presented method generalizes to higher and lower dimensions easily and may apply to curve data in other domains. We believe that a

number of technical aspects could be improved. Specifically, a more robust slicing and projection method that explicitly down-weights outliers could help reduce uncertainty in the per-slice quantities. Also, while we are not aware of any "gold-standard" DTI fiber-track clustering data set, it would be useful to validate our method more quantitatively. A more challenging extension of the idea of slicing that could be pursued is to build a semi-parametric generative model for a cluster.

In Chapter 8, we presented a global analysis of brain fiber pathways based on persistent homology. This analysis has two main advantages: First, it captures the effects of axonal dropout, which creates "holes" in the white matter space, caused by diffuse neurological disorders. Initial results indicate that the proposed analysis can capture global atrophic effects of neuropathologies in the brain. Second, it provides a topological signature of the structural brain connectivity. The latter can be useful in understanding and classification of the topological structure of the neural pathways in the brain.

In Chapter 9, we introduced a simple and practical addition to existing tract-based multi-subject diffusion data analysis techniques, which are essentially based on collapsing bundles into a representative subset and sampling diffusion measures along this subset. Our approach considers all the tracts in a given bundle and treats the associated diffusivity measures as samples from a PDF characterizing the subject. We estimate this PDF using the histogram and kernel estimators. While using the histogram density estimator is fast and simple, the kernel estimator provides smoother results with useful differential properties. Given a reasonable choice of variables, estimated PDFs are likely to be descriptive and rich with information, and therefore well-suited for analysis of individual and group differences in diffusion data sets. In the case examined here, both the PDFs themselves and their entropy measures appear to be good indicators of group differences, although additional work is required to understand how this generalizes to other cases. Similarly, our initial results on the classification of subjects based on densities suggest the potential utility of PDFs in detecting individual differences. Note that using natural distance measures on PDFs such as the Bhattacharyya distance and Kullback-Leibler (KL) divergence could further improve the analysis discussed above.

Computational brain connectivity is an emerging field that is rich with important and interesting problems. Fusion and multimodal analysis of brain connectivity data (structural, functional, and causal alike) acquired from different sources at different scales, however, will be the most important challenge in the field in years to come. The dissertation presented here, a slice from this new field, establishes that characterizing patterned structures of diffusion MRI-derived brain connectivity improves exploration of the brain's connectivity space and provides measures quantifying normal and pathological variation in the brain.

Bibliography

[1] D. Akers, A. Sherbondy, R. Mackenzie, R. Dougherty, and B. Wandell. Exploration of the brain's white matter pathways with dynamic queries. In *Proc. of Visualization*, pages 377–384, 2004.

[2] David Akers. Wizard of Oz for participatory design: Inventing an interface for 3D selection of neural pathway estimates. In *Proceedings of CHI 2006 Extended Abstracts*, pages 454–459, 2006.

[3] Andrew L. Alexander, Khader M. Hasan, Mariana Lazar, Jay S. Tsuruda, and Dennis L. Parker. Analysis of partial volume effects in diffusion-tensor mri. *Magnetic Resonance in Medicine*, 45(5):770–780, 2001.

[4] Daniel C. Alexander. Multiple-fibre reconstruction algorithms for diffusion mri. *Annals of the New York Academy of Sciences*, 1064:113–133, December 2005.

[5] Pavel Sergeyevich Alexandrov. *Einfachste Grundbegriffe der Topologie*. Julius Springer, 1932.

[6] Bowen Alpern, Larry Carter, Matt Grayson, and Chris Pelkie. Orientation maps: techniques for visualizing rotations (a consumer's guide). In *VIS '93: Proceedings of the 4th conference on Visualization '93*, pages 183–188, 1993.

[7] Francois Apery. *Models of the Real Projective Plane*. Friedrich Vieweg & Sohn Verlag, 1987.

[8] K. Arakawa, S. Tamaki, N. Kono, N. Kido, K. Ikegami, R. Ogawa, and M. Tomita. Genome Projector: zoomable genome map with multiple views. *BMC bioinformatics*, 10(1):31, 2009.

[9] Yaniv Assaf, Raisa Z. Freidlin, Gustavo K. Rohde, and Peter J. Basser. New modeling and experimental framework to characterize hindered and restricted water diffusion in brain white matter. *Magnetic Resonance in Medicine*, 52(5):965–978, 2004.

[10] Suyash P. Awate, Paul A. Yushkevich, Zhuang Song, Daniel J. Licht, and James C. Gee. Multivariate high-dimensional cortical folding analysis, combining complexity and shape, in neonates with congenital heart disease. In *IPMI*, pages 552–563, 2009.

[11] Angelos Barmpoutis, Bing Jian, and Baba C. Vemuri. Adaptive kernels for multi-fiber reconstruction. In *Proceedings of the 21st International Conference on Information Processing in Medical Imaging*, IPMI '09, pages 338–349, Berlin, Heidelberg, 2009. Springer-Verlag.

