Brain Mapping Methods: Segmentation, Registration, and Connectivity Analysis

A dissertation submitted in partial satisfaction

of the requirements for the degree

Doctor of Philosophy in Computer Science

by

Gautam Prasad

2013

ABSTRACT OF THE DISSERTATION

Brain Mapping Methods: Segmentation, Registration, and Connectivity Analysis

by

Gautam Prasad

Doctor of Philosophy in Computer Science

University of California, Los Angeles, 2013

Professor Demetri Terzopoulos, Chair

We present a collection of methods that model and interpret information represented in structural magnetic resonance imaging (MRI) and diffusion MRI images of the living human brain. Our solution to the problem of brain segmentation in structural MRI combines artificial life and deformable models to develop a customizable plan for segmentation realized as cooperative deformable organisms. We also present work to represent and register white matter pathways as described in diffusion MRI. Our method represents these pathways as maximum density paths (MDPs), which compactly represent information and are compared using shape based registration for population studies. In addition, we present a group of methods focused on connectivity in the brain. These include an optimization for a global probabilistic tractography algorithm that computes fibers representing connectivity pathways in tissue, a novel maximum-flow based measure of connectivity, a classification framework identifying Alzheimer's disease based on connectivity measures, and a statistical framework to find the optimal partition of the brain for connectivity analysis. These methods seek to advance our understanding and analysis of neuroimaging data from crucial pre-processing steps to our fundamental understanding of connectivity in the brain.

The dissertation of Gautam Prasad is approved.

Alan L. Yuille

Stanley Osher

Paul M. Thompson

Demetri Terzopoulos, Committee Chair

University of California, Los Angeles

2013

Table of Contents

LIST OF FIGURES

LIST OF TABLES

me with the idea that there is a life outside of academia. I'm also thankful for Nina's love, support, and guidance and concede that a fifteen year PhD may not be practical in the long run.

<h1 style="text-align:center">Vita</h1>

June, 2007 B.S. (Computer Science), UC Davis, Davis, California.

June, 2010 M.S. (Computer Science), UCLA, Los Angeles, California.

March–June, 2011 Teaching Assistant, Computer Science Department, UCLA. Teaching Computer Science 35L Software Construction Laboratory.

June–Sept., 2011 Software Development Engineer Intern, Amazon.com, Seattle, Washington. Developed and implemented algorithms for forecasting product demand.

Sept., 2007–present Graduate Student Researcher, Laboratory of Neuro Imaging, UCLA School of Medicine, Los Angeles, CA. Developed algorithms and applications for brain classification, segmentation, and analysis using magnetic resonance (MR) and diffusion imaging images.

<h1 style="text-align:center">Publications</h1>

G. Prasad, S.H. Joshi, T.M. Nir, A.W. Toga, P.M. Thompson, *Machine Learning for Connectivity-based Alzheimer's Disease Classification*, Organization for Human Brain Mapping (OHBM) Meeting, 2013.

G. Prasad, S.H. Joshi, T.M. Nir, A.W. Toga, P.M. Thompson, *Machine Learning for Connectivity-based Alzheimer's Disease Classification*, Organization for Human Brain Mapping (OHBM) Meeting, 2013.

G. Prasad, S.H. Joshi, T.M. Nir, A.W. Toga, and P.M. Thompson, *Brain Connectivity based on Maximum Flow in Alzheimer's Disease: The EMFATIC Method*, Organization for Human Brain Mapping (OHBM) Meeting, 2013.

G. Prasad, S.H. Joshi, T.M. Nir, A.W. Toga, and P.M. Thompson, *Refining Brain Connectivity Networks to Optimally Identify Brain Disease*, Organization for Human Brain Mapping (OHBM) Meeting, 2013.

G. Prasad, T.M. Nir, A.W. Toga, and P.M. Thompson, *Fiber Density and Connectivity in Alzheimer's Disease*, Organization for Human Brain Mapping (OHBM) Meeting, 2013.

G. Prasad, S.H. Joshi, T.M. Nir, A.W. Toga, P.M. Thompson, *Flow-Based Network Measures of Brain Connectivity in Alzheimer's Disease*, IEEE International Symposium on Biomedical Imaging (ISBI), 2013.

G. Prasad, T.M. Nir, A.W. Toga, P.M. Thompson, *Tractography Density and Network Measures in Alzheimer's Disease*, IEEE International Symposium on Biomedical Imaging (ISBI), 2013.

T.M. Nir, **G. Prasad**, S.H. Joshi, J. Villalon, N. Jahanshad, A.W. Toga, M.A. Bernstein, B.J. Borowski, C.R. Jack, M.W. Weiner, P.M. Thompson, and the Alzheimer's Disease Neuroimaging Initiative (ADNI), *Predicting Future Brain Atrophy from DTI-based Maximum Density Path Analysis in Mild Cognitive Impairment and Alzheimer's Disease*, Medical Image Computing and Computer-Assisted Intervention (MICCAI) Workshop on Novel Imaging Biomarkers for Alzheimer's Disease and Related Disorders, 2012.

J. Villalon, **G. Prasad**, S.H. Joshi, M. Jalbrzikowski, A.W. Toga, C.E. Bearden, P.M. Thompson, *Statistical Analysis of Maximum Density Path Deformation Fields in White*

Matter Tracts, Medical Image Computing and Computer-Assisted Intervention (MICCAI) Workshop on Computational Diffusion MRI, 2012.

N. Jahanshad, **G. Prasad**, A.W. Toga, K.L. McMahon, G.I. de Zubicaray, N.G. Martin, M.J. Wright, P.M. Thompson, *Genetics of Path Lengths in Brain Connectivity Networks: HARDI-Based Maps in 457 Adults*, Medical Image Computing and Computer-Assisted Intervention (MICCAI) Workshop on Multimodal Brain Image Analysis, 2012.

G. Prasad, S. Joshi, N. Jahanshad, J. Villalon, G.I. de Zubicaray, K.L. McMahon, N.G. Martin, M.J. Wright, I. Aganj, G. Sapiro, A.W. Toga, P.M. Thompson, *Genetic Analysis of Fibers in White Matter Pathways from HARDI Images*, Organization for Human Brain Mapping (OHBM) Meeting, 2012.

G. Prasad, S. Joshi, N. Jahanshad, A.W. Toga, P.M. Thompson, *White Matter Tract Analysis in 367 Adults using Fiber Clustering, Maximum Density Paths, and Curve Registration*, Medical Image Computing and Computer-Assisted Intervention (MICCAI) Workshop on Computational Diffusion MRI, 2011.

G. Prasad, A.A. Joshi, A.W. Toga, D. Terzopoulos, P.M. Thompson, *Brain Segmentation using Deformable Organisms and Error Learning*, Medical Image Computing and Computer-Assisted Intervention (MICCAI) Workshop on Mathematical Foundations of Computational Anatomy, 2011.

J. Villalon, **G. Prasad**, S. Joshi, N. Jahanshad, A. Quintero, A.W. Toga, P.M. Thompson, T. Simon, *White Matter Tract Abnormalities in Fragile-X and 22q11.2 Deletion Syndromes*, Proc. of the Society for Neuroscience (SFN), 2011.

G. Prasad, N. Jahanshad, I. Aganj, C. Lenglet, G. Sapiro, A.W. Toga, P.M. Thomp-

son, *Atlas-Based Fiber Clustering for Multi-Subject HARDI Tractography*, Organization for Human Brain Mapping (OHBM) Meeting, 2011.

G. Prasad, N. Jahanshad, I. Aganj, C. Lenglet, G. Sapiro, A.W. Toga, P.M. Thompson, *Atlas-based Fiber Clustering for Multi-subject Analysis of High Angular Resolution Diffusion Imaging Tractography*, IEEE International Symposium on Biomedical Imaging (ISBI), 2011.

G. Prasad, A.J. Joshi, P.M. Thompson, A.W. Toga, D.W. Shattuck, D. Terzopoulos, *Skull-stripping With Deformable Organisms*, IEEE International Symposium on Biomedical Imaging (ISBI), 2011.

D.W. Shattuck , **G. Prasad**, M. Mirza, K.L. Narr, A.W. Toga, *Online Resource for Validation of Brain Segmentation Algorithms*, NeuroImage, 45(2):431-439, April 2009.

K.R. Beutner III, **G. Prasad**, E. Fletcher, C. DeCarli, O.T. Carmichael, *Estimating Uncertainty in Brain Region Delineations*, Proc. of the Information Processing in Medical Imaging (IPMI), 2009.

G. Prasad, O.T. Carmichael, *Automated Brain Region Delineation in Structural Magnetic Resonance Images*, UC Davis Undergraduate Research Conference, 2007.

O.T. Carmichael, **G. Prasad**, *Robust, Fully-Automated Delineation of Brain Regions in Structural MR*, Biomedical Image Segmentation Conference, UC Davis Cancer Research Center, 2006.

CHAPTER 1

Introduction

High resolution structural magnetic resonance imaging (MRI) and diffusion weighted imaging (DWI) images provide an unprecedented view of the human brain *in vivo*. Structural MRI images show a detailed grayscale three-dimensional picture of tissue organization, containing tens of millions of voxels. DWI builds on this information by capturing the diffusion of water molecules in tissue. It samples the diffusion strength in multiple directions at each point in the brain to generate a profile of diffusion. This information furnishes a macroscopic view of the brain's tissue, which restricts diffusion because of interactions with axons, cell membranes and vascular structures. In subjects, these images along with the genome sequence, cognitive tests, age, and sex comprise a detailed picture of the human condition for the study of disease, genetic effects, developmental changes, and aging.

Researchers have designed methods that capture and interpret medical images using tools from computer science, statistics, and mathematics. They used an array of techniques (Klein et al., 2009) to register MRI images into the same space such as the discrete cosine transform with Levenberg-Marquardt optimization (Andersson et al., 2008) or a Eulerian velocity framework with a multigrid method (Ashburner, 2007). Others have tackled brain segmentation with a trained Markov random field model (Fischl et al., 2002), a machine learning feature based discriminative model combined with a generative model (Tu et al., 2008), and recently a geodesic curvature flow on the cortical surface using level sets (Joshi et al., 2012). In DWI, scientists first captured the diffusion information using a tensor (Basser et al., 1994) and then introduced the orientation distribution function (ODF) (Aganj et al., 2010; Tuch, 2004), a parametrization rich with information designed for high angular

resolution diffusion imaging (HARDI) (Tuch et al., 2002). The knowledge represented by these models was used by investigators to generate a representative set of fibers for pathways of white matter tissue in the brain, a process referred to as *tractography*. Researches initially defined tractography as fibers following the direction of highest local diffusion (Basser et al., 2000), a method referred to as *streamlines*, but more recently introduced methods such as a global probabilistic method (Aganj et al., 2011) that searches all possible fibers to find those that optimize a score function. The fibers from tractography give a way to quantify connectivity between regions and there has been a plethora of research on how to represent (Zalesky et al., 2011) and understand connectivity (Bullmore and Sporns, 2009), this mapping deemed the *connectome* (Sporns et al., 2005). In addition to standard image processing problems, scientists have created disease classification methods of image features based on support vector machines (Cuingnet et al., 2011) or nearest neighbor classification with a self-smoothing operator (Iglesias et al., 2011a). The availability of quick and affordable genome sequencing give researches the ability to perform genome-wide association studies (GWAS) (Stein et al., 2012) in tens of thousands of subjects to discover relationships between our brain structure and our genetic code.

In this dissertation, we present a collection of methods to understand the structure of the human brain and compare this information across a group of subjects. Our first method (Prasad et al., 2011c,b) addresses brain segmentation with a deformable organism framework paired with a machine learning error identification algorithm. We then present a method (Prasad et al., 2011a,d, 2012) that compactly represents white matter tissue in the brain and uses this representation to understand differences in a cohort of subjects. The latter part of this dissertation presents methods that capture and define connectivity in the brain: an optimized tractography algorithm (Prasad et al., 2013f,g), a flow based connectivity metric (Prasad et al., 2013a,b), a connectivity based classification function (Prasad et al., 2013c), and a framework to define the architecture of connectivity (Prasad et al., 2013d,e).

1.1 Brain Segmentation

Brain segmentation involves delineating the brain in full head MRI images, the process removes extraneous material such as the skin, cerebrospinal fluid (CSF), eyes, and nose from the image and is referred to as *skull-stripping*. The segmented brain is an important pre-processing step in most neuroimaging studies, necessary for registration algorithms (Ardekani et al., 2005), cortical analysis (Fischl et al., 1999), segmentation methods (Morra et al., 2008), and inhomogeneity correction (Zhang et al., 2001). The resulting masks are also applied to DWI images to focus subsequent analysis on exclusively brain tissue. Skull-stripping was traditionally done manually by trained experts who went through each slice and traced the boundary of the brain, a process that is prone to error and has difficulties with repeatability. In addition, manual tracing can take up to one week per subject, making it intractable for large studies with thousands of images. There have been a wide array of proposed automatic skull-stripping methods that can quickly and consistently identify the brain. Some of the most widely used methods use a deformable model expanding from the center of the image (Smith, 2002), the watershed algorithm in combination with a template (Segonne et al., 2004), and edge detection combined with morphological operations (Shattuck and Leahy, 2002). However, depending on the input data, these algorithms may be deficient and require manual clean-up, a process taking hours per subject. We proposed a segmentation method (Prasad et al., 2011c,b) that allows the user to specify a high level plan unique to their data. Our algorithm draws from the principals of artificial life (Langton and Shimohara, 1997) and deformable models (McInerney and Terzopoulos, 1996) to create *deformable organisms* that cooperatively work together to segment the brain.

Segmenting brain from non-brain tissue within magnetic resonance (MR) images of the human head, also known as skull-stripping, is a critical processing step in the analysis of neuroimaging data. Though many algorithms have been developed to address this problem, challenges remain. In this paper, we apply the "deformable organism" framework to the skull-stripping problem. Within this framework, deformable models are equipped with higher-

level control mechanisms based on the principles of artificial life, including sensing, reactive behavior, knowledge representation, and proactive planning. Our new deformable organisms are governed by a high-level plan aimed at the fully-automated segmentation of various parts of the head in MR imagery, and they are able to cooperate in computing a robust and accurate segmentation. We applied our segmentation approach to a test set of human MRI data using manual delineations of the data as a reference "gold standard." We compare these results with results from three widely used methods using set-similarity metrics.

Segmentation methods for medical images may not generalize well to different data sets or tasks, hampering their utility. We attempt to remedy these issues using deformable organisms to create an easily customizable segmentation plan. This plan is developed by borrowing ideas from artificial life to govern a set of deformable models that use control processes such as sensing, proactive planning, reactive behavior, and knowledge representation to segment an image. The image may have landmarks and features specific to that dataset; these may be easily incorporated into the plan. We validate this framework by creating a plan to locate the brain in 3D magnetic resonance images of the head (skull-stripping). This is important for surgical planning, understanding how diseases affect the brain, conducting longitudinal studies, registering brain data, and creating cortical surface models. Our plan dictates how deformable organisms find features in head images and cooperatively work to segment the brain. In addition, we use a method based on Adaboost to learn and correct errors in our segmentation. We tested our method on 630 T1-weighted images from healthy young adults, evaluating results using distance and overlap error metrics based on expert gold standard segmentations. We compare our segmentations with and without the error correction step; we also compare our results to three other widely used methods: BSE, BET, and the Hybrid Watershed algorithm. Our method had the least Hausdorff distance to expert segmentations on this dataset, but included slightly more non-brain voxels (false positives). Our framework captures diverse categories of information needed for skull-stripping, and produces competitive segmentations.

1.2 White Matter Representation & Registration

The white matter tissue described by DWI is vital to understanding connectivity in the brain (Jones et al., 1999), but researchers are still in need of an effective way to represent and compare the tissue across a population of subjects. The composition, shape, and differences in white matter allow scientists to find changes in connectivity depending on sex (Shin et al., 2005), age (Andrews-Hanna et al., 2007), or disease (Rose et al., 2000) and have been frequently analyzed by comparing a scalar measure of connectivity (fractional anisotropy; FA) (Basser and Pierpaoli, 1996) represented as volumetric images (Holzapfel et al., 2006) or skeletons (Smith et al., 2006) of these images. However, nonlinear volumetric registration and the metric used may have difficulties in identifying corresponding structures (Tustison et al., 2012). In addition to FA, the fibers from tractography represent pathways connecting different areas of the brain, but the set of fibers sometimes numbering greater than one million per subject and may be difficult to bring into correspondence (Zöllei et al., 2010), even for a single white matter bundle (Corouge et al., 2006). We introduce our own representation called the maximum density path (MDP) (Prasad et al., 2011a,d, 2012), which compactly represents the fibers in a white matter regions by representing their density as a path. Our method brings these paths into correspondence for a cohort of subjects using geodesic curve registration, enabling population-wise analysis. The resulting registration based on the shape of the white matter pathways avoids many of pitfalls of volumetric registration.

High angular resolution diffusion imaging (HARDI) allows *in vivo* analysis of the white matter structure and connectivity. Based on orientation distribution functions (ODFs) that represent the directionality of water diffusion at each point in the brain, tractography methods can recover major axonal pathways. This enables tract-based analysis of fiber integrity and connectivity. For multi-subject comparisons, fibers may be clustered into bundles that are consistently found across subjects. To do this, we scanned 20 young adults with HARDI at 4 T. From the reconstructed ODFs, we performed whole-brain tractography with a novel Hough transform method. We then used measures of agreement between the extracted 3D

curves and a co-registered probabilistic DTI atlas to select key pathways. Using median filtering and a shortest path graph search, we derived the maximum density path to compactly represent each tract in the population. With this tract-based method, we performed tract-based analysis of fractional anisotropy, and assessed how the chosen tractography algorithm influenced the results. The resulting method may expedite population-based statistical analysis of HARDI and DTI.

We introduce a framework for population analysis of white matter tracts in diffusion weighted images. Our pipeline computes fibers from high angular resolution diffusion imaging (HARDI) images; it clusters the fibers incorporating prior knowledge from an atlas; the method then represents each fiber bundle compactly using a path following points of highest density (maximum density path; MDP); and registers these paths together using geodesic curve registration to find local correspondences across a population. Our framework is tested with 454 subjects images using 4-Tesla HARDI to find localized statistics across 50 white matter tracts based on fractional anisotropy (FA) and sex differences. In addition, we compared the results with those derived with streamline tractography.

1.3   Connectivity

The connectivity between areas of the cortex is often quantified as their intersection with tractography fibers, with these measures commonly organized into a connectivity matrix (Hagmann et al., 2007). The connectivity matrix is a type of adjacency matrix where each element represents the number of fibers connecting two regions. These matrices have been used to understand differences in connectivity depending on hemisphere (Skudlarski et al., 2008), sex (Jahanshad et al., 2011), age (Fair et al., 2007), disease (Stam et al., 2007), heritability (Schmitt et al., 2008), and our genes (Dennis et al., 2011). Conventionally, researches have relied on streamline tractography to generate fibers for connectivity analysis, but streamlines may compound noise in DWI images (Tuch et al., 2000). In response to these issues our work focused on generating fibers using a global probabilistic tractography

algorithm (Aganj et al., 2011) based on the Hough transform, that optimized a measure of FA along with a probability based on diffusion. Due to the added complexity and lengthy running time of the Hough transform method, we introduced an optimization (Prasad et al., 2013f,g) based on an ODF lookup table and random traversal of the parameter space. In addition, we explored the relationship of fiber density or count on measures of connectivity between healthy elderly controls versus Alzheimer's disease patients. Our optimization of the tractography make it practical for researches to use as an alternative to streamlines and our work stresses the importance that additional fibers may have to connectivity analysis.