[12] Peter J. Basser, James Mattiello, and Denis LeBihan. Estimation of the effective self-diffusion tensor from the NMR spin echo. *J Magn Reson B*, 103(3):247–254, March 1994.

[13] Petter J. Basser and Derek K. Jones. Diffusion-tensor mri: theory, experimental design and data analysis - a technical review. *NMR in Biomedicine*, 15:456–467, 2002.

[14] Mark F. Bear, Barry W. Connors, and Michael A. Paradiso. *Principles of Neural Science*. Lippincott Williams & Wilkins, 3rd edition, 1985.

[15] Christian Beaulieu. The biological basis of diffusion anisotropy. In Heidi Johansen-Berg and Timothy E. J. Behrens, editors, *Diffusion MRI*, pages 105–126. Academic Press, 2009.

[16] T E Behrens, H Johansen-Berg, M W Woolrich, S M Smith, C A Wheeler-Kingshott, P A Boulby, G J Barker, E L Sillery, K Sheehan, O Ciccarelli, A J Thompson, J M Brady, and P M Matthews. Non-invasive mapping of connections between human thalamus and cortex using diffusion imaging. *Nat Neurosci*, 6(7):750–7, 2003. 1097-6256 (Print)Journal Article.

[17] T E Behrens, M W Woolrich, M Jenkinson, H Johansen-Berg, R G Nunes, S Clare, P M Matthews, J M Brady, and S M Smith. Characterization and propagation of uncertainty in diffusion-weighted mr imaging. *Magn Reson Med*, 50(5):1077–88, 2003. Journal ArticleResearch Support, Non-U.S. Gov'tUnited Statesofficial journal of the Society of Magnetic Resonance in Medicine / Society of Magnetic Resonance in Medicine.

[18] Denis Le Bihan, Jean-Francois Mangin, Cyril Poupon, Chris A. Clark, Sabina Pappata, Nicolas Molko, and Hughes Chabriat. Diffusion tensor imaging: Concepts and applications. *Journal of Magnetic Resonance Imaging*, 13:534–546, 2001.

[19] C. M. Bishop. *Pattern Recognition and Machine Learning*. Springer, 2006.

[20] C. M. Bishop. *Pattern Recognition and Machine Learning*. Springer, 2006.

[21] Davi D. Bock, Wei-Chung A. Lee, Aaron M. Kerlin, Mark L. Andermann, Greg Hood, Arthur W. Wetzel, Sergey Yurgenson, Edward R. Soucy, Hyon S. Kim, and R. Clay Reid. *Nature*, 471(7337):177–182, 2011.

[22] Benedetta Bodini and Olga Ciccarelli. Diffusion mri in neurological disorders. In Heidi Johansen-Berg and Timothy E. J. Behrens, editors, *Diffusion MRI*, pages 175–203. Academic Press, 2009.

[23] M. Bostock and J. Heer. Protovis: A graphical toolkit for visualization. *IEEE Transactions on Visualization and Computer Graphics*, 15(6):1121–1128, 2009.

[24] Werner Boy. *Über die Curvatura Integra und dieTopologie der Geschlossener Flachen*. PhD thesis, Universität Göttingen, Göttingen, 1901.

[25] James N. Breckenridge. Replicating cluster analysis: Method, consistency, and validity. *Multivariate Behavioral Research*, 24(2):147–161, 1989.

[26] KL Briggman, M Helmstaedter, and W Denk. Wiring specificity in the direction-selectivity circuit of the retina. *Nature*, 471(5):183–8, 2011-03-10 00:00:00.0.

[27] A. Brun, H. Knutsson, H.-J Park, M.E. Shenton, and C.-F. Westin. Clustering fiber traces using normalized cuts. In *MICCAI*, pages 368–375, 2004.

[28] Anders Brun, Hae-Jeong Park, Hans Knutsson, and Carl-Fredrik Westin. Coloring of dt-mri fiber traces using laplacian eigenmaps. In *EUROCAST'03*, pages 564–572, 2003.

[29] Robert L. Bryant. Surfaces in conformal geometry. In *The mathematical heritage of Hermann Weyl (Durham, NC, 1987)*, volume 48 of *Proc. Sympos. Pure Math.*, pages 227–240. Amer. Math. Soc., Providence, RI, 1988.

[30] William Elwood Byerly. *An Elementary Treatise on Fourier's Series, and Spherical, Cylindrical, and Ellipsoidal Harmonics, with Applications to Problems in Mathematical Physics*. Dover, New York, 1959.

[31] G. Carlsson, T. Ishkhanov, V. de Silva, and A. Zomorodian. On the local behavior of spaces of natural images. *International Journal of Computer Vision*, 76(1):1–12, 2008.

[32] M. Catani, R.J. Howard, S. Pajevic, and D.K. Jones. Virtual in vivo interactive dissection of white matter fasciculi in the human brain. *NeuroImage*, 17:1:77–94, 2002.

[33] F. Cazals, J. Giesen, M. Pauly, and A. Zomorodian. The conformal alpha shape filtration. *The Visual Computer*, 22(8):531–540, 2006.

[34] Çağatay Demiralp, Gregory Shakhnarovich, Song Zhang, and David H. Laidlaw. Slicing-based coherence measure for refining clusters of 3D curves. In *Proceedings of MICCAI*, 2008.

[35] H. Chabriat, Sabina Pappata, C. Poupon, C. A. Clark, K. Vahedi, F. Poupon, J.-F. Mangin, M. Pachot-Clouard, Antoinette Jobert, Denis LeBihan, and M.-G. Bousser. Clinical severity in CADASIL related to ultrastructural damage in white matter: in vivo study with diffusion tensor mri. *Stroke*, 30(12):2637–2643, 1999.

[36] M. Chalmers. A linear iteration time layout algorithm for visualising high-dimensional data. In *Proceedings of the 7th conference on Visualization'96*. IEEE Computer Society Press Los Alamitos, CA, USA, 1996.

[37] Wei Chen, Ziang Ding, Song Zhang, Anna MacKay-Brandt, Stephen Correia, Huamin Qu, John Allen Crow, David F. Tate, Zhicheng Yan, and Qunsheng Peng. A novel interface for interactive exploration of DTI fibers. *IEEE TVCG (Proc. of Visualization)*, 2009.

[38] Ali Civril, Malik Magdon-Ismail, and Eli Bocek-Rivele. Sde: Graph drawing using spectral distance embedding. In *GD'05*, Limerick, Ireland, 2005.

[39] Isabelle Corouge, P. Thomas Fletcher, Sarang Joshi, Sylvain Gouttard, and Guido Gerig. Fiber tract-oriented statistics for quantitative diffusion tensor MRI analysis. *Medical Image Analysis*, 10(5):786–798, 2006.

[40] Isabelle Corouge, Sylvain Gouttard, and Guido Gerig. Towards a shape model of white matter fiber bundles using diffusion tensor mri. In *International Symposium on Biomedical Imaging*, pages 344–347, 2004.

[41] Stephen Correia, Stephanie Y. Lee, Thom Voorn, David F. Tate, Robert H. Paul, Song Zhang, Stephen P. Salloway, Paul F. Malloy, and David H. Laidlaw. Quantitative tractography metrics of white matter integrity in diffusion-tensor MRI. *NeuroImage*, 42(2):568–581, 2008.

[42] Thomas M. Cover and Joy A. Thomas. *Elements of Information Theory*. Wiley-Interscience, New York, NY, USA, 1991.

[43] C.M. Danis, F.B. Viegas, and M. Wattenberg. Your place or mine? Visualization as a community component. In *Proceedings of CHI*, 2008.

[44] Jan de Leeuw. Applications of convex analysis to multidimensional scaling. In J.R. Barra, F. Brodeau, G. Romier, and B. Van Cutsem, editors, *Recent Developments in Statistics*, pages 133–146. North Holland Publishing Company, 1977.

[45] Çağatay Demiralp, John F. Hughes, and David H. Laidlaw. Using Boy's real projective plane immersion for coloring DT-MRI slices. In *Proceedings of ISMRM*, 2008.

[46] Çağatay Demiralp, John F. Hughes, and David H. Laidlaw. Coloring 3d line fields using boy's real projective plane immersion. *IEEE Transactions on Visualization and Computer Graphics*, 15(6):1457–1464, 2009.

[47] Çağatay Demiralp, Radu Jianu, and David H. Laidlaw. Exploring brain connectivity with two-dimensional maps. In *New Developments in the Visualization and Processing of Tensor Fields*. LNCS-Springer, 2012.

[48] Çağatay Demiralp and David H. Laidlaw. Similarity coloring of DTI fiber tracts. In *Proceedings of DMFC Workshop at MICCAI*, 2009.

[49] Çağatay Demiralp and David H. Laidlaw. Tract-based probability densities of diffusivity measures in dt-mri. In *Proceedings of the 13th international conference on Medical image computing and computer-assisted intervention: Part I*, MICCAI'10, pages 542–549, Berlin, Heidelberg, 2010. Springer-Verlag.

[50] Çağatay Demiralp and David H. Laidlaw. Generalizing diffusion tensor model using probabilistic inference in Markov random fields. In *Proceedings of MICCAI Computational Diffusion MRI Workshop*, 2011.

[51] Çağatay Demiralp and David H. Laidlaw. Generalizing diffusion tensor model using probabilistic inference in Markov random fields. In *Proceedings of ISMRM*, 2011.

[52] Çağatay Demiralp, Song Zhang, David Tate, Stephen Correia, and David H. Laidlaw. Connectivity-aware sectional visualization of 3DDTI volumes using perceptual flat-torus coloring and edge rendering. In *Proceedings of Eurographics*, 2006.