We expanded on the traditional fiber count based measure by introduceding a flow based measure of connectivity between regions (Prasad et al., 2013a,b). Our method modeled the DWI image space as a graph (a set of nodes and edges), where the edge weight was the diffusion strength and represented the capacity between two nodes. We ran a modified maximum flow algorithm that incorporated the biologically viable pathways from tractography between two nodes to quantify the connectivity and compared it with the standard fiber count method to differentiate Alzheimer's disease. Our name for the method was evaluating maximum flow across tracts indicating connectivity (EMFATIC).

Scientists have developed many graph measures (Rubinov and Sporns, 2010) derived from connectivity matrices in hopes to provide scalar measures that assess and describe networks. These measures are often used to understand how the connectome reflects the effects of disease, aging, or difference in hemisphere. We use these measures as features in a machine learning framework based on support vector machines to classify Alzheimer's disease (Prasad et al., 2013c). Our framework computes hundreds of thousands of these features, reduces their dimension using principal components analysis, and feeds them into a 10-fold cross-validated design to evaluate the information contained in the features based on accuracy, specificity, and sensitivity. We found that classification based on connectivity is a useful addition to current frameworks for identifying Alzheimer's disease that rely on MRI measures of volume, CSF biomarkers, ApoE genotype, age, sex, body mass index, and diagnostic tests (Kohannim et al., 2010; Klöppel et al., 2008).

In most cases the regions of the cortex or nodes in the connectivity matrix are based on an automatic segmentation into functional areas (Bullmore and Sporns, 2009). A few researchers have explored how the size, shape, and location of these nodes affects the connectivity matrix (Tzourio-Mazoyer et al., 2002; Wig et al., 2011). We address this issue by introducing a Bayesian statistical framework that finds the optimal segmentation of the cortex so that the resulting connectivity matrix optimizes an objective function (Prasad et al., 2013d,e), a method we brand evolving partitions to improve connectomics (EPIC). We accomplish this by posing the problem as a partition of a set of small random patches on the cortex and sample it based on probability dictated by an objective function using markov chain monte carlo (MCMC). In our experiments we defined our objective function as the accuracy of a classification of Alzheimer's disease based on graph measures derived from the connectivity matrix. We compared our method with a matrix following the standard function segmentation of the cortex and a matrix based on the small random patches. Our method allows researchers to choose the nodes or architecture of a connectivity matrix to best capture connectivity information for any purpose whether that be disease, aging, or genetics.

Brain connectivity declines in Alzheimer's disease (AD), both functionally and structurally. Connectivity maps and networks derived from diffusion-based tractography offer new ways to track disease progression and to understand how AD affects the brain. Here we set out to identify (1) which fiber network measures show greatest differences between AD patients and controls, and (2) how these effects depend on the density of fibers extracted by the tractography algorithm. We computed brain networks from diffusion-weighted images (DWI) of the brain, in 110 subjects (28 normal elderly, 56 with early and 11 with late mild cognitive impairment, and 15 with AD). We derived connectivity matrices and network topology measures, for each subject, from whole-brain tractography and cortical parcellations. We used an ODF lookup table to speed up fiber extraction, and to exploit the full information in the orientation distribution function (ODF). This made it feasible to compute high density connectivity maps. We used accelerated tractography to compute

a large number of fibers to understand what effect fiber density has on network measures and in distinguishing different disease groups in our data. We focused on global efficiency, transitivity, path length, mean degree, density, modularity, small world, and assortativity measures computed from weighted and binary undirected connectivity matrices. Of all these measures, the mean nodal degree best distinguished diagnostic groups. High-density fiber matrices were most helpful for picking up the more subtle clinical differences, e.g. between MCI and normals, or for distinguishing subtypes of MCI (early versus late). Care is needed in clinical analyses of brain connectivity, as the density of extracted fibers may affect how well a network measure can pick up differences between patients and controls.

We present a new flow-based method for modeling brain structural connectivity. The method uses a modified maximum-flow algorithm that is robust to noise in the diffusion data and guided by biologically viable pathways and structure of the brain. A flow network is first created using a lattice graph by connecting all lattice points (voxel centers) to all their neighbors by edges. Edge weights are based on the orientation distribution function (ODF) value in the direction of the edge. The maximum-flow is computed based on this flow graph using the flow or the capacity between each region of interest (ROI) pair by following the connected tractography fibers projected onto the flow graph edges. Network measures such as global efficiency, transitivity, path length, mean degree, density, modularity, small world, and assortativity are computed from the flow connectivity matrix.

We show the benefits and differences of using different types of connectivity features to distinguish disease states in Alzheimer's disease. Our features come from connectivity matrices constructed from tractography fiber density between regions and using a flow based connectivity method. In addition, we compute 28 network measures from the connectivity matrices. We reduce the dimensionality of these features using principal components analysis (PCA) and incorporate them into a 10-fold cross-validation using support vector machines (SVMs), a supervised learning algorithm. Measures of accuracy, sensitivity, and specificity show the differences between feature sets in discriminating between normal healthy controls, early- and late-stage mild cognitive impairment, and Alzheimer's disease in 110 subjects.

We propose a method for choosing nodes in a connectivity matrix of the brain that optimizes feature-based classification performance for disease. Given N cortical parcellations, the network measures are stored in a $N \times N$ matrix that represents the pairwise connectivity between those regions. We formulate the problem by sampling over the space of possible partitions of a random set of cortical patches. These partitions are assigned a value using a support vector machine (SVM) classifier of disease, based on graph measures derived from the connectivity matrix specified for a group of subjects. We characterize the resulting high dimensional optimization problem using Markov Chain Monte Carlo (MCMC), which provides information on the optimal partition for classification accuracy and a sample variance to understand the distribution. Ultimately, the method yields a quantitative connectivity measurement to determine the importance of selecting the optimal network architecture for classification. We demonstrate our method on the ADNI-2 dataset, where we obtain an optimal parcellation of the cortex corresponding to an 80% classification accuracy using the network measures alone.

CHAPTER 2

Brain Segmentation

2.1 Skull-Stripping with Deformable Organisms

2.1.1 Introduction

Skull-stripping is the process of segmenting brain from non-brain tissues (e.g., skull, scalp, eyes, or neck) in whole-head magnetic resonance (MR) images. Delineating the brain region is important for applications such as surgical planning (Gering et al., 2001); analysis of brain images, where removing non-brain structures allows all subsequent analysis to focus on the brain voxels specifically; or in brain registration (Woods et al., 1999; Klein et al., 2009). Skull-stripping can also perform an important role in extracting cortical surface models (Dale et al., 1999), in analyzing how the brain changes over time in longitudinal studies (Resnick et al., 2003), and in examining how the brain is affected by disease (Thompson et al., 2003).

Skull-stripping can be performed manually, and manually delineated brain masks are typically used as the "gold standard" to validate automatic brain extraction algorithms. Manual skull-stripping takes a substantial amount of time to complete for a MRI volume, and many groups have developed automated algorithms to address this problem. Still, users often resort to manual clean-up of automated skull-stripping results or will forgo the use of automated algorithms entirely. Achieving better reliability in automated skull-stripping algorithms remains an important research problem, as this would be less labor-intensive while also reducing rater variation across sets of images. Figure 2.1 shows a manual identification of the brain (*green*) in a T1 MR image of the head (*red*).
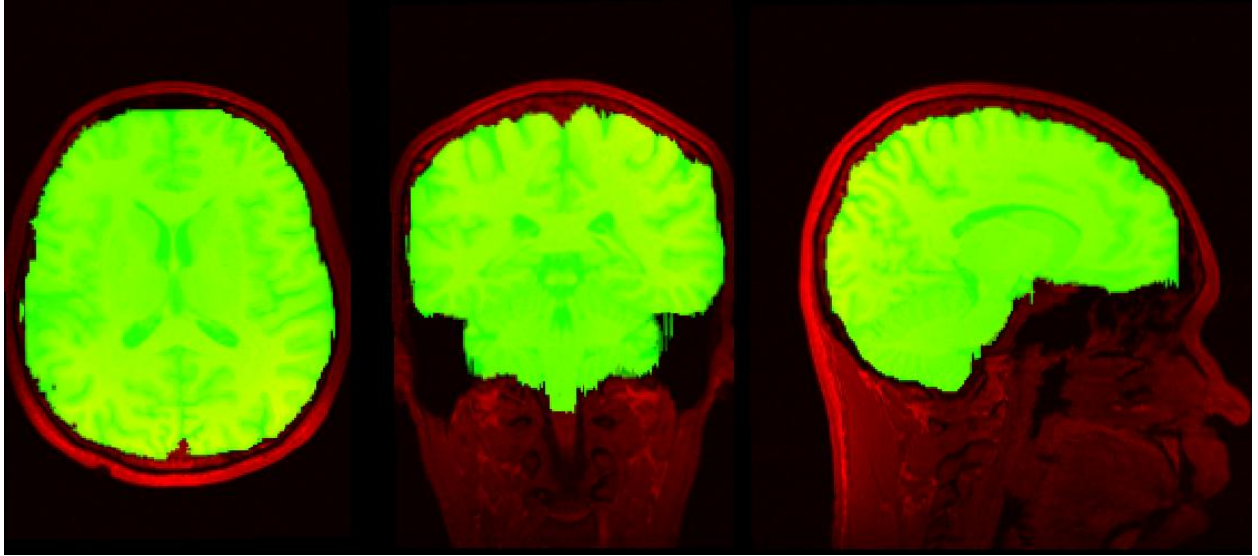
Figure 2.1: T1 MR image (*red*) with the brain location manually identified in *green*.

Many automated approaches have been described in the literature and made available publicly. The Brain Surface Extractor (BSE) applies Marr-Hildreth edge detection to the image to obtain a connected component that represents the brain; it then applies morphological operations to remove erroneous connected regions and to produce a smooth brain mask (Shattuck et al., 2001). In some cases, the erosion may fail to completely separate the brain from the surrounding tissues, which can lead to attached skull, scalp, or dura. The Brain Extraction Tool (BET) (Smith, 2002) uses a deformable model that evolves into the shape of the brain by relying on local information in the image. This method is more robust for images where certain areas do not have a clear boundary between the brain and non-brain tissue, but it does not incorporate anatomical tissue regions and their relative positions. The Hybrid Watershed Algorithm (HWA) (Segonne et al., 2004), applies the watershed algorithm to the MR image to obtain a rough estimate of the brain region, then fits a deformable surface to the region, and allows it to deform based on geometric constraints and a statistical atlas. Comparisons of these and other skull-stripping algorithms have appeared in (Fennema-Notestine et al., 2006).

In this paper, we introduce the application of "deformable organisms" (McInerney et al.,

2002) to the problem of skull-stripping[1]. Within the deformable organisms framework, deformable models are equipped with higher-level control mechanisms based on the principles of artificial life. These mechanisms include sensing, reactive behavior, knowledge representation, and proactive planning. Our new deformable organisms are governed by a high-level plan aimed at fully-automated segmentation of various parts of the head in MR images, and they cooperate in computing the segmentation. The organisms make use of local information, such as edges, along with global image tissue classifications from K-means clustering. Their intermediate goals of finding easily recognized features in the image make the final segmentation more robust.

2.1.2 Methods

2.1.2.1 Image Processing

The subject T1 MR of the head is first processed so that its results are available to the deformable organisms as they segment the brain. The organisms are embedded in the image space and "sense" the processed images by analyzing the intensities around them.

The images are processed in various ways to make different types of information available to the sensors of the organisms. The basic T1 weighted MR image consists of a volume grayscale intensity image and is used as a base for three different types of processing.

- The gradient of the base MR image is computed to create an image that emphasizes the edges.

- A threshold is applied to the base image by finding a histogram of the intensities and processing it to classify voxels lying within the head from those of the surrounding air in the image.

- Once a threshold has been applied to the image, a K-Means classification is used to

---

[1](Bovenkamp et al., 2004) proposed a framework similar to deformable organisms for segmenting IntraVascular UltraSound (IVUS) images. Their work, however, focused on establishing a set of rules that multiple agents can use to communicate in order to solve particular medical image segmentation tasks.

classify the head voxels into K different classes. In our case we set K to 2 and 3 (Figure 2.2). The brain region is actually composed of two tissue types, but the rest of the head adds interference into the classification.
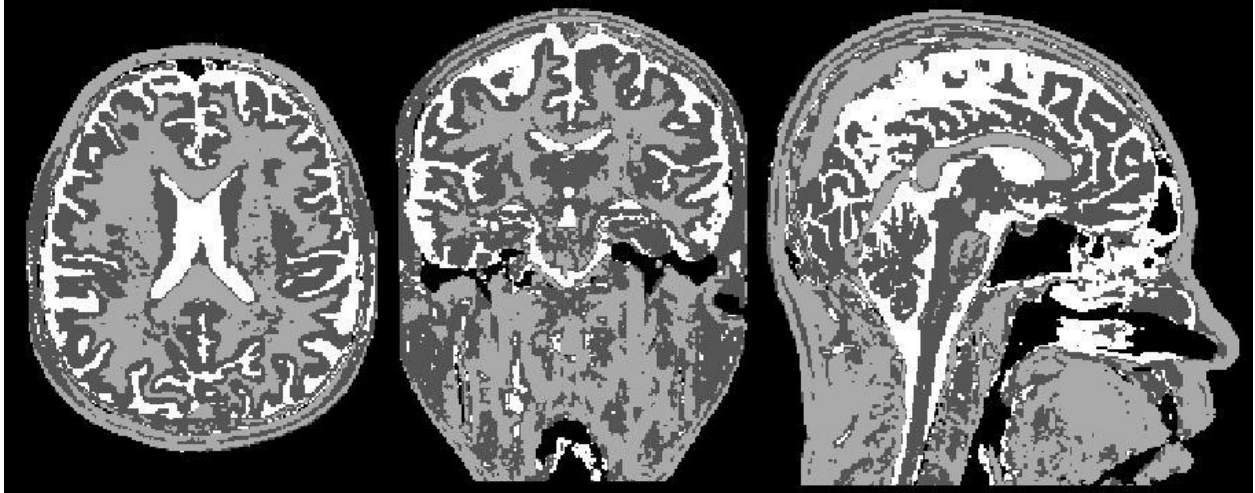
Figure 2.2: 3-Means classification of a T1 MR image. The intensities are mapped to three different labels, segmenting it into three areas.

2.1.2.2 Deformable Organisms

Our deformable organisms combine representations and control mechanisms of various different types, as outlined below. Each layer is customized to whatever needs to be segmented, in our case we focus on the brain.

Geometry and Physics Each deformable organism is represented geometrically as a 3D triangulated mesh. The models are initialized in the shapes of spheres that either contract or expand to find the boundary of the object being modeled. Each geometric model is deformed iteratively to model different structures in the MR images. Each vertex on the mesh is moved either inwards or outwards along the direction of the normal vector at that point. At each iteration, Laplacian smoothing is applied to the mesh to constrain the movement of each vertex in order to maintain a smooth mesh that does not pass through itself.

Perception The perception layer enables the organism to sense the medical image in which it is embedded. The vertices of the triangulated mesh are represented in real coordinates embedded in a volume image represented by a set of voxels. Hence, nearest neighbor interpolation was used to sample the image intensities at the location of a certain point in the mesh. The locations of all the organisms in the image are computed by the perception layer. A 3D rasterization method followed by dilation is used to calculate which voxels in the volume are located within each organism. The dilation is required to deal with the discrepancy between the low resolution of the volume image in which meshes are embedded. When the mesh is rasterized it will conservatively choose a voxel that is only partly intersected by the mesh and thus remove too many voxels from the boundary. The dilation helps to fit the border voxels better to the boundary of the mesh and resolve this discrepancy. This information is used by the motor control layer to regulate where an organism is allowed to deform by restricting organisms from intersecting each other.

2.1.2.3 Motor Control

The motor control of the organism is a function of intensities along the line normal to the mesh surface going through each vertex. The intensities along this line are from the images available to the perception layer. This layer looks for or avoids a particular intensity or relative intensity or tries to fit a certain model or statistic to the data along these lines. The intensities along the normal lines are sampled from different types of sensors and can each be processed using a different set of constraints.

Behavior The organism has a repertoire of behaviors. Translation is a behavior that moves a particular organism rigidly without any deformation to the mesh, as does the rotation behavior.

These behaviors can take into account the organism's relationship to other organisms and use information about their locations to decide how to move rigidly. Another behavior is the local deformation of the mesh. This behavior also depends on sensing different infor-

mation, various motor controls, and the locations of surrounding organisms. Thus high level information about locating different regions in the image can be used to create a plan for segmentation.

Cognition The cognitive layer is created by putting together a set of behaviors to accomplish certain goals. Different behaviors can be activated dynamically depending on what goals have been accomplished or what features have or need to be located.

2.1.2.4 Skull-Stripping Plan

Figure 2.3 shows the image processing steps and the organism plan. The *red* arrow points to the processing of the T1 MR image. The black arrows show the image and location dependencies at each step. The images are sensed at various stages in the organism plans. When a structure is found, its location can be helpful when locating other structures. For instance, the deformation of the brain utilizes the location of the skin and the eyes to figure out the bounds of where it can expand.

The Skull-Stripping plan dictates which organisms are used, the goals for each organism and how the organisms interact. It begins by finding the location of a skin organism that surrounds the head. Then using the skin organism's location and shape, two eye organisms are spawned. Finally, a brain organism is created that interacts with the three other organisms to refine its own location.

Figure 2.4 shows an example of the organism deformations and interactions during the skull-stripping process. The skin organism is initialized as a large spherical triangulated mesh that is deformed into the surface of the head using the threshold of the initial MR image (Fig. 2.4A and 2.4B). The skin organism is then processed to locate the nose. This information is used to locate the eyes (Fig. 2.4C and 2.4D). Figure 2.4D shows the two eye organisms expanded to the full size of the eye by sensing the 3-Means classification of the MR image. Once this is complete, the skin organism again deforms to locate the area surrounding the brain by deforming through the eyes by sensing their locations and by using
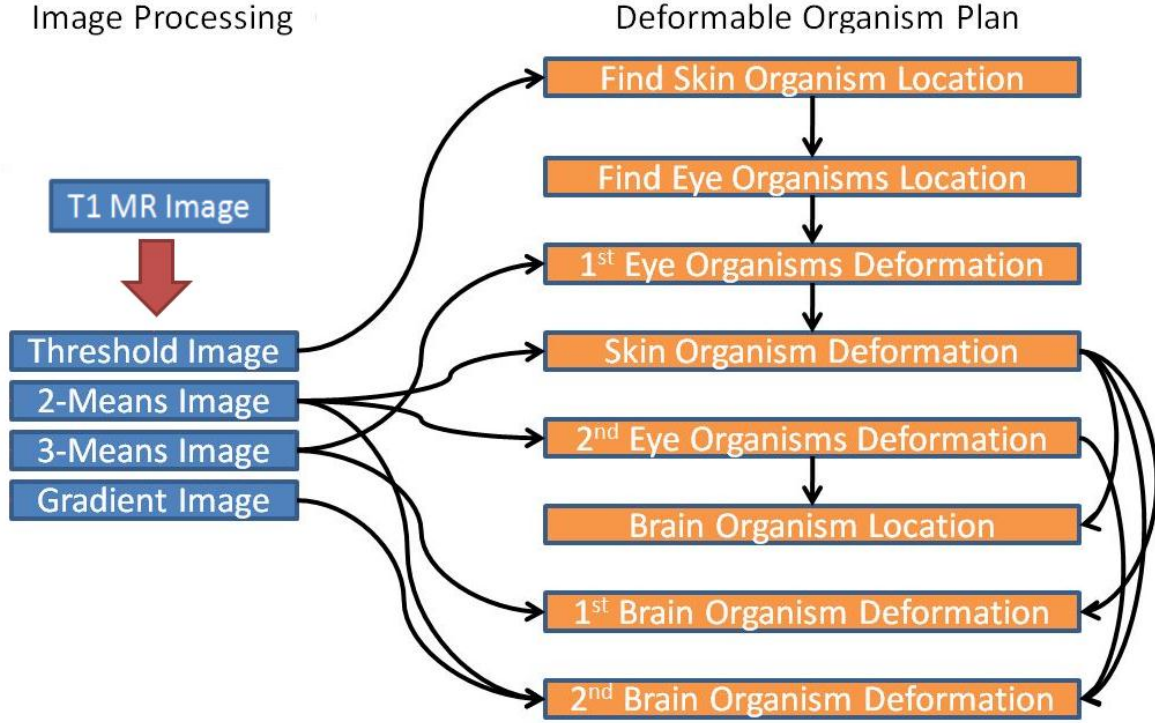
Figure 2.3: The flow of data and steps of the algorithm. The *red* arrow shows processing steps for the image and the black arrows represent information that is being passed.

edge information from the 2-Means classified image as shown in Fig. 2.4E. In Fig. 2.4F, the eyes are again deformed by sensing the 2-Means image to take into account the surrounding tissues and to restrict more areas that the brain organism may try to expand into. Then using the location of the eyes and skin meshes a brain organism is spawned (Fig. 2.4G) that deforms itself to match the classification of tissues in the 3-Means image and to stay within the skin mesh and complete the segmentation of the brain (Fig. 2.4H).