[53] Çağatay Demiralp and David H. Laidlaw. Tract-based probability densities of diffusivity measures in dt-mri. In *Proceedings of the 13th international conference on Medical image computing and computer-assisted intervention: Part I*, MICCAI'10, pages 542–549, 2010.

[54] Rachid Deriche, Jeff Calder, and Maxime Descoteaux. Optimal real-time q-ball imaging using regularized kalman filtering with incremental orientation sets. *Medical Image Analysis*, 13(4):564 – 579, 2009.

[55] Zhaohua Ding, John C. Gore, and Adam W. Anderson. Case study: reconstruction, visualization and quantification of neuronal fiber pathways. In *VIS '01: Proceedings of the conference on Visualization '01*, pages 453–456. IEEE Computer Society, 2001.

[56] Manfredo P. do Carmo. *Differential Geometry of Curves and Surfaces*. Prentice-Hall, 1976.

[57] A. T. Du, N. Schuff, D. Amend, M. P. Laakso, Y. Y. Hsu, W. J. Jagust, K. Yaffe, J. H. Kramer, B. Reed, D. Norman, H. C. Chui, and M. W. Weiner. Magnetic resonance imaging of the entorhinal cortex and hippocampus in mild cognitive impairment and alzheimer's disease. *J Neurol Neurosurg Psychiatry*, 71(4):441–447, October 2001.

[58] R. O. Duda, P. E. Hart, and D. G. Stork. *Pattern Classification*. Wiley-Interscience, 2nd edition, 2000.

[59] P. Eades. A heuristic for graph drawing. *Congressus Numerantium*, 42(149160):194–202, 1984.

[60] Herbert Edelsbrunner and John L. Harer. *Computational Topology: An Introduction*. AMS, 2010.

[61] Herbert Edelsbrunner and Ernst P. Mücke. Three-dimensional alpha shapes. *ACM Trans. Graph.*, 13:43–72, January 1994.

[62] Julia M. Edgar and Ian R. Griffiths. White matter structure: A microscopist's view. In Heidi Johansen-Berg and Timothy E. J. Behrens, editors, *Diffusion MRI*, pages 75–103. Academic Press, 2009.

[63] B. Efron and R. J. Tibshirani. *An Introduction to the Bootstrap*. Chapman & Hall, New York, 1993.

[64] C. M. Ellis, A. Simmons, D. K. Jones, J. Bland, J. M. Dawson, M. A. Horsfield, S. C. Williams, and P. N. Leigh. Diffusion tensor mri assesses corticospinal tract damage in als. *Neurology*, 53(5):1051–1058, September 1999.

[65] Mark D. Fairchild. *Color Appearance Models*. WileyBlackwell, 2004.

[66] Christopher G. Filippi, Aziz M. Ulug, Elizabeth Ryan, Steven J. Ferrando, and Wilfred van Gorp. Diffusion tensor imaging of patients with hiv and normal-appearing white matter on mr images of the brain. *American Journal of Neuroradiology*, 22:277–283, February 2001.

[67] George K. Francis. *A Topological Picture Book*. Springer, 2006.

[68] David Freedman and Persi Diaconis. On the histogram as a density estimator:L_2 theory. *Probability Theory and Related Fields*, 57(4):453–476, December 1981.

[69] Jane Fridlyand and Sandrine Dudoit. Applications of resampling methods to estimate the number of clusters and to improve the accuracy of a clustering method. Technical report, Division of Biostatistics, School of Public Healty, UC Berkeley, 2001.

[70] G3D. `http://g3d-cpp.sourceforge.net/`.

[71] Stuart Geman and Christine Graffigne. Markov random field image models and their applications to computer vision. In *Proceedings of the International Congress of Mathematicians*, 1986.

[72] Peter J. Giblin. *Graphs, Surfaces and Homology: An Introduction to Algebraic Topology*. Chapman and Hall, 2nd edition, 1981.

[73] Phillip I. Good. *Permutation, Parametric, and Bootstrap Tests of Hypotheses (Springer Series in Statistics)*. Springer-Verlag New York, Inc., Secaucus, NJ, USA, 2004.

[74] Casey B. Goodlett, P. Thomas Fletcher, John H. Gilmore, and Guido Gerig. Group statistics of DTI fiber bundles using spatial functions of tensor measures. In *Proceedings of MICCAI*, pages 1068–1075, Berlin, Heidelberg, 2008. Springer-Verlag.

[75] GoogleMapsAPI. `http://code.google.com/apis/maps/`.

[76] J. M. Graham, N. Papadakis, J. Evans, E. Widjaja, C.A.J. Romanowski, M.N.J. Paley, L. I. Wallis, I.D. Wilkinson, P. J. Shaw, and P. D. Griffiths. Diffusion tensor imaging for the assessment of upper motor neuron integrity in als. *Neurology*, 63:2111–2119, 20-4.

[77] J. M. Hammersley and P. Clifford. Markov fields of finite graphs and lattices. Technical report, University of California-Berkeley, preprint, 1968.