Figure 2.5 shows how the brain organism interacts with the skin (*yellow*) and eye (*red*) organisms. Figure 2.5A shows how the brain organism (*cyan*) is expanding by sensing the image forces from the 3-Means classification of the image. The *blue* arrows show the direction in which the brain organism is deforming. Figure 2.5B shows that the organism is restricted by the skin and eye organisms as it pushes outwards. It is able to sense the locations of the
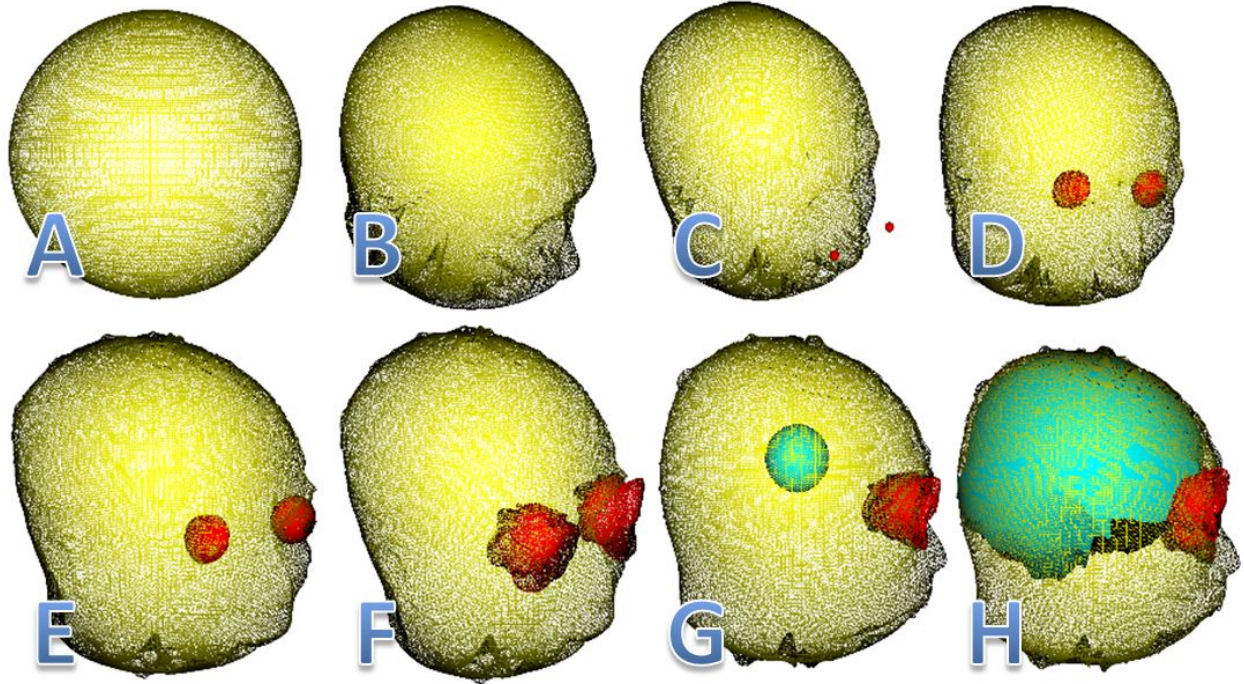
Figure 2.4: This figure shows the sequential steps that the skin (*yellow*), eye (*red*), and brain (*cyan*) organisms use to skull-strip the head image.

other organisms and is forced to stay within the boundary as shown by the red arrows.

2.1.2.5 Implementation

Operations and processing of the MR images was implemented using the Insight Toolkit (ITK)[2], an open-source C++ library for medical image analysis. The mesh operations and the visualization of the deformable organisms was completed using the Visual Toolkit (VTK)[3], an open-source C++ library for 3D graphics, visualization, and image processing.

The entire deformable organisms method for skull-stripping takes less than three minutes on an Intel 2GHz machine with 1.50 GB of RAM to segment an MR image.

---

[2]http://www.itk.org/
[3]http://www.vtk.org/
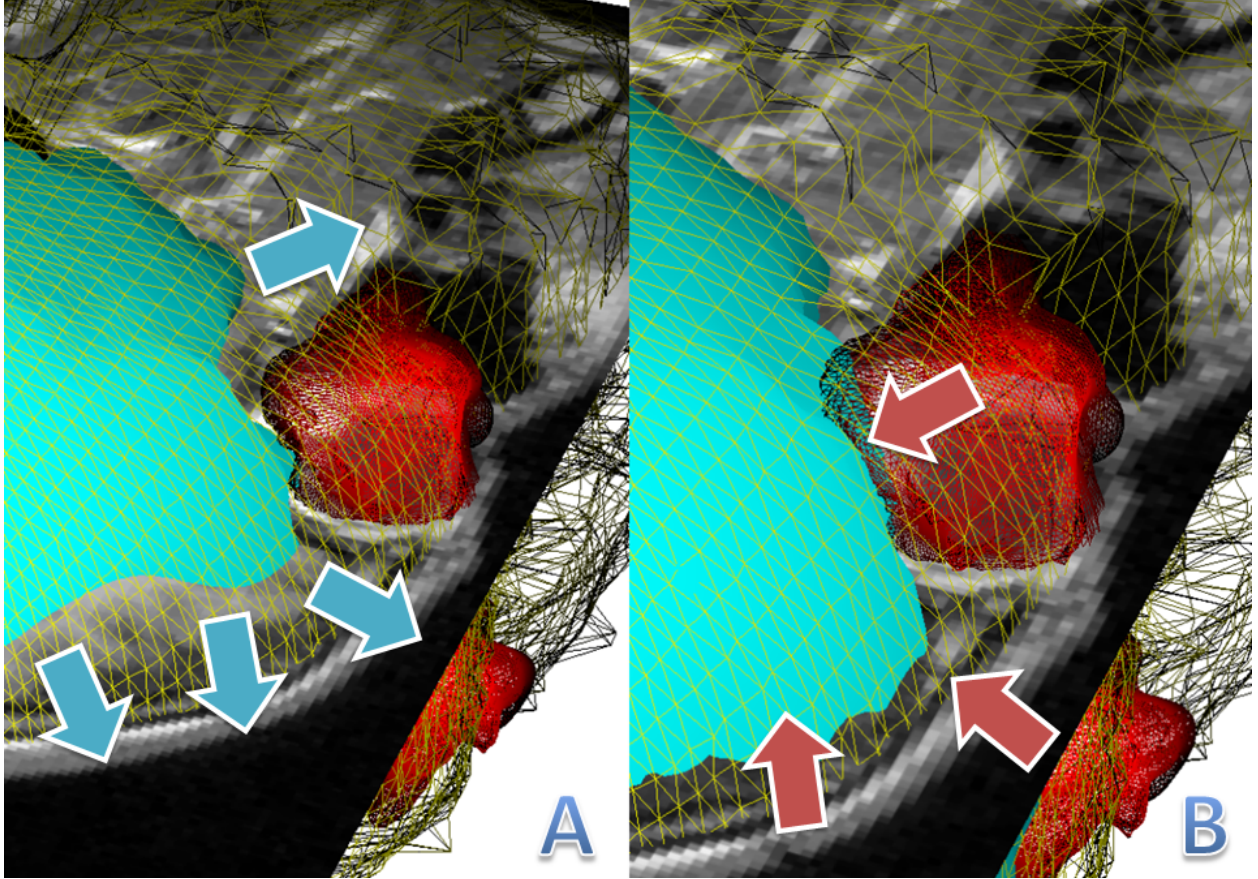
Figure 2.5: Interactions between the brain (*cyan*), eye (*red*), and skin (*yellow*) organisms. The arrows (*blue*) in Subfigure A show how the brain organism is expanding and Subfigure B has arrows (*red*) showing how its movement is restricted by the other organisms.

2.1.2.6 Evaluation

We evaluated the performance of our algorithm using the Segmentation Validation Engine (SVE; http://sve.loni.ucla.edu), a web-based resource that we developed previously (Shattuck et al., 2009). SVE provides a set of 40 human whole-head T1-weighted MR images of $256 \times 124 \times 256$ voxels (voxel size $0.86 \times 1.50 \times 0.86$ mm^3) that were delineated manually as part of the LONI Probabilistic Brain Atlas (LPBA40) (Shattuck et al., 2008). The 40 images were downloaded and processed with our algorithm. The brain masks were then uploaded to the SVE server, which computed a series of measures comparing our brain masks with the manually-delineated masks. The overlap metrics used were the Jaccard Similarity, the Dice Coefficient, Sensitivity, and Specificity; these are described in (Shattuck et al., 2009).

19

Validation results are archived, allowing comparison with previously evaluated methods.

2.1.3   Results

We applied the algorithm to a set of data, compared it to three existing algorithms, and validated the results.

2.1.3.1   Other Algorithms

The deformable organism framework for skull-stripping was compared with FSL's Brain Extraction Tool (BET) (Smith, 2002) (version: BETv2.1 settings: -B), the Brain Surface Extractor (BSE) (Shattuck et al., 2001) (version: BSEv08b settings: -n 5 -d 15 -s 0.65 -p –noneck), and the Hybrid Watershed (HWA) (Segonne et al., 2004) (version: HWA3 settings: -less).

Table 1 shows the mean for each of the metrics across the 40 images used to test the algorithms [mean±S.D.]. The results from the SVE false positive and false negative projection maps showed that the method includes too many voxels around the cerebellum lobe and the inferior and anterior portions of the temporal lobe. These can be improved through refining the parameters in the behavior layer. The method does do a good job at delineating the area around the medial longitudinal fissure (the groove separating the two hemispheres of the brain) and the anterior part of the parietal lobe when compared to the other algorithms.

Method	Jaccard	Dice	Sensitivity	Specificity
DO	0.8954±0.0288	0.9446±0.0163	0.9616±0.0129	0.9864±0.0082
BSE	0.9431±0.0282	0.9705±0.0158	0.9747±0.0334	0.9941±0.0019
BET	0.9310±0.0089	0.9642±0.0048	0.9875±0.0117	0.9892±0.0014
HWA	0.8537±0.0184	0.9210±0.0107	0.9992±0.0003	0.9695±0.0053

Table 2.1: Results from the application on the LPBA40 dataset of deformable organisms (DO), the brain extraction tool (BET), and the brain surface extractor (BSE). [mean±S.D.]

2.2 Deformable Organisms and Error Learning for Brain Segmentation

2.2.1 Introduction

Deformable organisms label objects in images by integrating high level control mechanisms into a segmentation plan. They combine sensing, knowledge representation, reactive behavior, and proactive planning so our devised organisms may cooperatively segment an image. Deformable organisms were introduced into medical imaging by (McInerney et al., 2002) who combined ideas from artificial life (Steels, 1993) and deformable models (McInerney and Terzopoulos, 1996; Terzopoulos et al., 1987). Since their introduction, deformable organisms have been used for limb delineation (McIntosh et al., 2007), and segmentation of the spinal cord (McIntosh and Hamarneh, 2006b), vasculature (McIntosh and Hamarneh, 2006c), and corpus callosum in the brain (Hamarneh and McIntosh, 2005). (McIntosh and Hamarneh, 2006a) created a deformable organisms framework using the Insight Toolkit (ITK) (Ibanez et al., 2005), but we did not use it here, as its capabilities were too basic for our application. In contrast to several brain segmentation methods that work with low-level image processing and computer vision techniques, our deformable organisms can incorporate high-level knowledge and expectations regarding image data. In addition, almost every step of the plan presented here differs significantly from the one we previously presented (Prasad et al., 2011c).

Segmenting brain from non-brain tissues (such as the eyes, skull, scalp, and neck) in magnetic resonance (MR) images of the head is a vital pre-processing step for many types of image analysis. Accurate masks of the brain are helpful for longitudinal studies (Resnick et al., 2003), for multi-subject analyses of brain structure and function (Thompson et al., 2003), and as a pre-processing step prior to cortical surface modeling (Thompson et al., 2004), surgical planning (Gering et al., 2001), and brain registration (Woods et al., 1999).

The process of segmenting brain versus non-brain tissue on MRI is commonly referred

to as "skull-stripping" (although, strictly speaking, the skull generates almost no signal on T1-weighted MRI and the scalp and meninges are the main tissues removed). This has traditionally been done manually by trained experts, or by automated methods that are subsequently corrected by hand. Manually-created masks may also be used as gold standard delineations to validate performance of skull-stripping methods based on different principles. Though many approaches have been developed for this task, time consuming manual clean-up of these generated masks is almost always required. Most published methods do not perform well on all datasets, making improvements over existing methods critical.

There are a variety of existing skull-stripping methods. The Brain Extraction Tool (BET) (Smith, 2002) evolves a deformable model to find the boundary of the brain. It provides a robust way to find the boundary in unclear regions but does not incorporate prior knowledge of the brain's shape. The Brain Surface Extractor (BSE) (Shattuck and Leahy, 2002) uses edge detection and morphological operations to find the brain/non-brain boundary. BSE quickly extracts the brain from an image but may include extra material in the mask, as it sometimes fails to remove connections between the brain and surrounding tissue. The Hybrid Watershed Algorithm (HWA) (Segonne et al., 2004) uses the watershed algorithm to find the brain region, then fits a deformable model to the region, and finally deforms it based on a statistical atlas and geometric constraints. These methods have also been analyzed in (Boesen et al., 2004). We chose these methods as they are the most widely used and are part of larger neuroimaging toolkits.

We create a deformable organism plan that governs a collection of organisms to segment different parts of the head and brain. The organisms evolve dynamically in the images and cooperatively compute an accurate and robust segmentation of the brain. We then use a learning method, based on an Adaboost wrapper, (Wang et al., 2011), to classify the error in our method. We evaluate the effectiveness of this additional error correction step in improving our segmentations. We test our method with 630 T1-weighted MR images from healthy young adults, aged 20-30 years. We compare our approach to three widely used methods and we validate our results using distance, overlap, and error metrics.

2.2.2   Methods

Our deformable organisms method aims to segment the brain in T1-weighted MR images of the head. We describe our deformable organism plan to segment the brain, a way to learn and correct errors in our method, validation metrics to compare our results to the gold standard and to other widely-used methods, and our experimental results.

2.2.2.1   Deformable Organisms

Deformable organisms are organized in five different layers that combine control mechanisms and different representations to segment an image. We adapt this general approach for segmenting the brain.

Geometry and Physics   We represent each organism as a 3D triangulated mesh. These meshes are initialized on a standard brain template image. Our template was selected from the 40 images in the LONI Probabilistic Brain Atlas (LPBA40) (Shattuck et al., 2008), which have corresponding manual segmentations for 56 structures, and have manual delineations of the brain boundary. In the image we selected from this set, the voxels lying in each of our regions of interest are labeled. We fit our organisms to these labels to create a mesh using a marching cubes method (Lorensen and Cline, 1987) that goes through the image. The mesh is made up of polygons representing the border of the regions, which are then fused together.

These meshes deform to fit the 3D region that their corresponding organism is modeling. This iterative process moves each of the mesh's vertices along its normal direction with respect to the mesh surface. The surface is smoothed at every iteration using curvature weighted smoothing (Desbrun et al., 1999; Ohtake et al., 2000). This smoothing technique attenuates noise through a diffusion process as

$$\frac{\partial \vec{S}}{\partial t} = \lambda L(\vec{S}),$$ (2.1)

where S is the mesh, surface, or manifold and L is the Laplacian, which is equivalent to the

total curvature of the surface. This Laplacian is linearly approximated by

$$L(\vec{v_i}) = w_{ji} \sum \vec{v_j} - \vec{v_i}_{j \in N(i)}, \tag{2.2}$$

where v_i is the vertex i in the mesh, $N(i)$ are the neighbors of i, and w_{ji} is a weight proportional to the curvature between vertices j and i. We smooth the mesh to constrain its deformations to prevent intersections and artifacts from corrupting the boundary. On the boundary, we sample from the surrounding image.

Our prior based on an image from LPBA40 precludes us from using the Segmentation Validation Engine tool (Shattuck et al., 2009) to validate our method as it would lead to an unfair bias that would be favorable to our method; instead we use other metrics (below).

Perception Our organisms "sense" the encompassing image by sampling its intensities at vertices of the mesh. The vertices composing the mesh have real-valued coordinates, so we used nearest neighbor interpolation to find the intensities that correspond to them in the discrete grid of voxels in the image. The images may be any of the subject-derived volumes, which include the threshold image, 2-means classified image, 3-means classified image, or gradient image.

Paramount to the perception layer is the organisms' ability to sense each other's locations. We locate the voxels an organism resides on using a 3D rasterization algorithm (Pineda, 1988; Gharachorloo et al., 1989) that efficiently computes these values. These locations allow our organisms to dynamically change the way they deform based on their own positions and intensities of the original T1-weighted subject image.

Motor Control We move the vertices of the mesh along their normal direction with respect to the mesh surface by analyzing a set of intensities along this normal line. We describe the evolution of our mesh or surface $S(i, t)$ with respect to time t, where i is a vertex or point

on the surface, as

$$\frac{\partial \vec{S}(i,t)}{\partial t} = F(\mathbf{P}, \vec{n}, \mathbf{I}_d, b(c,l)_j)\vec{n}, \tag{2.3}$$

with F being the speed of evolution. F samples a set of positions P along point i's normal and interpolates these values from any of the derived images I_d, where d specifies the set of derived image. That set consists of the threshold image (t), 2-means image (2), 3-means image (3), and gradient image (g). The function $b(c)_j$ specifies any of a number of behaviors and decides how to move the point i on the surface by analyzing the sampled intensities l subject to a set of constraints, and weights the movement by the scalar c.

In practice, we evolve each vertex by the amount specified by F along its normal and progress through time by iterating through all vertices in the mesh until there is no longer any significant movement.

Behavior Our behaviors are a higher level of abstraction to indicate how the organisms function and what information they need to find. Behaviors may prescribe a function for organisms to be attracted to or repelled from landmarks or help converge on the boundary of an object. The functions for our behaviors had specific tasks in mind in the context of skull-stripping but are general and simple enough for repeated use in any segmentation task.

We developed six behaviors that play an important role in almost every step of our skull-stripping plan.

1. We create a behavior that analyzes a binary image and locates a boundary in these images. It contracts if a vertex corresponds to an off value and expands if it corresponds to an on value, and may be described as

$$b(c,l)_1 = \begin{cases} -c & \text{if } l_i = 0 \\ c & \text{otherwise} \end{cases} \tag{2.4}$$

In this case the set l consists of a single value l_i, the value of the binary image that corresponds to the vertex i.

2. Our second behavior moves a vertex outwards if its corresponding intensity value is q and may be described as

$$b(c,l)_2 = \begin{cases} c & \text{if } l_i = q \\ 0 & \text{otherwise} \end{cases} \tag{2.5}$$

Its purpose is to expand into an area of an image with voxels having intensities q. In addition q may be a set of labels that are appropriate for expansion.

3. Our third behavior is customized to move our skin mesh through the skull and skin into the cerebrospinal fluid (CSF). It contracts itself further if the boundary intersects another organism (in our plan, the other organisms are the eyes and we wish to deform through those areas) and will check the intensities (l) along the normal for those that correspond to CSF. If CSF markers m are found (signified by very low intensity values and a specific label in our k-means images) we contract the mesh, and if they are deficient, we expand the mesh. This behavior may be represented as

$$b(c,l)_3 = \begin{cases} -c & \text{if } l_i \text{ intersects other organism} \\ -c & \text{if } m \in l \\ c & \text{if } m \notin l \\ 0 & \text{otherwise} \end{cases} \tag{2.6}$$

In our framework, the intensities, l, are sampled from the k-means images and m is the label corresponding to CSF. The sampled points are locations inside the surface with respect to i.

4. the fourth behavior we created was designed for the brain organism to locate the boundary of the brain. It contracts vertices away from other organisms, contracts if there are CSF marker values (m) present within the surface, expands if the label value at i is not q, and expands if the gradient intensity at l_i is greater than or equal to a

26

threshold t. The precedence of these constraints is ordered as

$$
b(c, l)_4 = \begin{cases} -c & \text{if } l_i \text{ intersects other organism} \\ -c & \text{if } m \in l \\ c & \text{if } l_i \neq q \\ c & \text{if } l_i \leq t \\ 0 & \text{otherwise} \end{cases} \tag{2.7}
$$

The values at l are sampled from either the 3D rasterization image of the other organisms' meshes, the 3-means classified image, or the gradient image.

5. Our final behavior is the same as behavior 2 but instead of expanding it contracts.

$$
b(c, l)_5 = \begin{cases} -c & \text{if } l_i = q \\ 0 & \text{otherwise} \end{cases} \tag{2.8}
$$

Cognition   We create a plan of different behaviors to perform a segmentation task. The plan may dynamically activate different behaviors depending on what features the organisms were able to find in the image. Our plan to skull-strip the brain is one such plan.

2.2.2.2   Skull-Stripping Plan

Our skull-stripping plan combines our image processing and deformable organisms to create objectives in the image to extract the locations of different regions, culminating in extracting the boundary of the brain. In what follows, we describe each step in detail and how it depends on previous knowledge obtained by organisms. This is just one plan and may be fashioned for any type of segmentation or specifics of the data. Table 2.2 summarizes the steps each organism takes during the segmentation.

1. We begin by registering the subjects T1-weighted MR image to the template we selected from the LPBA40. This registration step is important to transform subject images into

a standard coordinate space as our organisms are tuned (iterations for deformations and labels for k-means classification) to images roughly corresponding to our template. Our template incorporates prior information and may be changed by users who need something closer to their own data. It provides initial locations and shapes for our skin, eye, and brain organisms.

We used an affine transformation for registration provided by FMRIB's Linear Image Registration Tool (FLIRT) (Jenkinson and Smith, 2001). FLIRT uses the correlation ratio (Roche et al., 1998) as the metric between the two images that takes the form

$$\frac{1}{Var(Y)} \sum_k \frac{n_k}{N} Var(Y_k). \tag{2.9}$$
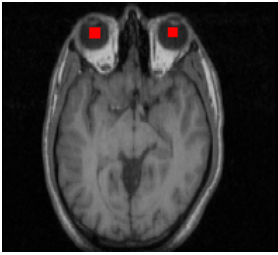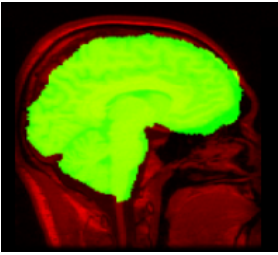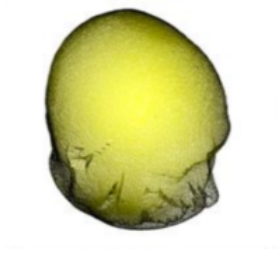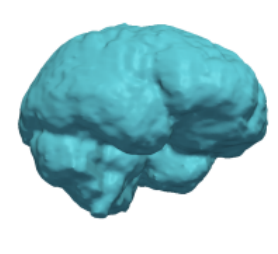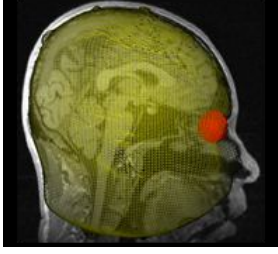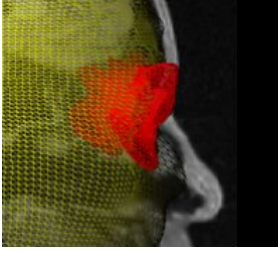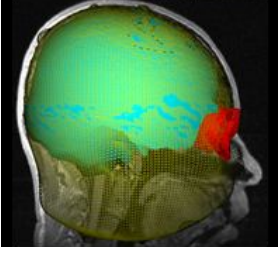
Y represents one of the images, $Var(Y)$ is the variance of Y, Y_k is the k-th iso-set i.e. the set of intensities in Y at positions where the other image X has intensities in the kth bin, n_k is the number of values in Y_k with $N = \sum_k n_k$. This cost is optimized to find a 12-parameter affine transformation.

In addition, we compute the inverse transformation to take the subject image back to its native space at the end of the segmentation.

2. We find the location or boundary of the skin with the skin organism. Its initial shape is of the skin of our template image found using the marching cubes method. We dilated this substantially to ensure we encompass the head of any subject registered to the template. Our template-fitted mesh needs to be further refined to fit our subject. To do this, we analyze the subject's intensities and apply a threshold to create a binary image masking out the head. We also use behavior 1 to sense the threshold image and evolve our skin organism's mesh to find this perimeter. We iterate the deformations dictated by behavior 1 (applying smoothing at every step) until there is no significant movement of the surface or we reach a maximum iteration bound. We specify this bound based on images being reasonably aligned to the template, an approach we used for all our deformations. The adjacent eyes are handled in a similar manner.

3. Our eye organisms find the eye boundary by sensing the 3-means classified image. We initialize the eye organisms' meshes by fitting them to our template and eroding them to make sure they lie within a subject image's eyeballs. The eyeball locations are found by the organisms sensing the 3-means image with behavior 2, which chooses a label found in the eyeball region.

4. Our next step locates the cerebrospinal fluid (CSF) that surrounds the brain. We achieve this using our skin organism by contracting its mesh into the head through the skin and skull. The skin and skull locations are roughly classified in our 2-means image and we apply behavior 3 to sense it and find the CSF boundary. To further refine this boundary the behavior also makes the skin organism deform through the eye organisms because there is more information about the CSF location by the eyes.

5. This step finds the tissues surrounding the eyeballs that need to be excluded from the brain delineation. We attain this goal by expanding the eye organisms further by sensing the 2-means image along with behavior 2 again, this time behavior 2 looks for a different label in the classified image, one that gives an understanding of tissues around the eyes. The eyes now furnish a better understanding to locate the brain.

6. We complete our plan by finding the brain using our brain organism. Every step in the plan supports of this step and all of our knowledge up to this point will come into play. Our brain organism begins by sensing the 3-means image and gradient image with behavior 4. The behavior is cognizant of the other organisms' locations and uses them to constrain its evolution. With the completion of behavior 4 we further refine the boundary by sensing the 3-means image again with behavior 5, which results in the brain being encapsulated by our brain mesh. The mesh is then converted by the 3D rasterization scheme to a binary volume to which we apply our inverse transform from FLIRT to bring the subject's delineation back into its native space to complete the segmentation.

Table 2.2: Plan for each organism to skull-strip an MRI image.

	Skin Organism	Eye Organism	Brain Organism
Initialization			
Deformation 1			
Deformation 2			

We make a binary volume of the brain organism and apply the inverse transformation back to the subject image space, completing the segmentation.

2.2.2.3 Error Learning

We are able to learn the types of errors our method makes, by comparing the masks it generated with expert manual delineations. (Wang et al., 2011) introduced an algorithm using Adaboost (Freund and Schapire, 1995) to learn a weighting of a set of features used to classify if a voxel has been correctly classified by a prior algorithm. This 'Adaboost wrapper' algorithm uses a set of corresponding automated and manual segmentations, as well as intensity images to find features that lie in regions that the first-pass method incorrectly classifies. We use this algorithm to learn situations in which our method makes errors and

thereby improve the segmentation.

2.2.2.4 Validation

The masks from the deformable organism method are compared with the gold standard manual segmentations using standard distance, overlap, and error metrics. We use the Hausdorff distance measure (Huttenlocher et al., 1993) to find the distance from the furthest point in the deformable organism method mask to the closest point of the mask in the manual delineation. We also compare the expert and automated masks using the Jaccard coefficient (union overlap), Dice coefficient (mean overlap), false negative rate, and false positive rate described in (Klein et al., 2009).

We compared the differences in metric values across methods using paired-sample t-tests to understand if their results were statistically different.

2.2.2.5 Experiments

We tested our deformable organism skull stripping method on our set of 630 manually-labeled subject images. In addition, we also ran BET (BET2, FSL 4.1.5, default parameters), BSE (BSE 10a, default parameters), and the Watershed algorithm (Freesurfer 5.0.0, default parameters), and assessed their errors using standard distance-based, overlap, and error metrics.

Typically patterns of error in our method were learned by selecting a subset of our segmentation results and using the error learning algorithm (the 'Adaboost wrapper' approach). We then segmented a new subset of images with the error classified and corrected. We repeated this experiment 10 times, using 10 random images from our results to train and 10 random (but non-overlapping) images to test. Masks were then compared to expert ground truth before and after error correction; note that the test set of images was independent of those used for training the error correction step.

31

2.2.3 Results

2.2.3.1 Subject Data

Our subject data consisted of 630 T1-weighted magnetic resonance (MR) images from healthy young adults, between 20 and 30 years of age. These images are from Australian twins, and have been used in numerous prior analyses (Joshi et al., 2011).

Each of the images had been manually skull-stripped by a neuroanatomically trained expert. These manual labels were used as the gold standard to compare with automatic segmentation results of our method and the other 3 widely-used methods. The subjects were scanned with a 4-Tesla Bruker Medspec whole-body scanner. 3D T1-weighted images were acquired using a magnetization-prepared rapid gradient echo sequence, resolving the anatomy at high resolution. Acquisition parameters were: inversion time (T_I)/repetition time (T_R)/echo time (T_E)=700/1500/3.35 ms, flip angle=8°, slice thickness=0.9 mm with a $256 \times 256 \times 256$ acquisition matrix.

In addition, we used one of the 40 images from the LONI Probabilistic Brain Atlas (LPBA40) (Shattuck et al., 2008). Each image had 56 different structures manually labeled, including a mask of the brain.

Table 2.3 shows the distance, overlap, and error metrics for the automated skull-stripping algorithms compared to manual segmentation. We compare BET, BSE, and the Watershed method. Average metrics over the 630 subject images are shown. Our paired-sample t-tests showed the metric values across all the methods were significantly different. These test results showed that the deformable organisms approach was statistically better than the others in the Hausdorff distance and false negative error for our dataset.

The deformable organisms method took a few minutes to run on the subject images we used on a machine with dual 64-bit 2.4 gigahertz AMD Opteron 250 CPU with 8 gigabytes of memory.

We list average results of deformable organisms with and without error correction versus

Table 2.3: Distance, overlap, and error metrics comparing automated results with manual skull-stripping.

	Hausdorff Distance	Jaccard Overlap	Dice Overlap	False Negative Error	False Positive Error
Deformable Organisms	36.3475±24.3842	0.8478±0.0242	0.9175±0.0143	0.0253±0.0115	0.1328±0.0280
BET	41.7997±24.7972	0.8860±0.0183	0.9395±0.0104	0.0711±0.0218	0.0491±0.0189
BSE	64.7451±25.7714	0.8348±0.1295	0.9040±0.0842	0.1045±0.1456	0.0720±0.0323
Watershed	67.4672±8.7859	0.3650±0.0599	0.5321±0.0617	0.4511±0.0613	0.4831±0.0648

Table 2.4: Distance, overlap, and error metrics comparing the deformable organisms segmentation with and without error correction versus manual skull-stripping.

	Hausdorff Distance	Jaccard Overlap	Dice Overlap	False Negative Error	False Positive Error
Basic DO	35.4780±2.7890	0.8485±0.0009	0.9178±0.0005	0.0256±0.0006	0.1318±0.0010
DO with correction	35.5194±3.1299	0.8858±0.0013	0.9393±0.0007	0.0253±0.0004	0.0929±0.0012

manual training in Table 2.4. Random samples of 20 images from the 630 were selected, using 10 to train and 10 to test the error correction. We repeat this 10 times and average the results. These average results were used a for two-sample t-test that found statistically significant improvement in Jaccard, Dice, and false positive error metrics.

CHAPTER 3

White Matter Representation & Registration

3.1 Atlas-Based Fiber Clustering for Multi-Subject Analysis of High Angular Resolution Diffusion Imaging Tractography

3.1.1 Introduction

Fiber integrity and connectivity in the living brain may be measured using high angular resolution diffusion imaging (HARDI), which reconstructs the local profile of water diffusion at each point in the brain. Following the dominant directions of local water diffusion throughout the white matter, is one way tractography methods can recover the geometry and connectivity of the major white matter fiber pathways.

In clinical research, tractography may be used to reconstruct white matter tracts for surgical planning (Nimsky et al., 2005). Analysis of structural connectivity is also useful for understanding coherent activity in functional networks. There are several ongoing efforts to map the human connectome, and to detect altered patterns of connectivity in disease. Additionally, fiber tracts may be extracted to study hemispheric asymmetries and functional lateralization (Powell et al., 2006), and to identify genetic effects and sex differences in neural network organization (Jahanshad et al., 2011). However, before group differences in specific tracts can be studied, the large collection of 3D curves generated by tractography needs to organized into bundles that correspond to well-known white matter tracts (Mettler, 1964). Corresponding tracts and bundles also need to be matched across subjects, as a basis for statistical comparisons.

A wealth of methods have been developed to cluster fibers extracted using diffusion-based tractography (Visser et al., 2011; Wang et al., 2011). Some methods embed the problem in $\mathbb{R}^n$ (O'Donnell and Westin, 2007); they compute geometrical invariants for each curve; these are then stacked into high-dimensional vectors that are clustered. The resulting clusters of points may correspond to known white matter commissures such as the corpus callosum, or major fasciculi. Some methods are inspired by standard unsupervised pattern recognition methods. They divide a set of n observations into g groups, to maximize some metric of similarity among members of a group versus between members of different groups. Algorithms based on K-means, fuzzy clustering, hierarchical or agglomerative methods, and self-organizing maps can all be adapted, in principle, to cluster tracts. One such approach, spectral clustering (O'Donnell et al., 2006; O'Donnell and Westin, 2007) groups fibers based on pair-wise distance measures among fibers as well as their spatial locations. Probabilistic clustering has also been used, without anatomical constraints, based on polynomial regression mixture models (Gaffney and Smyth, 2005). One problem with clustering methods that do not use prior anatomical information is that in many cases, known fiber bundles may be split into two or combined by the algorithm. A consistent partition across subjects is often difficult to find, and results may not correspond to known tracts in neuroanatomical atlases. Overall, the goal of this work is to (1) generate representative curves for tract based analysis of HARDI; (2) derive them guided by regions of interest in a brain atlas, so that the resulting tracts better reflect known anatomy, and (3) to make the tracts analyzed robust to discrepancies between the atlas and each individual subject, which tend to make atlas-based methods harder to apply to new subjects. To assess the utility of the method, we show an illustrative application assessing hemispheric asymmetry in tract-based FA, and we assess how the results depend on the tractography method that provides a basis for the analyses.

To create consistent fiber maps across subjects, we extracted fibers with a whole-brain tractography method, based on the Hough transform, and developed an atlas-based method to cluster them. First we retain all curves that are consistent, according to a similarity metric, with those in a probabilistic atlas. We then remove spurious curves by applying a

median filter to a binary map representing the preliminary fiber clusters. We create a graph from a density image of the clustered fibers in each subject. With the Dijkstra shortest path algorithm (Dijkstra, 1959), we compute a representative path through this graph. This path follows areas with the highest fiber density. These maximum-density shortest paths offer a compact way to compare fiber bundles across a population, and compute multi-subject statistics. We assessed how the results depend on the choice of tractography algorithm, using one algorithm that used the full HARDI ODF, and a more standard streamline method, that follows the principal eigenvector of the diffusion tensor. We concluded that both methods gave satisfactory results for tract-based statistical analyses, with some characteristic differences in the paths extracted; we comment on these in the Discussion.

3.1.2 Methods

3.1.2.1 Image data

We analyzed 105-gradient high-angular resolution diffusion images (HARDI) acquired from 20 healthy young adults (De Zubicaray et al., 2008), on a 4 T Bruker Medspec MRI scanner. 3D volumes consisted of 55 2-mm thick axial slices with a 1.79×1.79 mm^2 in-plane resolution. 94 diffusion-weighted images ($b = 1159$ s/mm^2) were acquired with a uniform distribution of gradient directions on the hemisphere. 11 b_0 (non-diffusion encoding) images were also acquired. Images were corrected for motion and eddy current distortions using FSL (www.fmrib.ox.ac.uk/fsl/).

3.1.2.2 Tractography methods

We analyzed two different tractography methods to explore the robustness of our clustering technique: (1) a novel Hough transform method to compute optimal paths through a field of constant solid angle orientation distribution functions (CSA-ODFs) derived from the HARDI data (Aganj et al., 2010), and (2) a streamline-based method based on the diffusion tensor (Mori et al., 1999). We used *TrackVis* (Wang et al., 2007) to visualize fibers and clusters

from both tractography methods.

The Hough transform method (Aganj et al., 2011) tests 3D fibers that pass through seed points throughout the image space. It assigns a score to each fiber, based on the log-probability of the existence of the fiber, according to a cost function. The cost function measures the generalized fractional anisotropy (GFA) (Tuch, 2004) along the fiber path, as well as the probability of a tract propagating in any given direction, derived from the ODF. The algorithm applies a Hough transform based voting process to a vast set of 3D curves in the volume. It assigns a score to each of them, and chooses those with the highest scores as potential tracts. The Hough transform is used to assign the highest scoring curve through a seed point as the potential fiber tract by having the voxels vote for the possible curves. For a single subject, each curve's score is defined as:

$$\int (\log[ODF(\vec{x}(s), \vec{t}(s)) GFA(\vec{x}(s))] + \lambda) ds \tag{3.1}$$

where $\vec{x}(s)$ and $\vec{t}(s)$ are respectively the location and the unit tangent vector to the curve under consideration, at the arc length s, and λ is a positive constant used as a prior on fiber length.