[78] H. Hanyu, T. Asano, H. Sakurai, Y. Imon, T. Iwamoto, M. Takasaki, H. Shindo, and K. Abe. Diffusion-weighted and magnetization transfer imaging of the corpus callosum in alzheimer's disease. *Journal of the neurological sciences*, 167(1):37–44, August 1999.

[79] Allen Hatcher. *Algebraic Topology*. Cambridge Press, 2002.

[80] Christopher G. Healey. Choosing effective colours for data visualization. In *VIS '96*, pages 263–ff., 1996.

[81] Suzana Herculano-Houzel. The remarkable, yet not extraordinary, human brain as a scaled-up primate brain and its associated cost. *Proceedings of the National Academy of Sciences*, 109(Supplement 1):10661–10668, 2012.

[82] Danny Holten. Hierarchical edge bundles: Visualization of adjacency relations in hierarchical data. *IEEE Transactions on Visualization and Computer Graphics*, 12:741–748, 2006.

[83] M. A. Horsfield, M. Lai, and S. Webb. Apparent diffusion coefficients in benign and secondary progressive multiple sclerosis by nuclear magnetic resonance. *Magn. Reson. Med.*, 36:393–400, 1996.

[84] Viren Jain, H Sebastian Seung, and Srinivas C Turaga. Machines that learn to segment images: a crucial technology for connectomics. *Current Opinion in Neurobiology*, 20(5):653 – 666, 2010.

[85] Won-Ki Jeong, Johanna Beyer, Markus Hadwiger, Amelio Vazquez, Hanspeter Pfister, and Ross T. Whitaker. Scalable and interactive segmentation and visualization of neural processes in EM datasets. *IEEE Transactions on Visualization and Computer Graphics*, 15:1505–1514, 2009.

[86] R. Jianu and D. Demiralp, cC.and Laidlaw. Exploring 3D DTI fiber tracts with linked 2D representations. *IEEE TVCG (Proc. of Visualization)*, 15(6):1449–1456, 2009.

[87] Radu Jianu, Çağatay Demiralp, and David H. Laidlaw. Exploring brain connectivity with two-dimensional neural maps. In *IEEE Visualization 2010 Poster Compendium*, 2010.

[88] Radu Jianu, Çağatay Demiralp, and David H. Laidlaw. Visualizing and exploring tractograms via two-dimensional connectivity maps. In *Proceedings of ISMRM'10*, 2010.

[89] Radu Jianu, Çağatay Demiralp, and David H. Laidlaw. Exploring brain connectivity with two-dimensional neural maps. *IEEE Transactions on Visualization and Computer Graphics*, 18(6):978–987, June 2012.

[90] Heidi Johansen-Berg and Timothy E. J. Behrens, editors. *Diffusion MRI*. Academic Press, 2009.

[91] D.W. Johnson and TJ Jankun-Kelly. A scalability study of web-native information visualization. In *Proceedings of Graphics Interface*, pages 163–168, 2008.

[92] I. T. Joliffe. *Principal Component Analysis*. Springer, 2nd edition, 2002.

[93] Derek K. Jones, editor. *Diffusion MRI*. Oxford University Press, 2011.

[94] Eric R. Kandel, J. H. Schwartz, and Thomas M. Jessell. *Principles of Neural Science*. Elsevier Science Publishing, 2nd edition, 1985.

[95] Peter M. Kasson, Afra Zomorodian, Sanghyun Park, Nina Singhal, Leonidas J. Guibas, and Vijay S. Pande. Persistent voids: a new structural metric for membrane fusion. *Bioinformatics*, 23(14):1753–1759, 2007.

[96] Gordon Kindlmann, Erik Reinhard, and Sarah Creem. Face-based luminance matching for perceptual colormap generation. In *Procs. of Vis'02*, pages 299–306, 2002.

[97] Kitware, Inc. *The Visualization Toolkit User's Guide*, January 2003.

[98] Jon M. Kleinberg. An impossibility theorem for clustering. In *NIPS*, pages 446–453, 2002.

[99] Rob Kusner. Conformal geometry and complete minimal surfaces. *Bull. Amer. Math. Soc. (N.S.)*, 17(2):291–295, 1987.

[100] Tilman Lange, Volker Roth, Mikio. L. Braun, and Joachim M. Buhmann. Stability-based validation of clustering solutions. *Neural Computation*, pages 1299–1323, 2004.

[101] Ann B. Lee, Kim S. Pedersen, and David Mumford. The nonlinear statistics of high-contrast patches in natural images. *Int. J. Comput. Vision*, 54(1-3):83–103, 2003.

[102] Hyekyoung Lee, Moo K. Chung, Hyejin Kang, Bung-Nyun Kim, and Dong Soo Lee. Discriminative persistent homology of brain networks. In *ISBI*, pages 841–844, 2011.

[103] Haim Levkowitz and Gabor T. Herman. Color scales for image data. *IEEE CG&A*, 12(1):72–80, 1992.