We also used the Diffusion Toolkit (Wang et al., 2007) to calculate standard streamline based fibers from the diffusion tensor model. Tracking follows the local principal diffusion direction, from a seed point in the image, to form a fiber. The stopping criterion is based on a strong change in fiber direction,

$$R = \sum_i^s \sum_j^s |v_i \cdot v_j| / s(s-1), \tag{3.2}$$

where s is the number of data points referenced on the fiber, v_i is the direction of maximal diffusion at point i. *TrackVis* was used to visualize fibers and filter them based on regions of interest (ROI).

3.1.2.3 Probabilistic atlas clustering

Each subjects T1-weighted image was aligned by 9-parameter transformation to the Colin27 high-resolution brain template (Holmes et al., 1998) after manual skull stripping. The 11 b_0 images from the DWI were averaged and aligned to the corresponding registered T1-weighted image, after they were masked using BET (Smith, 2002). In this space, FA maps were generated for each subject. We also created a geometrically-centered study-specific mean FA template (or minimal deformation template; MDT) (Jahanshad et al., 2010).

To cluster fibers from the tractography algorithm we used the JHU white matter tractography atlas (Wakana et al., 2004), which contains 17 prominent white matter tracts based on a set of fibers traced in 4 healthy subjects. The FA image from the JHU atlas was affinely transformed to the same MDT using FMRIB's Linear Image Registration Tool (FLIRT) (Jenkinson and Smith, 2001). The transformation was also applied to each of the 3D probabilistic tract images.

Figure 3.1 shows the streamline tractography results with *TrackVis*. The fibers shown are those that intersect any of six selected regions from the fiber tract atlas: the *forceps major* (*shown in green*); *forceps minor* (*orange*); left superior longitudinal fasciculus (*purple*); right superior longitudinal fasciculus (*red*); left inferior fronto-occipital fasciculus (*burgundy*); and right inferior fronto-occipital fasciculus (*blue*). Fibers were retained if they passed through any of the six regions. The fibers are colored by the direction of their middle segment.

This figure gives insight into how we can use the atlas to select consistent sets of fibers in the brain. *TrackVis* allows set-based operations (logical AND/OR/NOT) with ROIs to select fibers, but generally requires manual effort to interact with the fiber sets on an individual fiber level.

To allow atlas-based fiber labeling independent of *TrackVis*, we fitted a 3D cubic spline curve to each fiber, and regularly sampled it. We found the set of voxels that it intersects with in the encompassing image space, then we computed their overlap with the white matter tract of interest in the atlas. The number of voxels in the intersection measures how well a

tract fits into an atlas region. We then selected all tracts that intersected the atlas. This automatically selects fibers from a population of images, or from each individual subject. Figure 3.2(a) shows the part of the fiber tract atlas that represents the *forceps major* region of the corpus callosum, which includes the splenium. In this region, fibers arch backwards into the occipital lobes at the back of the brain. Figure 3.2(b) shows all the tracts that intersect it. The tracts shown here were computed using the Hough method. The fibers are colored by the number of voxels they intersect that are within the atlas region.

As shown in Figure 3.2(b), choosing all possible fibers that intersect a given white matter atlas region does not necessarily yield a good representation of that particular fiber bundle. Spurious curves may be included. Because of this, a threshold based on the voxel intersection count was used to further constrain which fibers were assigned to that fiber bundle (Figure 3.2(c)). *TrackVis* enables the selection of fibers that intersect an ROI. In the case of the JHU atlas, *TrackVis* allows the user to reduce the size of the ROI. This helps constrain the number of fibers selected. We employed a threshold that differed from *TrackVis* in that it enabled us to select fibers based on how much fiber passed through the atlas. We chose the threshold to be the thickness of the ROI, although varying this did not affect the results drastically.

3.1.2.4   Median filtering of tract density

Once a set of tracts has been found in one subject using the JHU atlas, we create a binary volume image where a voxel is 1-valued if a fiber intersects it and 0 otherwise. We computed all the voxels that the selected fibers pass through, to create one binary volume image. We then apply median filtering (Tyan, 1981) to this volume image. Median filtering replaces each voxels value by the median of the voxel values in its neighborhood (which for our binary data, is like a voting process). In our experiments, we used a $5 \times 5 \times 5$ box as the neighborhood centered on each voxel. The median filter changes any 1-valued voxels far from the main cluster to 0, to suppress noise and erroneous fibers. We use the median-filtered binary image to further screen the fibers in the cluster by calculating the intersection of the binary image

voxels with fibers in the cluster. We remove a fiber if all the voxels it travels through are not included in the filtered binary image. This removes fibers far from the main cluster, and creates a dense set of tracts with a coherent shape.

The resulting set of fibers after this median filtering step is the final set included in the fiber bundle and serve as the basis for the white matter tract. Representative results of median filtering are shown in Figure 3.3(a).

3.1.2.5 Representing the cluster by the maximum density path

To compactly represent the fiber bundle in each subject, we create a path through the bundle to follow the highest density of fibers. We first create a fiber density volume image by counting, at each voxel, how many fibers pass through it. We then create a graph (a set of nodes and undirected edges connecting them) from this density image. Nodes in the graph represent the voxel locations that have a non-zero density value, and the edges connect each voxel to its surrounding 26 neighboring voxels. The value of the edge connecting nodes i and j is dictated by:

$$\frac{1}{d_i + d_j} \tag{3.3}$$

where d_i is the density at the voxel corresponding to node i. This makes the cost of traveling from node i to node j inversely proportional to their density.

We then select a start and end node in the graph by first assigning a start and end voxel location in the JHU atlas. The closest non-zero voxel locations in the density image (by Euclidean distance) to the start and end voxels in the atlas are used as the voxels that correspond to the start and end nodes in the graph. Dijkstra's algorithm (Dijkstra, 1959) is then run on the graph to find the shortest path connecting the start and end nodes. Dijkstras algorithm is a graph search algorithm that efficiently finds the shortest path from a start node to every other node in the graph. In this case, the shortest path will include edges connecting nodes with high density values. As a result, it follows the path with the highest density of fibers in the image.

This resulting path is an ordered set of voxels. The centers of the voxels in the path are then used as the coordinates that trace the path and a 3D cubic spline is fitted to them. The resulting path is used to compactly represent the fiber bundle. The maximum density path of the fiber bundle in 3.3(a) is shown in 3.3(b). By computing the shortest path in each subject, we can represent the fiber bundles in all subjects in the entire dataset. The FA along this fiber path may then be used to compare fiber bundles across subjects or groups (Fillard et al., 2003).

3.1.2.6 Paired t-test of FA

Next we illustrate how to use the method to obtain tract-based statistics (Corouge et al., 2005). Once the maximum density curves are computed, we performed a paired-sample t-test comparing the FA along the maximum density curves from the Hough transform method to those from the streamline method. This was done pointwise, by uniformly sampling 100 points across each curve and finding the FA value at each point from the corresponding FA image, via cubic interpolation. We then found the mean FA along each curve and used that as an attribute for statistical analysis along that curve.

Another paired-sample t-test was used to assess the degree of hemispheric (left/right) symmetry of the tracts using each method. We split the curves in two by separating them at the midsagittal plane. Then we found the mean FA along each of these curves and compared them. This test was applied to the group of maximum density curves for each tractography method.

3.1.3 Results

To test the proposed approach, we ran the methods on 20 subjects images. Fibers were found using the Hough transform (*green*) and streamline (*blue*) methods. Figure 3.4 shows 40 different maximum density paths, one for each of the two tractography methods, per subject. Paths are then readily compared across subjects and methods.

The streamline method generates many more fibers than the Hough transform method. This is partly because the cost function does not rate fibers based on their overall geometry, as the Hough method does. Its density image also covers a larger volume, and all maximum density paths originate at the user-specified start and end points. Fibers were manually selected using *TrackVis* by using the NOT logical function to remove fibers outside the JHU atlas region.

The paired-sample t-test comparing the FA along the maximum density curves from the Hough transform method compared to those from the streamline method gave a p-value of 8.8×10^{-8}. The paired-sample t-test assessing the hemispheric asymmetry of the maximum density curves from the Hough transform method gave a p-value of 0.56. The test for the symmetry of the curves from the streamline method gave a p-value of 9.4×10^{-3}. The difference may be because there are significantly more curves from the streamline method so that the maximum density curve has a somewhat different trajectory. The fibers generated from each method may also differ in shape. In future, it may be fairer to compare results using the same number of fibers for both methods, although there would have to be a principled decision about which fibers to keep. Also, the Hough method is somewhat more computationally intensive (given its more complex mathematical formulation), making it easier to sample a very high number of curves with the streamline method.

We then found the FA along each of these curves; results are plotted in Figure 3.5. Shown in green are the curves from the fibers generated using ODF/Hough tractography. In blue, we show the curves generated from the streamline single-tensor method. In general, the FA along the curves from the Hough method was greater, though both approaches had the same general pattern (higher at midline).

3.2 White Matter Tract Analysis in 454 Adults using Maximum Density Paths

3.2.1 Introduction

Diffusion weighted imaging (DWI) captures the diffusion of water in the brain *in vivo*. The brain's structure may be understood by analyzing the local diffusion information these images provide. These images are used by tractography algorithms to follow the dominant directions of diffusion to reconstruct fibers oriented along axons that correspond to the brain's major white matter pathways. We can then study these white matter regions in individuals and populations to gain localized understanding of the effects of disease (Takahashi et al., 2002), changes with age (Abe et al., 2002), hemispheric differences (Jahanshad et al., 2010), and sex differences (Peled et al., 1998).

The wealth of information from high angular resolution imaging (HARDI) images affords us a better understanding of the brain's structure when compared to the single-tensor model (Basser and Pierpaoli, 1996). The single-tensor model does not account for fiber crossing or mixing, while the orientation distribution function (ODF) (Tuch, 2004) derived from HARDI images allow us to discriminate multiple fibers passing through a voxel with different orientations. In contrast to previous tractography approaches our method leverages all the information rendered by our HARDI images to generate fibers using a novel tractography algorithm.

Tractography algorithms generate a huge number of fibers in the image space and need to be clustered for analysis. A wealth of clustering methods have been applied to tractography results including fuzzy clustering (Shimony et al., 2002), normalized cuts (Brun et al., 2004), k-means (O'Donnell and Westin, 2005), spectral clustering (O'Donnell et al., 2006), Dirichlet distributions (Maddah et al., 2008), and hierarchical clustering (Visser et al., 2011), and a Gaussian process framework (Wassermann et al., 2010). Many of these methods will benefit from prior anatomical information provided by an atlas of likely locations of the tracts in

the brain, helping to provide insight into when to split or combine clusters to conform to known anatomy.

In (Corouge et al., 2006) the authors model fiber results with parametrized curves enabling coordinate based population analysis. This analysis has also been widely done using tract based spatial statistics (TBSS) (Smith et al., 2006), which uses an alignment-invariant tract skeleton representation. Our approach provides additional analysis power because it finds homologies across our representations through registration allowing localized comparisons in a population.

We create a framework to generate fibers from HARDI images, cluster them by incorporating prior anatomical knowledge, find correspondences across our white matter regions, and use them to compute tract-based measures of FA and geometric variation by examining hemispheric and sex differences in fifty major white matter regions. Our method is tested on 454 subjects captured with 4-Tesla HARDI, to our knowledge the largest whole-brain tractography study to date.

3.2.2 Methods

Our method computes population statistics of white matter tracts from HARDI subject images of the brain by finding fibers, clustering the fibers, representing the clusters compactly, and finding correspondences across these representations through registration. The framework is designed to use a novel tractography algorithm on each subject image to calculate a set of fibers. Those fibers that intersect an atlas of white matter tract regions of interest (ROI) are selected as a cluster. The clusters are represented by a path following points of highest density in the bundle and these paths are brought into correspondence using geodesic curve registration for point-wise comparisons. The pipeline uses this to find localized differences in white matter tract fractional anisotropy (FA) and geometry across hemispheres and sex. We summarize the method in Figure 3.6.

3.2.2.1 HARDI Tractography using the Hough Transform

We compute fibers in each HARDI subject image using a novel tractography algorithm (Aganj et al., 2011) based on the Hough transform. It utilized the extensive information provided by HARDI at each voxel, parametrized by the orientation distribution function (ODF).

Our tractography method selects fibers in the diffusion image space by generating scores for all possible curves. The method incorporates fractional anisotropy (FA) from the single-tensor model of diffusion (Basser and Pierpaoli, 1996) into a prior probability to generate seed points. These seed points are used to generate curves that receive a score estimating the probability of their existence, computed from the ODFs from the voxels the curve passes through.

The ODFs at each voxel from our HARDI images were computed using a normalized and dimensionless estimator derived from Q-ball imaging (QBI) (Aganj et al., 2010). In contrast to previous methods, this method views the Jacobian factor r^2 in computing the solid angle (CSA) ODF as

$$\frac{1}{4\pi} + \frac{1}{16\pi^2}\text{FRT}\left\{ \bigtriangledown_b^2 \ln\left(-\ln\frac{S(\hat{u})}{S_0} \right) \right\}. \tag{3.4}$$

In this equation, $S(\hat{u})$ is the diffusion signal, S_0 is the baseline image, FRT is the Funk-Radon transform, and $\bigtriangledown_b^2$ is the Laplace-Beltrami operator. This estimate is more accurate mathematically and outperforms the original QBI definition (Tuch, 2004), meaning it provides superior resolution for detecting multiple fiber orientations (Aganj et al., 2010; Fritzsche et al., 2010) and more information for our scoring function.

The Hough transform tractography chooses fibers from all curves generated in the image space. The method scores as many fibers as possible arising from a seed point and uses the voting process provided by the Hough transform to select the best fitting curve through each point. These filtered curves comprise the final fibers produced by the method.

The method was run on each subject image and generated 2000 to 5000 fibers (Fig. 3.7 shows a representative example with our data), which we subsequently clustered using a

white matter atlas.

3.2.2.2   Fiber Clustering with White Matter ROI Atlas

The fibers we calculated using tractography based on the Hough transform are clustered by selecting those that intersect an atlas of white matter tract ROI. We use the atlas to incorporate prior anatomical information into our clustering.

We utilize the Johns Hopkins University (JHU) atlas, which delineates 50 white matter regions of interest (ROI). This ROI atlas is registered to our subject space using an affine transform provided by FMRIB's Linear Image Registration Tool (FLIRT) (Jenkinson and Smith, 2001).

For our analysis pipeline, the intersection of fibers with the atlas specifies a fiber bundle or cluster. We select an ROI in the atlas and compute the locations our fibers intersect. The intersection is quantified by a set of voxels each fiber passes through and we count the number of voxels that coincide with the selected ROI to determine a fiber intersection score. This score is used to select fibers that belong to a ROI and thus a white matter tract. The ROI will select many fibers that spuriously intersect it. We address this by a threshold that removes fibers with a relatively low score, the threshold is dependent on how many and what type of fibers the tractography algorithm samples. In our case, we tuned the threshold to fibers computed with our Hough transform based method.

3.2.2.3   Cluster Representation using Maximum Density Path

Once we have selected a fiber bundle to represent a white matter tract we compute a compact representation. This representation follows points of maximum fiber density in the bundle to create a path. We refer to this as the maximum density path (MDP) and it provides a basis for population comparisons and allows localized statistics and analysis.

We constructed a graph (a set of nodes and undirected edges) to represent the voxel-wise density in our fiber bundle. We computed a density volume of our fiber bundle to characterize

our search space. This space specifies the count of fibers that intersect each voxel. These voxels are used as nodes in a graph (a set of nodes and undirected edges connecting them) with those of positive value connected to its 26 neighbors by undirected edges. Each edge was weighted inversely by the sum of the voxel densities it connected as

$$\frac{1}{d_i + d_j},\qquad(3.5)$$

with d_i and d_j as node i and j's corresponding voxel density value.

We found a path through this graph following nodes with highest density using Dijkstra's algorithm (Dijkstra, 1959). Dijkstra's algorithm is a graph search method that finds the shortest path from a source node to every other node. However, the number of nodes in the graph may be large and when the algorithm is used for a single destination node it may be stopped once that path is found.

To find the shortest path Dijkstra's algorithm requires the graph to have start and end nodes, these are specified beforehand in the ROI atlas. The ROI points for the start and end locations may not have corresponding positive density values derived from our bundle, we find the closest voxels in the density volume as the corresponding start and end nodes in subsequent computation with Euclidean distance used as the metric.

Dijkstra's algorithm will find the path once the start and end nodes are finalized. It will first assign a distance value to every node in the graph, zero for the start node and infinity for the rest. The algorithm marks all nodes as unvisited and the start node as current. The current node's distance to its unvisited neighbors is calculated and if that distance is smaller then the neighbor's existing distance value it is replaced. The method will mark the current node visited and the current node's distance is final, meaning the distance from the start node to itself will remain zero. The new current node will be selected as any unvisited node and the algorithm will continue exploring nodes until the destination node is visited, resulting in a path to this node minimized in distance.

If Dijkstra's algorithm is unable to find the path between the start and end nodes our

method automatically identifies this situation and takes steps to remedy the graph and find the shortest path. The algorithm will be incapable of finding a connection between the two nodes if the graph structure is such that there are no edges from the subgraph containing the start node with the subgraph containing the end node, caused by scanner issues or a deficiency in fibers calculated from tractography. In this situation we add edges and nodes to the graph so all voxels within our ROI are fully connected with their neighbors. The edges are weighted with the smallest positive value our machine is capable of representing and allows gaps between the start and end nodes to be filled in without biasing the resulting shortest path.

Dijkstra's algorithm computes a set of voxel locations in our image space connecting the start and end nodes to create a path. We smooth the path so it is better conditioned for subsequent processing. We convolve the 3D coordinates of our path with a Gaussian kernel to achieve this, though fitting these points to a spline would also have sufficed.

Our resulting maximum density paths (MDPs) are then registered together to find correspondences across the set.

3.2.2.4   MDP Correspondences through Geodesic Curve Registration

The MDPs represent fiber bundles that characterize a tract in a low-dimension space. These compact descriptions of a tract's scale, location, and high-level geometric information are computed for all subjects in a population. We find correspondences across the paths by bringing them into the same space using geodesic curve registration. The registration's target is the population's mean MDP.

We use geometric features of the MDPs, based the on our method for matching 3D curves (Joshi et al., 2007a,b). The MDP path is represented by a parametrized curve in $\mathbb{R}^3$. Assuming a standard unit interval, the coordinate function of this path is denoted by $\beta(s) : [0, 1] \rightarrow \mathbb{R}^3$. The coordinate function is not a good choice to analyze the shape of the path, since it is confounded by global location, scale, and orientation. While the objective

of coordinate registration approaches is the establishment of spatial correspondences across tracts, our approach here is to establish shape correspondence between different tracts. As our goal is to analyze the tract geometry independent of the pose, we represent the shape of the coordinate function of the MDP using a vector valued function as

$$q(s) = \frac{\dot{\beta}(s)}{\sqrt{||\dot{\beta}(s)||}} \in \mathbb{R}^3 . \tag{3.6}$$

Here, $s \in [0,1]$, $|| \cdot || \equiv \sqrt{(\cdot, \cdot)_{\mathbb{R}^3}}$, and $(\cdot, \cdot)_{\mathbb{R}^3}$ is the standard Euclidean inner product in $\mathbb{R}^3$. We now denote $\mathcal{Q} \equiv \{q | q(s) : [0,1] \to \mathbb{R}^3 | \int_0^1 (q(s), q(s))_{\mathbb{R}^3} ds = 1\}$ as the space of all unit-length curves. On account of scale invariance, the space $\mathcal{Q}$ becomes a unit-sphere of functions, and represents all open elastic curves invariant to translation and uniform scaling. The tangent space of $\mathcal{Q}$ is given as $T_q(\mathcal{Q}) = \{w = (w_1, w_2, \ldots, w_n) | w(s) : [0,1] \to \mathbb{R}^3,$ $\forall s \in [0,1]$ such that $\int_0^1 (w(s), q(s))_{\mathbb{R}^3} \, ds = 0\}$. Here each w_i represents a tangent vector in the tangent space of $\mathcal{Q}$. We define a metric on the tangent space as follows. Given a curve $q \in \mathcal{Q}$, and the first order perturbations of q given by $u, v \in T_q(\mathcal{Q})$, respectively, the inner product between the tangent vectors u, v to $\mathcal{Q}$ at q is defined as,

$$\langle u, v \rangle = \int_0^1 (u(s), v(s))_{\mathbb{R}^3} ds. \tag{3.7}$$

Now given two shapes q_1 and q_2, the translation and scale-invariant shape distance between them is simply found by measuring the length of the geodesic connecting them on the sphere. To ensure rotational invariance, we search for an optimal geodesic path in a quotient space of the sphere under the special orthogonal group $SO(3)$.