[104] Jeff W. Lichtman and Winfried Denk. The big and the small: Challenges of imaging the brain's circuits. *Science*, 334(6056):618–623, 2011.

[105] J Livet, TA Weissman, H Kang, RW Draft, J Lu, RA Bennis, JR Sanes, and JW Lichtman. Transgenic strategies for combinatorial expression of fluorescent proteins in the nervous system. *Nature*, 450(5):56–62, 2007-11-01 00:00:00.0.

[106] John H. Martin. *Neuroanatomy*. McGraw Hill, 3rd edition, 2007.

[107] William S. Massey. *Algebraic Topology: An Introduction*. Number 56 in Graduate Texts in Mathematics. Springer-Verlag, New York, NY, 1967.

[108] Sylvain Merlet and Rachid Deriche. Compressed Sensing for Accelerated EAP Recovery in Diffusion MRI. In *MICCAI (1)*, 2010.

[109] NS Metwalli, M Benatar, S Usher, G Bhagavatheeshwaran, XP Hu, and JD Carew. Sensitivity of axial and radial diffusivity to the neuropathology of amyotrophic lateral sclerosis. *NeuroImage*, 47(Supplement 1):S128–S128, 2009.

[110] Oleg V. Michailovich and Yogesh Rathi. Fast and accurate reconstruction of hardi data using compressed sensing. In *MICCAI (1)*, pages 607–614, 2010.

[111] V.P. Mineev. Topological objects in nematic liquid crystals. In *Intuitive Combinatorial Topology*. Springer, July 2001.

[112] Bart Moberts, Anna Vilanova, and Jarke J. van Wijk. Evaluation of fiber clustering methods for diffusion tensor imaging. In *Procs. of Vis'05*, pages 65–72, 2005.

[113] N. Molko, S. Pappata, J.-F. Mangin, F. Poupon, D. Le Bihan, M. G. Bousser, and H. Chabriat. Monitoring disease progression in CADASIL with diffusion magnetic resonance imaging: a study with whole brain histogram analysis. *Stroke*, 33(12):2902–2908, 2002.

[114] S. Mori and P.C.M. Van Zijl. Fiber tracking: principles and strategies–a technical review. *NMR in Biomedicine*, 15(7-8):468–480, 2002.

[115] S. Mori and P.C.M. van Zijl. Fiber tracking: principles and strategies–a technical review. *NMR in Biomedicine*, 15(7-8):46e–480, 2002.

[116] Susumu Mori and Jiangyang Zhang. Principles of diffusion tensor imaging and its applications to basic neuroscience research. *Neuron*, 51:527–539, 2006.

[117] Bernard Morin. Équations du retournement de la sphère. *C. R. Acad. Sci. Paris Sér. A-B*, 287(13):A879–A882, 1978.

[118] M. E. Moseley, Y. Cohen, J. Mintorovitch, L. Chileuitt, H. Shimizu, J. Kucharczyk, M. F. Wendland, and P. R. Weinstein. Early detection of regional cerebral ischemia in cats: Comparison of diffusion- and t2-weighted mri and spectroscopy. *Magnetic Resonance in Medicine*, 176:330–346, 1990.

[119] J. Neil, J. Miller, P. Mukherjee, and P. S. Hüppi. Diffusion tensor imaging of normal and injured developing human brain - a technical review. *NMR in Biomedicine*, 15(7-8):543–552, 2002.

[120] A. Y. Ng, M. I. Jordan, and Y. Weiss. On spectral clustering: Analysis and an algorithm. In *Neural Information Processing Systems*, pages 849–856. MIT Press, 2001.

[121] Monica Nicolau, Arnold J. Levine, and Gunnar Carlsson. Topology based data analysis identifies a subgroup of breast cancers with a unique mutational profile and excellent survival. *Proceedings of the National Academy of Sciences*, 2011.

[122] L. O'Donnell, Kubicki M, M. E. Shenton, M. Dreusicke, W. E. L. Grimson, and C.-F. Westin. A method for clustering white matter fiber tracts. *AJNR*, 27(5), 2006.

[123] Lauren J. O'Donnell, Carl-Fredrik Westin, and Alexandra J. Golby. Tract-based morphometry. In *Proceedings of MICCAI*, volume 4792 of *Lecture Notes in Computer Science*, pages 161–168, Brisbane, Australia, 2007.

[124] Sinisa Pajevic and Carlo 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. *MRM*, 42:526–540, 1999.

[125] Sharon Peled, David G. Cory, Stephen A. Raymond, Daniel A. Kirschner, and Ferenc A. Jolesz. Water diffusion, t2, and compartmentation in frog sciatic nerve. *Magnetic Resonance in Medicine*, 42(5):911–918, 1999.

[126] Sharon Peled, Ola Friman, Ferenc Jolesz, and Carl-Fredrik Westin. Geometrically constrained two-tensor model for crossing tracts in DWI. *Magnetic Resonance Imaging*, 24:1263–1270, 2006.