In addition to computing geodesic paths on the space of shapes, one can also minimize the geodesic variance that consists of sum of squared pairwise geodesic distances for a given collection of paths. This optimization procedure enables one to compute statistical quantities, such as means and covariances of MDPs.

The MDP parametrization afforded by this registration method produces correspondences

49

that facilitate comparison across a population.

3.2.2.5 MDPs Compared across a Population

The registrations step provides features along each MDP that can be compared across the population in the form of points. We compute each point's variation of fractional anisotropy (FA).

3.2.3 Results

3.2.3.1 Subject Data

Our subjects comprised of 454 healthy young adult twins and their siblings. We captured these subjects with a 4-Tesla Bruker Medspec MRI scanner, collecting 3D 105-gradient high angular resolution diffusion images (HARDI) and standard structural T1-weighted magnetic resolution images (MRI). The images consisted of 55 slices 2-mm thick with a 1.79×1.79 mm^2 in-plane resolution. For each image, we collected 94 diffusion-weighted images ($b = 1159$ s/mm^2) using a uniform distribution of gradient directions on the hemisphere. Additionally, we collected 11 b_0 (non-diffusion encoding) images and corrected all images for eddy current distortions and motion with FSL (www.fmrib.ox.ax.uk/fsl). Our cohort consisted of 289 women and 165 men, ranging from 20 to 30 years of age.

We prepared our image data for our analysis pipeline. In each T1-weighted, we manually removed non-brain tissue and registered them to the Colin27 high resolution brain template using a 9-parameter transformation (Holmes et al., 1998). These skull-stripped and registered T1-weighted images each have a corresponding average b_0 diffusion weighted image (DWI), combining all eleven we collected. These average images were masked using BET (Smith, 2002) and this space was used to generate FA maps. Additionally, we used the FA images to compute a geometrically-centered study-specific mean template (mean deformation template; MDT).

We computed MDPs for 50 regions specified in the white matter ROI atlas for each of the

454 subjects. By registering the MDPs to their sample mean, we computed corresponding points across the MDPs. The FA at these locations was used for two-sample t-tests comparing MDPs in groups differing by sex. These tests were completed on the fibers derived using the Hough transform method and the widely used streamline method (Wang et al., 2011) that reduces the HARDI information to tensors and follows the direction of highest diffusion in the tensor i.e. the eigenvector with the largest eigenvalue.

Fig. 3.8 & 3.9 show the false discovery rate (FDR threshold < 0.05) (Storey and Tibshirani, 2003) adjusted p-values from two-sample t-tests on FA along each subject's MDP for the Hough transform and streamline based tractography methods respectively. p-values are overlaid in color (red being more significant) on the mean MDP for each white matter region.

Figure 3.1: Selected subsets of curves (extracted using *TrackVis*), based on their intersection with six different white matter regions from the JHU atlas. Atlas regions are shown in solid colors. These regions were used as a basis (ROIs) for curve selection and further analysis.
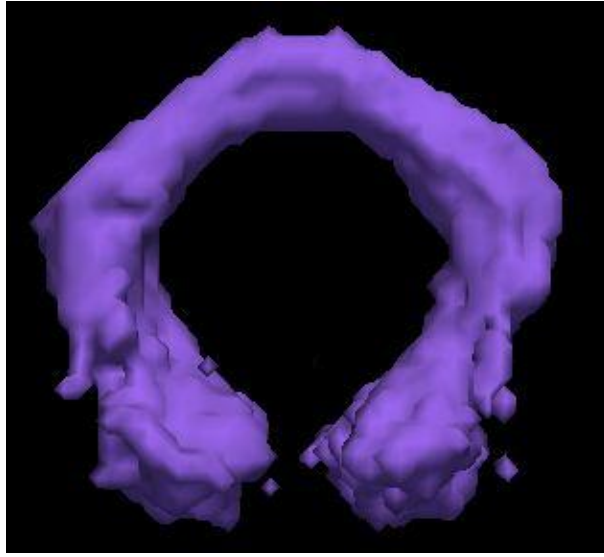
(a)
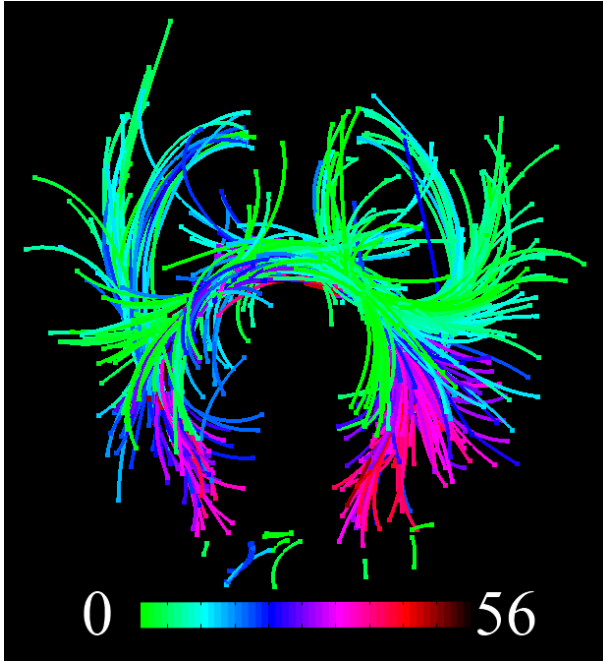
(b)
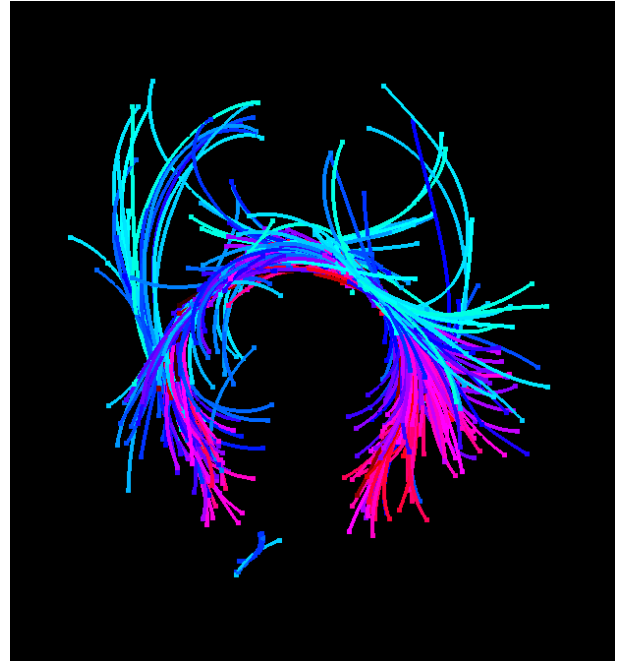
(c)

Figure 3.2: (a) *Forceps major* region of the corpus callosum from the JHU atlas. (b) Fibers from the Hough transform method that intersect the corpus callosum in the JHU atlas. (c) Fibers that remain after a threshold is applied to the number of points that must intersect the atlas. Fibers are colored by the number of voxels they intersect within the atlas.

(a)   (b)
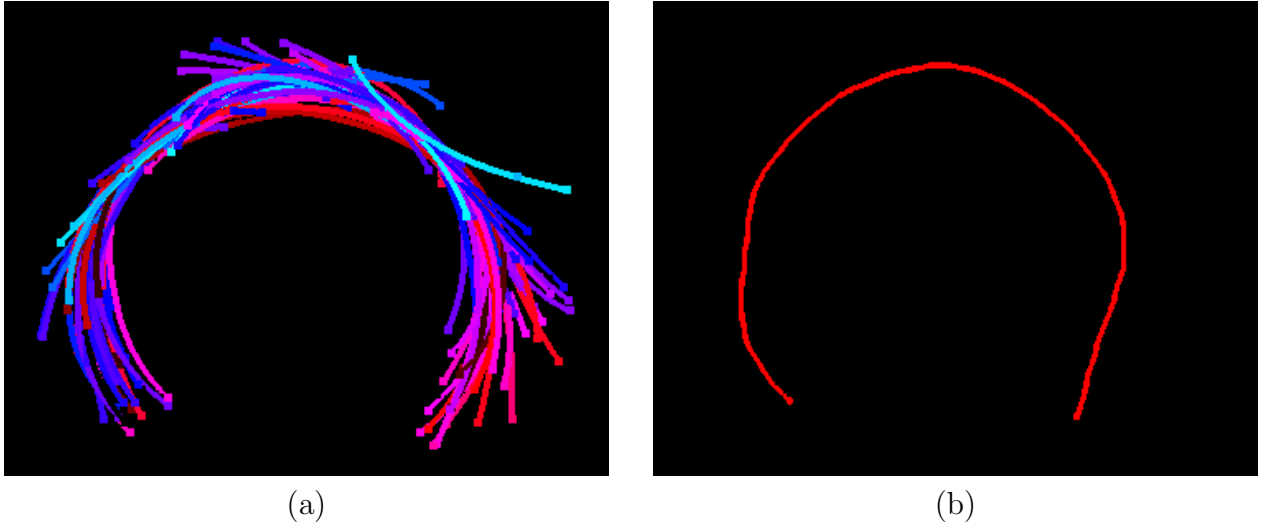
Figure 3.3: (a) Fibers included in the bundle after median filtering of the binary image. (b) Fiber path through the bundle, following the areas of highest density.

Figure 3.4: Maximum density curves from the Hough transform method (*green*) and the streamline method (*blue*). There are 40 curves in total: each method was run with the same 20 HARDI datasets.
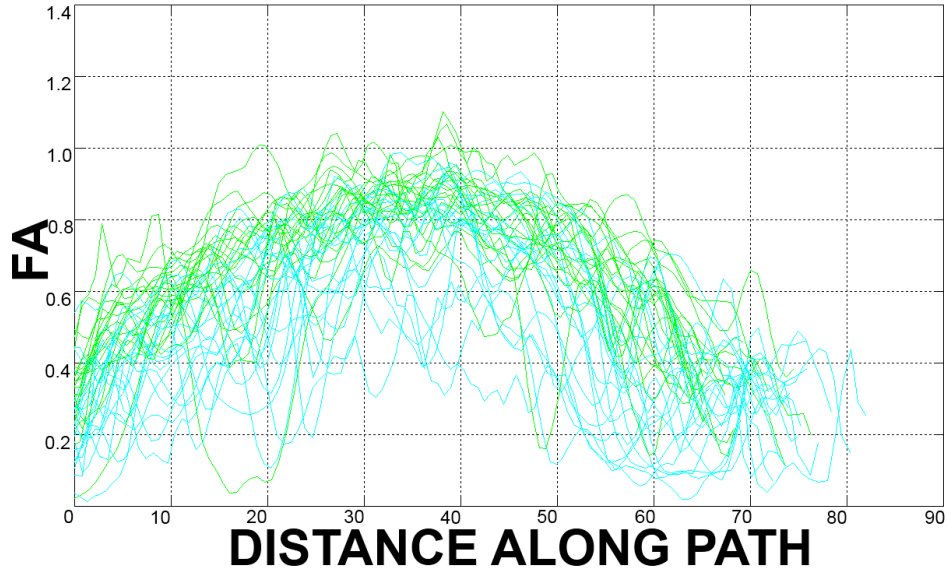
Figure 3.5: FA along the maximum density curves generated using fibers generated from the Hough transform and streamline method, for each of the 20 HARDI datasets. Interestingly, the Hough method picks curves with a higher FA, and there is some evidence that differences in FA are easier to detect when FA is higher, perhaps due to the higher SNR. For this reason, the Hough-based tracts may be more useful for studying factors that influence tract FA. The FA is slightly reduced away from the midline of the brain, partly due to a less degree of directional coherence among fibers within a voxel. This occurs because the fibers are all essential orthogonal to the midsagittal plane at midline, but mix with other tracts and fan out more as they deviate further from midline.
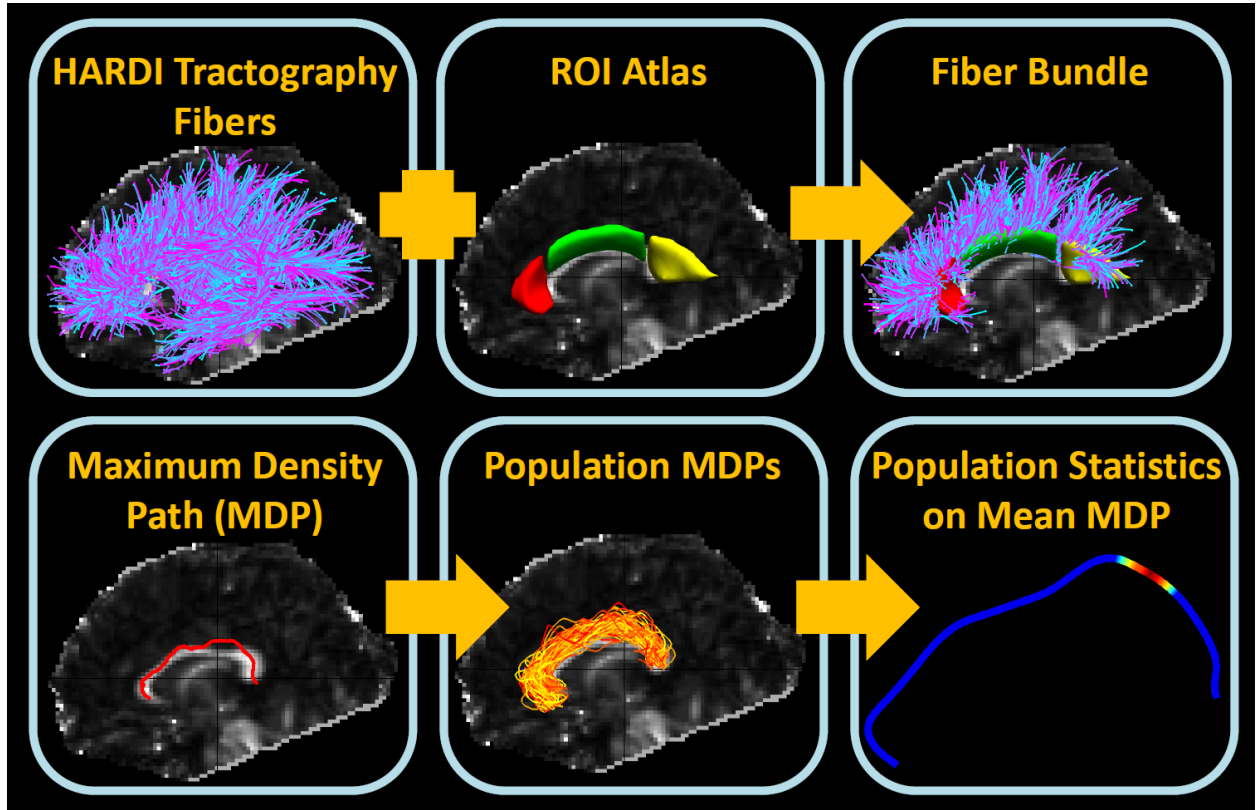
Figure 3.6: Schematic of the pipeline for statistical analysis of white matter tracts.
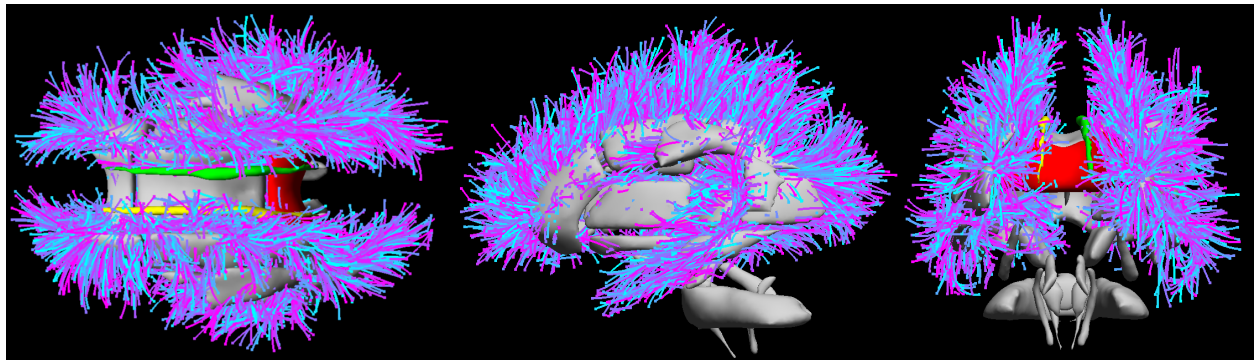
Figure 3.7: Fibers extracted using the Hough transform based method.
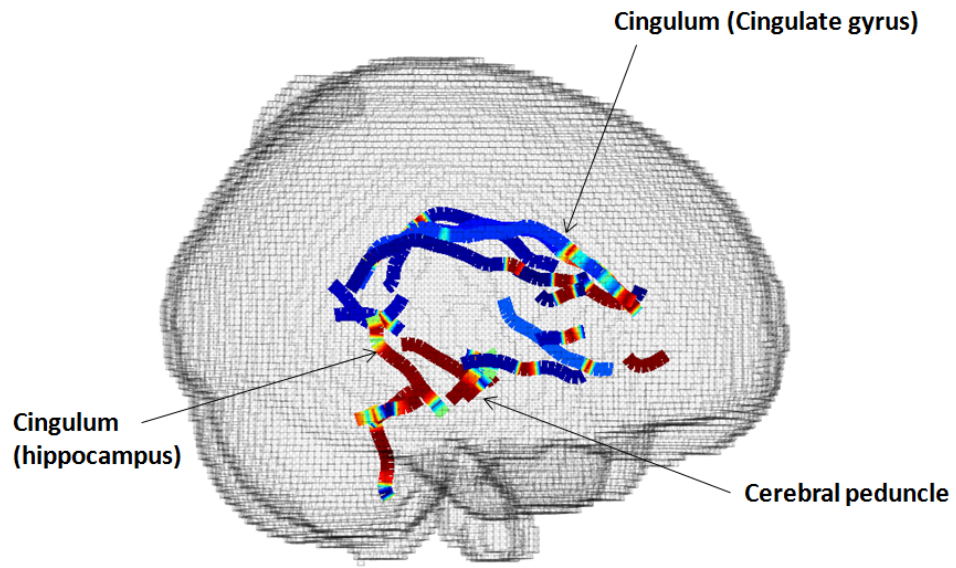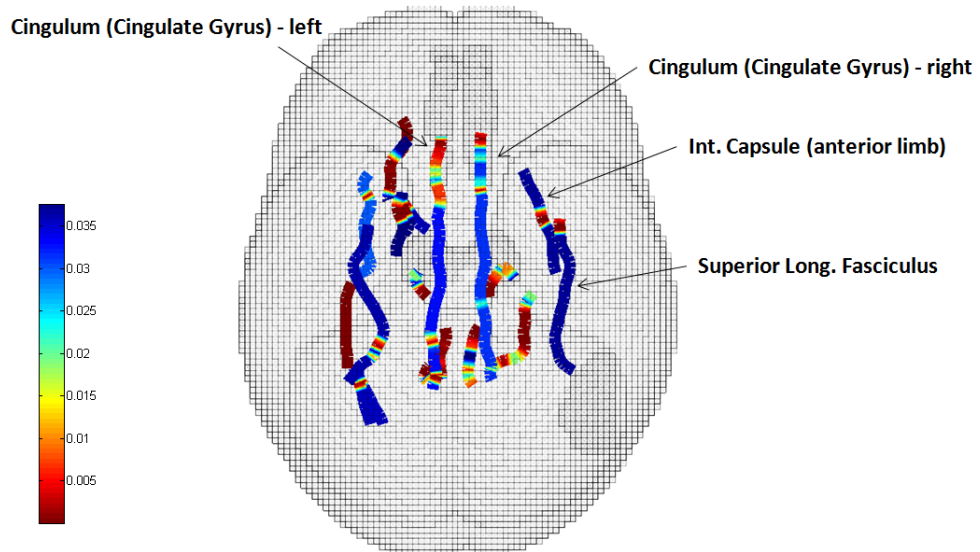
Figure 3.8: Significance test for FA along the mean density path for the Hough tractography analysis. Axial (a) and Sagittal (b) views of the mean Maximum Density Paths (red is more significant)
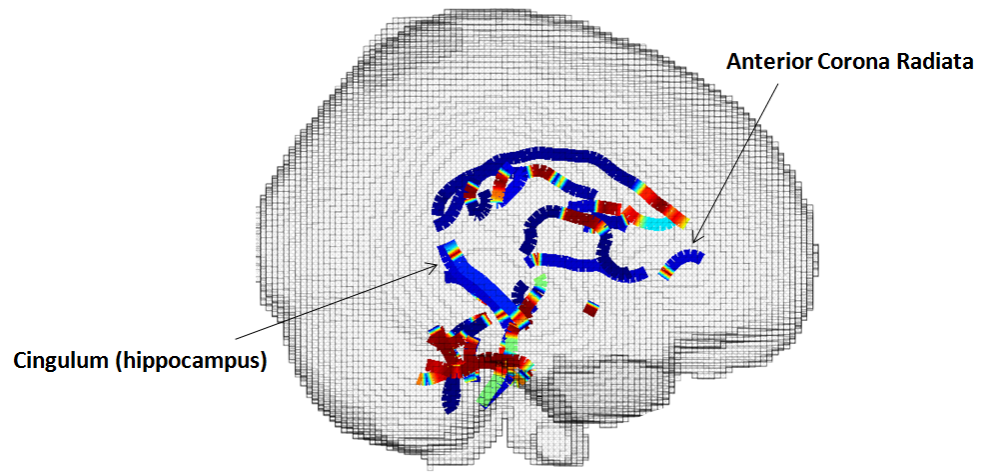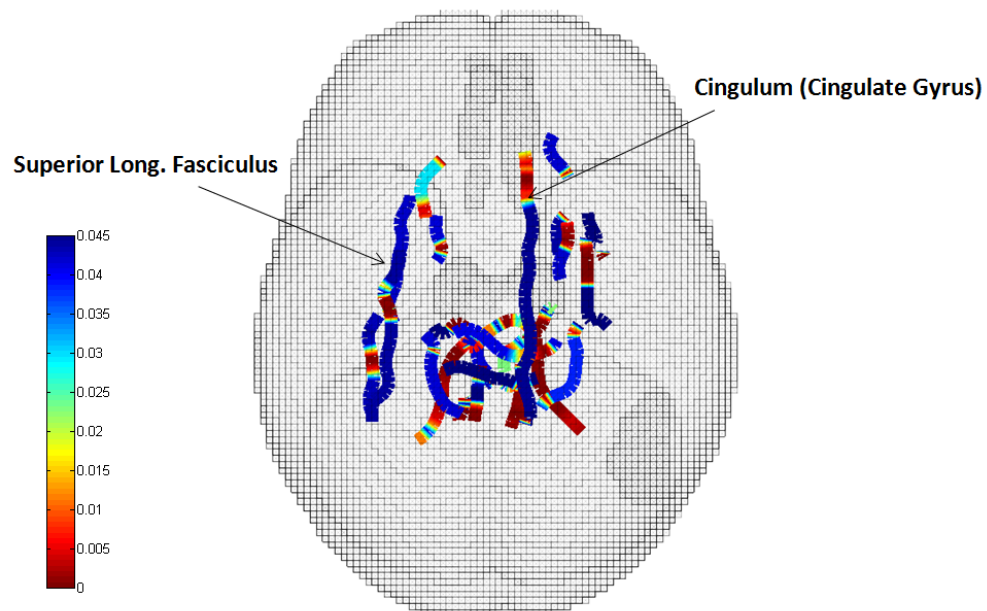
Figure 3.9: Significance test for FA along the mean density path for the streamline analysis. Axial (a) and Sagittal (b) views of the mean Maximum Density Paths (red is more significant)