[127] Binh Pham. Spline-based color sequences for univariate, bivariate and trivariate mapping. In *Procs. of Vis'90*, pages 202–208, 1990.

[128] Carlo Pierpaoli. Oh no! one more method for color mapping of fiber tract direction using diffusion mr imaging data. In *Proceedings of 5th ISMRM*, page 1741, 1997.

[129] N. Pomara, D.T. Crandall, S.J. Choi, G. Johnson, and K.O. Lim. White matter abnormalities in HIV-1 infection: a diffusion tensor imaging study. *Psychiatry Res*, 106(1):15–24, 2001.

[130] C. Poupon, C. A. Clark, V. Frouin, J. Rgis, I. Bloch, D. Le Bihan, and J. F. Mangin. Regularization of diffusion-based direction maps for the tracking of brain white matter fascicles. *NeuroImage*, 12(2):184 – 195, 2000.

[131] C. Poupon, J.-F. Mangin, C.A. Clark, V. Frouin, J. Rgis, D. Le Bihan, and I. Bloch. Towards inference of human brain connectivity from mr diffusion tensor data. *Medical Image Analysis*, 5(1):1 – 15, 2001.

[132] Qt. `http://www.qtsoftware.com/`.

[133] A. B. Ragin, P. Storey, B. A. Cohen, L. G. Epstein, and R. R. Edelman. Whole brain diffusion tensor imaging in hiv-associated cognitive impairment. *AJNR Am. J. Neuroradiol.*, 25(2):195–200, 2004.

[134] Philip K. Robertson. Visualizing color gamuts: A user interface for the effective use of perceptual color spaces in data displays. *IEEE CG&A*, 8(5):50–64, 1988.

[135] Stephen E Rosea, Fang Chena, Jonathan B Chalkb, Fernando O Zelayad, Wendy E Strugnellc, Mark Bensonc, James Semplee, and David M Doddrella. Loss of connectivity in alzheimer's disease: an evaluation of white matter tract integrity with colour coded mr diffusion tensor imaging. *J Neurol Neurosurg Psychiatry*, 69:528–530, 2000.

[136] `http://www.cs.brown.edu/~cad/rp2coloring/`.

[137] Miriam Sach, Gerhard Winkler, Volkmar Glauche, Joachim Liepert, Bernhard Heimbach, Martin A. Koch, Christian Buchel, and Cornelius Weiller. Diffusion tensor mri of early upper motor neuron involvement in amyotrophic lateral sclerosis. *Brain*, 127(2):340–350, 2004.

[138] Gideon Schwarz. Estimating the dimension of a model. *The Annals of Statistics*, 6(2):461–464, 1978.

[139] David W. Scott. *Multivariate Density Estimation: Theory, Practice, and Visualization*. Wiley-Interscience, 1992.

[140] Joshua Shimony, Avi Snyder, Nicholas Lori, and Thomas Conturo. Automated fuzzy clustering of neuronal pathways in diffusion tensor tracking. In *Proceedings of ISMRM*, 2002.

[141] Bernard W. Silverman. *Density Estimation for Statistics and Data Analysis*. Monographs on Statistics and Applied Probability. Chapman and Hall, London-New York, 1986.

[142] `http://www.cs.brown.edu/~cad/simcoloring/`.

[143] Gurjeet Singh, Facundo Memoli, Tigran Ishkhanov, Guillermo Sapiro, Gunnar Carlsson, and Dario L. Ringach. Topological analysis of population activity in visual cortex. *J. Vis.*, 8(8):1–18, 6 2008.

[144] Stephen M. Smith, Mark Jenkinson, Heidi Johansen-Berg, Daniel Rueckert, Thomas E. Nichols, Clare E. Mackay, Kate E. Watkins, Olga Ciccarelli, M. Zaheer Cader, Paul M. Matthews, and Timothy E.J. Behrens. Tract-based spatial statistics: voxelwise analysis of multi-subject diffusion data. *NeuroImage*, 31(4):1487–1505, 2006.

[145] Sheng-Kwei Song, Shu-Wei Sun, Michael J. Ramsbottom, Chen Chang, John Russell, and Anne H. Cross. Dysmyelination revealed through MRI as increased radial (but unchanged axial) diffusion of water. *NeuroImage*, 17(3):1429–1436, 2002.

[146] E. O. Stejskal. Use of spin echoes in a pulsed magnetic-field gradient to study anisotropic, restricted diffusion and flow. *The Journal of Chemical Physics*, 43(10):3597–3603, 1965.

[147] Richard Szeliski, Ramin Zabih, Daniel Scharstein, Olga Veksler, Vladimir Kolmogorov, Aseem Agarwala, Marshall Tappen, and Carsten Rother. A comparative study of energy minimization methods for markov random fields with smoothness-based priors. *IEEE Trans. Pattern Anal. Mach. Intell.*, 30:1068–1080, June 2008.