CHAPTER 4

Connectivity

4.1 Tractography Density and Network Measures in Alzheimer's Disease

4.1.1 Introduction

Alzheimer's disease (AD) is the most common form of dementia worldwide, and to date has no known cure. To evaluate new treatments, image-based measures of disease progression are urgently sought, and diffusion imaging has recently been added to several large-scale studies of AD in the hope of capturing breakdown in brain integrity and connectivity not detectable with standard anatomical MRI. The connectivity of the white matter progressively declines in AD (Rose et al., 2000; Bozzali et al., 2002), but it is not yet known which connectivity measures best distinguish patients from controls, and how robust these measures are to specific aspects of the methods used, such as the density of extracted fibers. Here we studied brain networks in people with AD, in those at heightened risk for AD (with an intermediate stage of the disease known as MCI), and in healthy elderly subjects using structural magnetic resonance imaging (MRI) and diffusion MRI. We aimed to identify (1) which network measures best distinguish AD from healthy controls, and (2) the influence of tractography density on these measures. Our data is from the Alzheimer's Disease Neuroimaging Initiative (ADNI) (Mueller et al., 2005), an endeavor to find and evaluate biomarkers that capture disease progression or help predict risk for AD. We studied the imaging data from 110 subjects in total - 28 controls (C), 11 with late-stage mild cognitive impairment (L-MCI), 56 early MCI (eMCI), and 15 AD.

We used global probabilistic tractography to capture the structure of the white matter. The method leverages all the diffusion direction information parameterized by the orientation diffusion function (ODF). We accelerated the method by using a lookup table of diffusion values from the ODF. Cook et al. (Cook et al., 2005) used a lookup table in diffusion data, but they focused on the probability density function while our approach is specifically designed to speed up tractography. Our optimized tractography algorithm allows us to quickly compute thousands of fibers per subject and study how this density affects our network connectivity measures.

We describe the specifics of our subject data and how we pre-process it for connectivity analysis. Our optimization of the global probabilistic tractography is detailed including a rapid method to access the pre-computed information in the look-up table. We then detail how we compute connectivity matrices and network measures, which are the features used to compare the different disease categories. The details of these results are analyzed and discussed to understand the interplay of fiber density, network measures, and disease.

4.1.2   Methods

4.1.2.1   Data

Our data consisted of 110 subjects from ADNI-2, a continuation of the ADNI project in which diffusion imaging was added to the standard MRI protocol. The dataset as of 11/01/12 included 28 cognitively normal controls (C), 56 early- and 11 late-stage MCI subjects (eMCI, LMCI), and 15 with Alzheimer's disease (AD).

The subjects were scanned on 3-Tesla GE Medical Systems scanners, which collected both T1-weighted 3D anatomical spoiled gradient echo (SPGR) sequences (256×256 matrix; voxel size = $1.2 \times 1.0 \times 1.0$ mm^3; TI=400 ms; TR = 6.98 ms; TE = 2.85 ms; flip angle = 11°), and diffusion weighted images (DWI; 256×256 matrix; voxel size: $2.7 \times 2.7 \times 2.7$ mm^3; scan time = 9 min). The DWI images were composed of 41 diffusion images with $b = 1000$ s/mm^2 and 5 T2-weighted b_0 images. This protocol was chosen after an effort to study trade-offs

between spatial and angular resolution in a tolerable scan time (Jahanshad et al., 2011).

Image Preprocessing   We processed the T1-weighted images to parcellate them into 68 cortical regions. We first automatically removed extra-cerebral tissues from the anatomical images using ROBEX (Iglesias et al., 2011b), a method that learned from manual segmentations of hundreds of healthy young adults. These skull-stripped brains were inhomogeneity corrected using the MNI N3 tool (Sled et al., 1998) and aligned to the Colin27 template (Holmes et al., 1998) with FSL Flirt (Jenkinson et al., 2002). The resulting images were segmented into 34 cortical regions (in each hemisphere) using FreeSurfer (Fischl et al., 2004). These segmentations were then dilated with an isotropic box kernel of $5 \times 5 \times 5$ voxels to make sure they intersected with the white matter for subsequent connectivity analysis.

We corrected head motion and eddy current distortion in each subject by aligning the DWI images to the average b_0 image with FSL's eddy correct tool. The brain extraction tool (BET) (Smith, 2002) was then used to skull-strip the brains. We EPI corrected these images with an elastic mutual information registration algorithm (Leow et al., 2007) that aligned the DWI images to the T1 scans. More detail on these preprocessing steps is given in (Nir et al., 2012).

4.1.2.2   Tractography

We employed a global probabilistic tractography method based on the Hough transform (Aganj et al., 2011). The method takes advantage of all the diffusion information provided at each voxel, parametrized by the orientation distribution function (ODF). We note that although the ADNI DTI protocol is not a high-angular resolution imaging protocol, the use of an explicit ODF model rather than a single-tensor model makes best use of the available angular resolution, which is inherently limited as many patients cannot tolerate a long scan.

The Hough method generates curves in the fiber space and scores them based on fractional

anisotropy (FA) and the ODF at each point along the curve. The scoring function is

$$\int \left(\log[\text{FA}(\vec{x}(s))\text{ODF}_{\vec{x}(s)}(\hat{t}(s))] + \lambda ds \right), \tag{4.1}$$

where s is the arc length along the curve, $\hat{t}$ is the unit tangent vector, $\vec{x}$ denotes the coordinates along the curve, and λ is a prior on the length of the fibers (it is tuned so that curves oriented towards the cortex will be long enough to reach the gray matter). FA was computed from the single-tensor model of diffusion (Basser and Pierpaoli, 1996). The method then computed the ODF at each voxel from the diffusion data with a normalized and dimensionless estimator derived from Q-ball imaging (QBI) (Aganj et al., 2010). This formulation includes the Jacobian factor r^2 in the constant solid angle (CSA) ODF, defined as

$$\frac{1}{4\pi} + \frac{1}{16\pi^2}\text{FRT}\left\{\bigtriangledown_b^2 \ln\left(-\ln\frac{S(\hat{u})}{S_0}\right)\right\}, \tag{4.2}$$

where $S(\hat{u})$ is the diffusion signal in unit direction $\hat{u}$, S_0 is the baseline image, FRT is the Funk-Radon transform, and $\bigtriangledown_b^2$ is the Laplace-Beltrami operator. This model is more accurate and outperforms the previous QBI definition (Tuch, 2004), giving greater resolution for detecting multiple fiber orientations (Aganj et al., 2010; Fritzsche et al., 2010) and additional information for the scoring function.

The Hough transform tractography method searches through a set of seed locations for the highest scoring curves. These 3D curves are parameterized by their arc length with $s = 0$ centered at a seed point. The curve's unit tangent vector, $\hat{t}$, is described by standard polar coordinates, $\theta(s)$ and $\phi(s)$, which are each approximated by an N order polynomial. In addition, there are parameters L_- and L_+ that decide the partial length along either side of the seed point. These parameters are set to include fibers that can span the white matter in the image space. The coordinates are found by integrating the tangent vector as

$$\vec{x}(s) = \vec{x}_0 + \int_0^s \hat{t}(s')ds', \tag{4.3}$$

where $s \in [-L_-, L_+]$, and $\vec{x}_0$ is the seed location. This characterization means that the method will represent each curve using $2N + 4$ parameters and uses the voting process provided by the Hough transform to select the best fitting curve for a seed. If we restrict our search space to polynomials of order $N = 3$ and the maximum length to be longest expected fiber length in an image, the complete search at each seed point can result in millions of curves and upwards of one minute of computation time per seed, taking weeks to complete tractography on a single subject.

Optimizations We optimized the Hough transform tractography method with an ODF lookup table and a random search through the curve parameter space.

The Hough tractography method spends most of its computational time finding diffusion values at tangent directions along possible curves by evaluating the ODF. ODFs are represented as a spherical harmonic series, and this consumes a great deal of time after millions of evaluations. We remedy this issue with a lookup table that stores diffusion data in a set of random directions and is pre-computed before the Hough method evaluates the scores of the sampled curves. This saves time because the method does not have to evaluate the spherical harmonic series every time a curve point is evaluated and has immediate access to the diffusion information. We created the table by selecting n random directions, evaluating the ODFs at each voxel in our image at those directions, and then storing the resulting diffusion values in an array.

In addition, we created a direction lookup table that stores the index of the closest direction for an arbitrary θ and ϕ. This table allows us the query the ODF lookup table without comparing a requested direction to our n random direction each time we need to evaluate the ODF. We created the table by first selecting the range 0 to 2π for both θ and ϕ and sample this range at m points, giving m^2 possible pairs of (θ, ϕ). For these pairs we compute the closest direction, α, from our set of n random directions that satisfies

$$\underset{\alpha}{\operatorname{argmax}}\left(\left| \frac{\vec{\alpha} \cdot \vec{\beta}}{\|\vec{\alpha}\| \|\vec{\beta}\|} \right| \right), \tag{4.4}$$

where β is a possible (θ, ϕ) combination. This distance is the absolute value of the cosine of the angle between the two directions α and β. The absolute value is necessary because of the symmetry found in the ODF. Fig. 4.1 shows an example table for $n = 1000$ random directions and $m = 1000$ points between 0 to 2π. Each color represents the index of the closest random direction for each (θ, ϕ) pair.

We also cut down the computation time by randomly selecting a subset of the curve parameters to search. In our experiments we found that a search that covered only 10% of the parameter space would give a very close approximation to the exhaustive results. Together the random search and lookup table reduced the computation time for a single seed to less than one second and for an entire image to a matter of hours.

In our experiments, this arrangement of relying on a large precomputed database accessed by many smaller threads is well suited to the MapReduce model (Dean and Ghemawat, 2008) to distribute computations to grid resources. In particular, our lookup table was designed to be accessed with low latency using the distributed file system provided by Apache Hadoop, an open-source implementation of MapReduce.

We used the optimized Hough transform tractography method to compute $50,000$ fibers in each of our subjects. Fig. 4.2 compares $5,000$ and $50,000$ fibers in the same subject.

4.1.2.3 Connectivity

From the tractography results, we created connectivity matrices that measure the number of detected fibers that connect cortical regions to each other. To create a subject's matrix, we went through each of the fibers and compared it to the dilated cortical labels. The labels that intersected the fiber were logged into the connectivity matrix. This results in a 68×68 connectivity matrix for each subject as shown in Fig. 4.3. The figure shows the intra-hemispheric connections in the top-left and bottom-right squares and the inter-hemispheric connections in the bottom-left and top-right areas. The colors signify how many fibers connect the two regions. We refer to this as the fiber density, while noting that
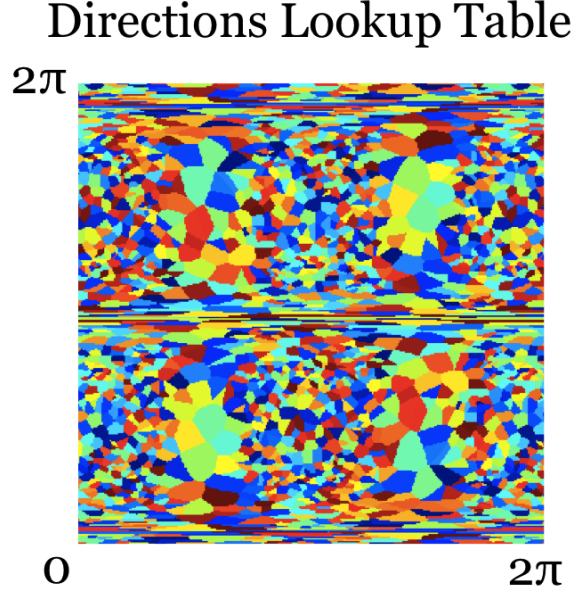
Figure 4.1: The directions lookup table affords immediate access to precomputed ODF information without having to compute the closest angle for a direction query. The colors represent the closest precomputed ODF direction for each θ and ϕ combination (each color represents the index or label for the closest pre-computed direction). This example shows $n = 1000$ random directions sampled on the sphere, with θ and ϕ ranging from 0 to 2π with $m = 1000$ intervals. This range is chosen to coincide with the query range from the Hough transform method.

the actual density of axonal fibers can only be determined histologically. Also the relation of our measure to the actual fiber density is complex as it depends on the seeding of the curves and the size of the ROIs on the cortex (for this reason, some other definitions normalize the fiber counts based on the areas of the cortical ROIs they interconnect).

We computed ten connectivity matrices for each subject. Each of these differed in the number of fibers used to create the matrix - ranging from 5,000 to 50,000, in increments of 5,000.

4.1.2.4 Connectivity Measures

We interpreted our connectivity matrices with eight network measures described in (Rubinov and Sporns, 2010), computed with the Brain Connectivity Toolbox. We chose global
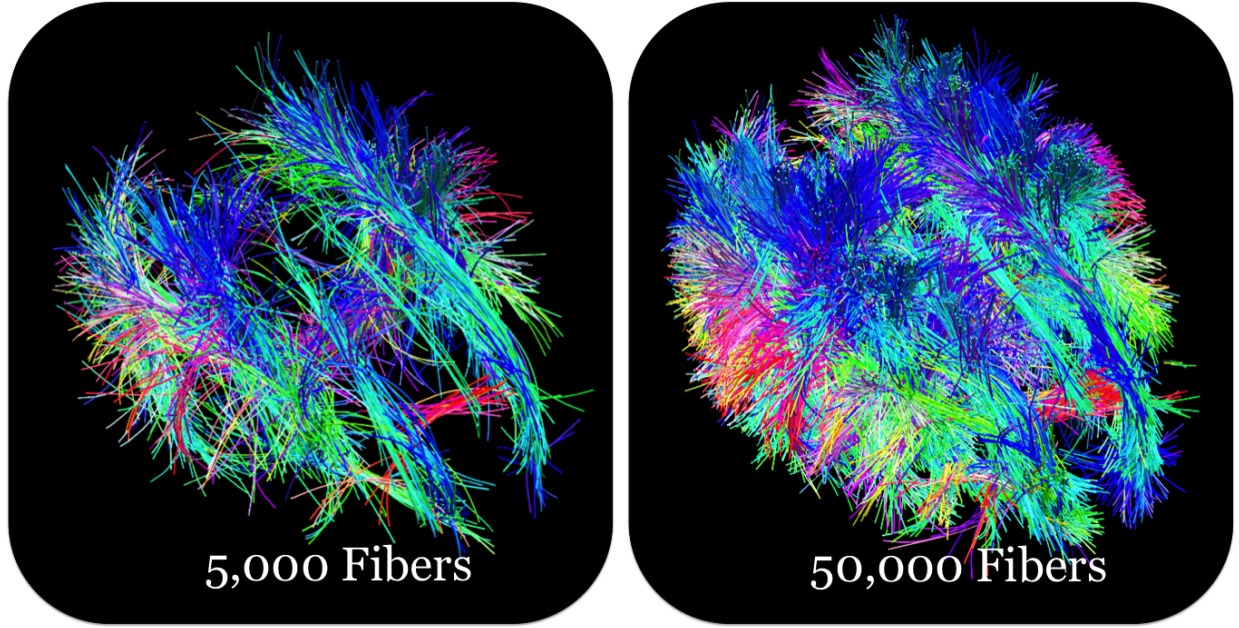
Figure 4.2: Fibers from the optimized Hough transform tractography method. We contrast two different densities of fibers.

efficiency, transitivity, path length, mean degree, density, modularity, small world, and assortativity. We computed global efficiency, transitivity, path length, density, modularity, and small world in both weighted and binary undirected networks. Ten different thresholds were applied to each connectivity matrix, to preserved a fraction of the weights ranging from .1 to 1 by intervals of .1. We note that such a thresholding operation is common in network analysis (cf. work on "k-cores" (Daianu et al., 2012) or Rips filtrations (Lee et al., 2011)), to retain subnetworks that contain the highest density of connections, and eliminate low density connections, to better understand overall network organization.