[148] J. E. Tanner. Transient diffusion in system partitioned by permeable barriers. Application to NMR measurements with a pulsed field gradient. *Journal of Chemical Physics*, 69(4):1748–1754, January 1978.

[149] J. B. Tenenbaum, V. de Silva, and J. C. Langford. A global geometric framework for nonlinear dimensionality reduction. *Science*, 290(5500):2319–2323, December 2000.

[150] David S. Tuch. Q-ball imaging. *Magnetic Resonance in Medicine*, 52(6):1358–1372, 2004.

[151] F.B. Viégas, M. Wattenberg, M. McKeon, F. Van Ham, and J. Kriss. Harry Potter and the meat-filled freezer: A case study of spontaneous usage of visualization tools. In *Proc. HICSS*, 2008.

[152] F.B. Viegas, M. Wattenberg, F. Van Ham, J. Kriss, and M. McKeon. ManyEyes: a site for visualization at internet scale. *IEEE Transactions on Visualization and Computer Graphics*, 13(6):1121, 2007.

[153] 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 1. *Radiology*, 230(1):77–87, 2004.

[154] M. P. Wand. Data-based choice of histogram bin width. *The American Statistician*, 51(1):59–64, February 1997.

[155] Yanshu Wang, Jiangyang Zhang, Susumu Mori, and Jeremy Nathans. Axonal growth and guidance defects in frizzled3 knock-out mice:a comparison of diffusion tensor magnetic resonance imaging, neurofilament staining, and genetically directed cell labeling. *J. Neurosci.*, 26(2):355–364, 2006.

[156] Colin Ware. Color sequences for univariate maps: Theory, experiments and principles. *IEEE CG&A*, 8(5):41–49, 1988.

[157] Van J. Wedeen, Patric Hagmann, Wen-Yih I. Tseng, Timothy G. Reese, and Robert M. Weisskoff. Mapping complex tissue architecture with diffusion spectrum magnetic resonance imaging. *Magnetic Resonance in Medicine*, 54(6):1377–1386, 2005.

[158] D. J. Werring, D. Brassat, A. G. Droogan, C. A. Clark, M. R. Symms, G. J. Barker, D. G. MacManus, A. J. Thompson, and D. H. Miller. The pathogenesis of lesions and normal-appearing white matter changes in multiple sclerosis: a serial diffusion mri study. *Brain*, 123 (Pt 8):1667–1676, Aug 2000.

[159] World Health Organization. *Neuralogical Disorders–Public Health Challenges*, 2006.

[160] Boris Yamrom, John A. Sutliff, and Andrew P. Woodfield. Visualizing polycrystalline orientation microstructures with spherical color maps. In *VIS '94: Proceedings of the conference on Visualization '94*, pages 46–51, Los Alamitos, CA, USA, 1994. IEEE Computer Society Press.

[161] T. Yates, M.J. Okoniewski, and C.J. Miller. X: Map: annotation and visualization of genome structure for Affymetrix exon array analysis. *Nucleic Acids Research*, 36(Database issue):D780, 2008.

[162] Hui Zhang, Suyash P. Awate, Sandhitsu R. Das, John H. Woo, Elias R. Melhem, James C. Gee, and Paul A. Yushkevich. A tract-specific framework for white matter morphometry combining macroscopic and microscopic tract features. In *Proceedings of MICCAI*, pages 141–149, Berlin, Heidelberg, 2009. Springer-Verlag.

[163] Hui Zhang, Torben Schneider, Claudia A. Wheeler-Kingshott, and Daniel C. Alexander. Noddi: Practical in vivo neurite orientation dispersion and density imaging of the human brain. *NeuroImage*, 61(4):1000 – 1016, 2012.

[164] J. Zhang, L. J. Richards, M. I. Miller, P. Yarowsky, P. van Zijl, and S. Mori. Characterization of mouse brain and its development using diffusion tensor imaging and computational techniques. In *Proceedings of IEEE EMBC*, pages 2252–2255, 2006.

[165] Jiangyang Zhang, Ying-Bei Chen, Marie J. Hardwick, Michael I. Miller, Celine Plachez, Linda J. Richards, Paul Yarowsky, Peter van Zijl, and Susumu Mori. Magnetic resonance diffusion tensor microimaging reveals a role for bcl-x in brain development and homeostasis. *J. Neurosci.*, 25(8):1881–1888, February 2005.

[166] Song Zhang, Çağatay Demiralp, and David H. Laidlaw. Visualizing diffusion tensor MR images using streamtubes and streamsurfaces. *IEEE TVCG*, 9(4):454–462, 2003.

[167] Song Zhang, Çağatay Demiralp, and David H. Laidlaw. Visualizing diffusion tensor MR images using streamtubes and streamsurfaces. *IEEE TVCG*, 9(4):454–462, 2003.

[168] A. Zomorodian and G. Carlsson. Computing persistent homology. *Discrete & Computational Geometry*, 33(2):249–274, 2005.

[169] `http://graphics.cs.brown.edu/research/sciviz/newbraininteraction/`.