4.1.2.5 Experiments

For each subject, we computed ten different connectivity matrices, for these matrices we computed a total of fourteen different network measures at ten different weight preservation thresholds. The resulting measures were then used for two-sample t-tests comparing controls vs. eMCI, controls vs. L-MCI, controls vs. AD, and eMCI vs. L-MCI across fiber count and
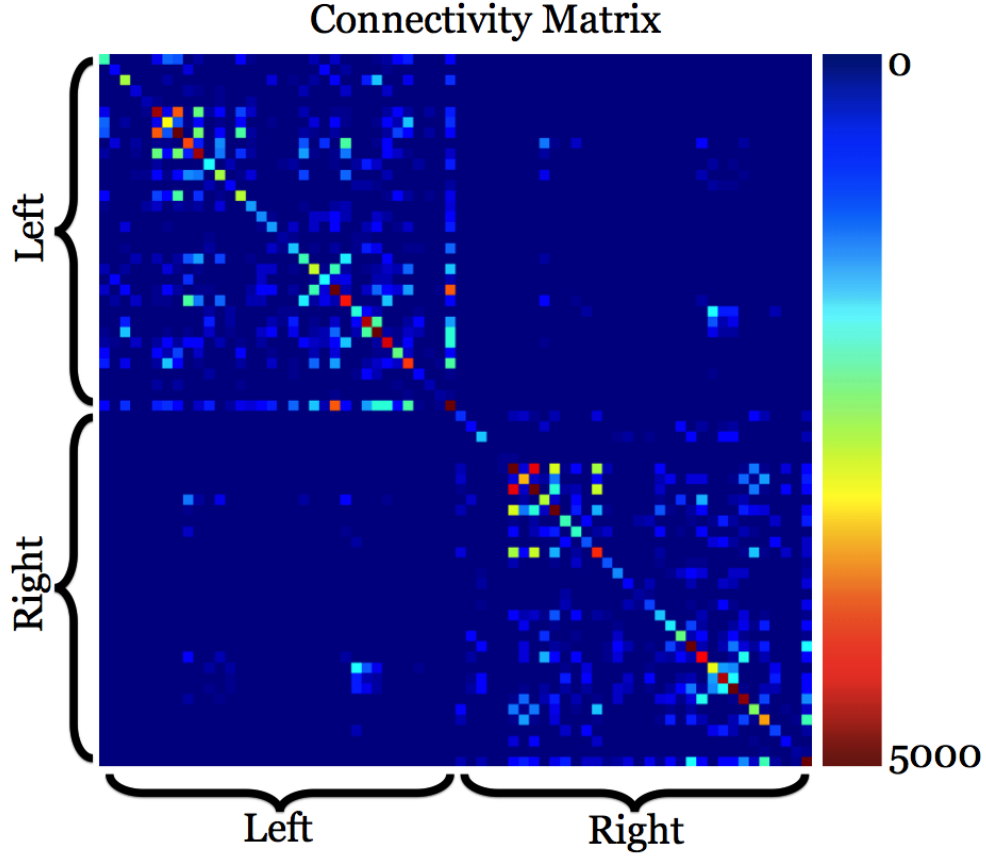
Figure 4.3: A sample connectivity matrix from our subject data showing connections in the left and right hemispheres. The colors signify how many fibers compose each connection.

threshold amount. This setup was intended to show how fiber density influences connectivity network measures.

4.1.3 Results

Fig. 4.4 shows p-values from the t-tests comparing different disease states in our data, as a function of the number of fibers used to create the connectivity matrices. We hypothesized that the group differences would show greater effect sizes (smaller P-values) as fiber density increased. The measures identified in the legend are global efficiency (GT), transitivity (T), path length (PL), mean degree (MD), density (D), modularity (M), small world (SM), and assortativity (A), with (b) and (w) signifying binary and weighted undirected connectivity
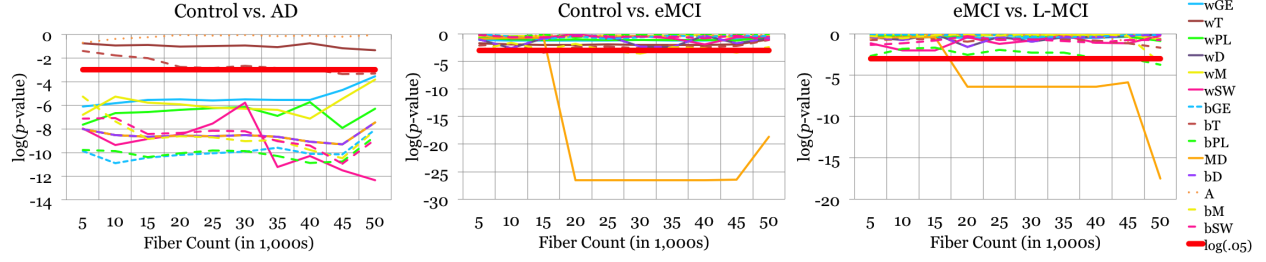
Figure 4.4: p-values from two-sample t-tests comparing controls vs. AD subjects, controls vs. eMCI, and eMCI vs. L-MCI using twelve network measures across different numbers of fibers used to create the connectivity matrix. The best network measure for distinguishing subjects in the eMCI tests was the mean nodal degree (MD). The differences were strong enough to survive multiple comparisons correction for the 12 network measures assessed, and for the search with respect to the number of fibers, and with respect to the network sparsity threshold.

matrices respectively. In the control vs. AD graph we retained 90% (based on searching through thresholds ranging from .1 to 1 at intervals of .1) of the weights for the connectivity matrices used to generate the figure and measures wSW, bSW, bPL, bM, bGE, bD, wPL, wM, wGE easily passed the $p = .05$ threshold. bT was able to pass the threshold when $45,000$ or greater fibers were used to create the connectivity matrix. Assortativity and Transitivity had a difficult time distinguishing the two classes, though binary transitivity passed the 0.05 threshold at $45,000$ fibers.

In the control vs. eMCI part of Fig. 4.4 we shows that bT passed the $p = .05$ threshold at $15,000$ fibers and bM passed it at $30,000$ and $45,000$ while MD becomes very highly significant with $20,000$ fibers or more. A threshold of 20% (this was the best performing threshold based on the ten we searched through) was applied to the connectivity matrix weights. Weighted modularity and path length were the worst performers in this test.

We also show results from eMCI versus L-MCI in Fig. 4.4. In this case, we present results where a 20% threshold is applied. bPL and bM became significant at 50,000 fibers. Again, MD becames highly significant at $20,000$ fibers or greater. Clearly, there is some minimal acceptable density to create an accurate network, but the addition of more samples also reduces noise, further boosting effect sizes. Weighted modularity and binary global efficiency were the worst performers and stable across the fiber counts.

Table 4.1: Most significant results, i.e. highest effect sizes, in our network measure comparison between disease states. C represents controls and % shows the percentage of weights retained in the connectivity matrices. The mean degree of the network was highly effective at distinguishing different phases of the disease.

Test	Measure	p-value	Fibers	%
C vs. AD	wSW	4.28×10^{-6}	50k	90
C vs. eMCI	MD	2.92×10^{-12}	40k	20
C vs. L-MCI	MD	1.17×10^{-3}	50k	20
eMCI vs. L-MCI	MD	2.50×10^{-8}	50k	20

Table 4.1 shows the most significant measure for each test along with its p-value, fiber count, and the optimal threshold used to create the connectivity matrices. It may be that the AD effect is so strong that it can be found no matter what the threshold is, but the MCI effect may be more subtle requiring a judicious choice of settings, in our case a 20% threshold. Mean density performed the best in three of the categories and is the same for weighted and binary connectivity matrices.

4.2 Flow-Based Network Measures of Brain Connectivity in Alzheimer's Disease

4.2.1 Introduction

Connectivity in the brain has commonly been measured from the relative count of tractography fibers connecting different areas. The pathways delineated by the fibers denote the connected structure between areas in the brain, but do not provide any information about the diffusion, or the 'capacity', of the fiber network.

Here we model the brain's anatomical network using whole-brain tractography, but we also compute a measure of the total 'capacity' or flow along these fibers using a flow network representation of our diffusion data. A flow network is defined as a dense connected network where each voxel center is a node, all neighbors of a voxel are connected to it, and the

connection strength of an edge depends on the orientation distribution function (ODF) in the direction of the edge. A natural measure of the 'flow' between regions is embedded into a flow connectivity matrix, which we compared with the standard fiber connectivity matrix to understand what new information flow provides.

Brain connectivity has been modeled from diffusion data using a maximum-flow inspired algorithm in (Zalesky and Fornito, 2009). Our method differs in that we also constrain paths between regions to follow tractography fibers instead of shortest paths between the two regions. The logic is to attempt to follow actual connections rather than other shorter paths through the diffusion functions that have no physical reality. In addition, we used the mean flow along the path instead of the minimum flow, to be more robust to noise in diffusion data. White matter connectivity is also modeled as a maximum flow problem by (O'Donnell et al., 2002), but unlike their optimization, we factor in the structure of the brain and guide paths using whole-brain tractography.

Specifically, we wanted to determine whether the new definition offered better effect sizes in discriminating diseased from normal networks, in the context of aging and Alzheimer's disease.

We apply our method to publicly available structural magnetic resonance imaging (MRI) and diffusion MRI from the Alzheimer's Disease Neuroimaging Initiative (ADNI) (Mueller et al., 2005). We studied the imaging data from 110 subjects: 28 normal healthy controls (C), 11 with late-stage mild cognitive impairment (LMCI), 56 early MCI (eMCI), and 15 AD.

We describe how we preprocess our data to generate whole-brain tractography fibers running between 68 automatically parcellated regions of interest. In parallel, we compute a very dense lattice, or flow graph, from our diffusion data. We project the tractography fibers onto the dense flow graph to form projected paths connecting any two regions of interest. The paths are used to estimate flow between the regions of interest and generate a flow-based connectivity matrix. We compare these matrices to the standard connectivity matrices (based on counting fibers) using network measures and test how well they discriminate

between disease groups in our data. We discuss how inclusion of flow or capacity changes the network landscape of the brain.

4.2.2 Methods
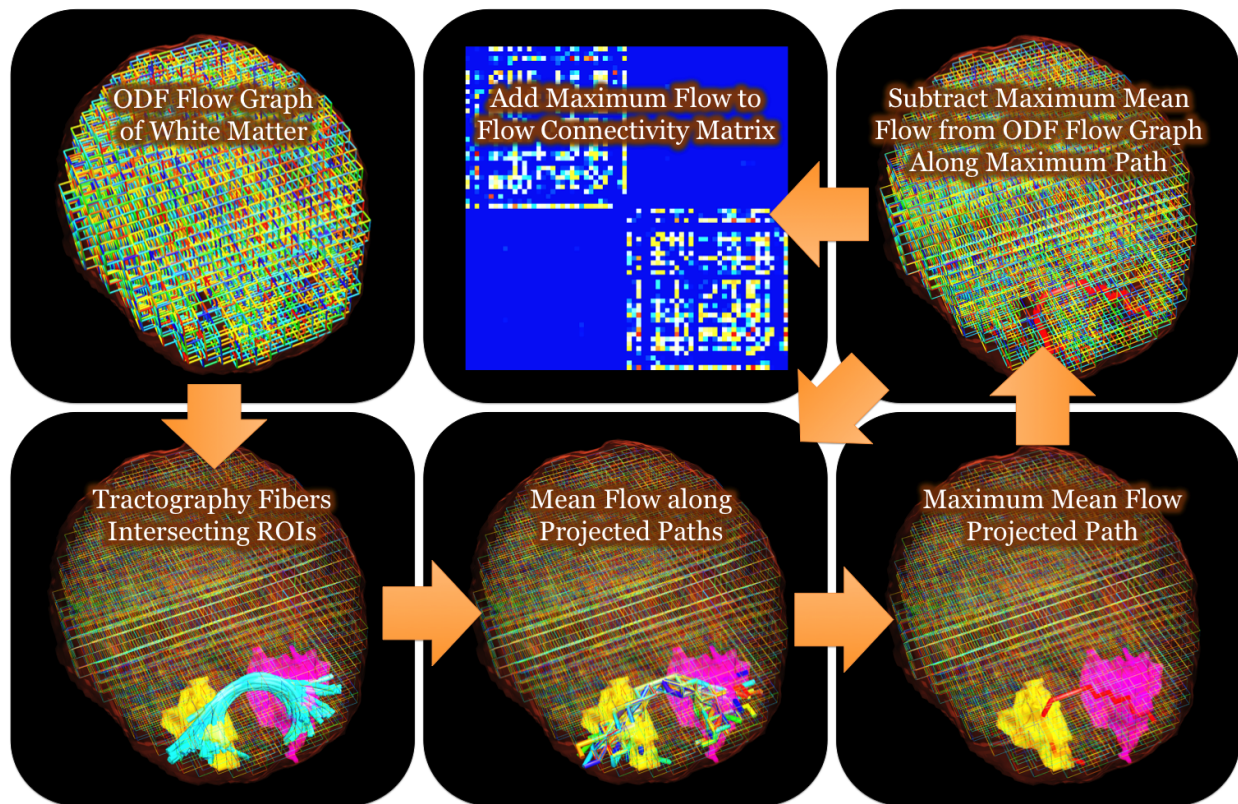
Figure 4.5: We summarize the method to compute a measure of flow between two regions in the brain. A dense flow graph, with all voxel centers as nodes, is created from diffusion data, with edge strengths to all neighbors derived from the orientation distribution function (ODF) evaluated in the edge direction. The graph shown (first panel) is downsampled for clarity; the true version is much denser. Next, fibers from a global probabilistic tractography algorithm are selected that connect two regions of interest. These fibers are projected onto the dense flow graph to form projected paths between the two ROIs and we compute the mean flow along each path. The projected path with the maximum mean flow is selected and we remove its mean flow from its edges in the graph. The maximum mean flow is added to the connectivity matrix and the process is repeated a set number of times.

4.2.2.1 Data

Our data consisted of 110 subjects scanned as part of the ADNI-2 (Trojanowski et al., 2010) collection, a continuation of the ADNI project in which diffusion imaging was added to the standard MRI protocol. The dataset at the time of writing this paper included 28 cognitively normal controls (C), 56 early- and 11 late-stage MCI subjects (eMCI, LMCI), and 15 with Alzheimer's disease (AD).

Subjects were scanned on 3-Tesla GE Medical Systems scanners, which collected both T1-weighted 3D anatomical spoiled gradient echo (SPGR) sequences (256×256 matrix; voxel size = $1.2 \times 1.0 \times 1.0$ mm^3; TI=400 ms; TR = 6.98 ms; TE = 2.85 ms; flip angle = 11°), and diffusion weighted images (DWI; 256×256 matrix; voxel size: $2.7 \times 2.7 \times 2.7$ mm^3; scan time = 9 min). For each subject, the DWI consisted of 41 diffusion images with $b = 1000$ s/mm^2 and 5 T2-weighted b_0 images. This protocol was chosen after an effort to study trade-offs between spatial and angular resolution in a tolerable scan time (Jahanshad et al., 2011).

Image Preprocessing   We processed the T1-weighted images as follows. We first automatically removed extra-cerebral tissues from the anatomical images using ROBEX (Iglesias et al., 2011b), a method that learned from manual segmentations of hundreds of healthy young adults. Skull-stripped brains were inhomogeneity corrected using the MNI N3 tool (Sled et al., 1998) and aligned to the Colin27 template (Holmes et al., 1998) with FSL Flirt (Jenkinson et al., 2002). The resulting images were segmented into 34 cortical regions (in each hemisphere) using FreeSurfer (Fischl et al., 2004). These segmentations were then dilated with an isotropic box kernel of $5 \times 5 \times 5$ voxels to make sure they intersected with the white matter for the subsequent connectivity analysis. We corrected head motion and eddy current distortion in each subject by aligning the DWI images to the average b_0 image with FSL's eddy correct tool. The brain extraction tool (BET) (Smith, 2002) was then used to skull-strip the brains. Further, the images were EPI corrected with an elastic mutual information registration algorithm (Leow et al., 2007) that aligned the DWI images to the T1 scans.

We used a global probabilistic tractography method based on the Hough transform (Aganj et al., 2011). While ADNI's scans are not high angular resolution, due to the need for a fast scan, this method takes advantage of all the diffusion information provided at each voxel, parametrized by the orientation distribution function (ODF). The Hough method generates curves in the fiber space and scores them based on fractional anisotropy (FA) and the ODF at each point along the curve. FA was computed from the single-tensor model of diffusion (Basser and Pierpaoli, 1996). ODFs at each voxel were computed with a recently proposed, normalized and dimensionless estimator derived from Q-ball imaging (QBI) (Aganj et al., 2010). This model is more accurate and outperforms the previous QBI definition (Tuch, 2004), offering better detection of multiple fiber orientations (Aganj et al., 2010; Fritzsche et al., 2010) and additional information for the scoring function. We used the tractography method to generated close to 50,000 fibers per subject.

4.2.2.2 Flow-Based Connectivity

We created a flow graph (Figure 1, first panel) constructed using a densely connected lattice, from our diffusion data to help compute connectivity between regions of the brain; its nodes consist of all the 3D voxel centers lying in the white matter (found using an FA mask with a threshold of 0.2). All nodes in the flow graph are connected to their neighbors with edge weights derived from the orientation distribution function (ODF) in the directions of edges connecting neighboring voxels. We projected fibers generated from tractography (parameterized polynomial curves) onto this lattice graph to find pathways in the lattice between different regions. The mean flow along these projected paths is used to estimate the relative rates of diffusion, or total integrated flow, between any pair of regions in the brain.

ODF Flow Graph   Our flow graph is a set of nodes and edges representing voxel locations and the diffusion or flow between them. These edges may be thought of as pipes or conduits, whose capacity or diameter is equal to the diffusion probability from the ODF. Each node is connected to its surrounding voxels by edges (the 6 neighborhood around the voxel per node).

The edges are weighted according to the value of the ODF at the angle of the neighboring voxel. The constant solid angle (CSA) ODF (Aganj et al., 2010) is used, which includes the Jacobian factor r^2 as

$$\frac{1}{4\pi} + \frac{1}{16\pi^2}\text{FRT}\left\{\triangledown_b^2 \ln\left(-\ln\frac{S(\hat{u})}{S_0}\right)\right\}, \tag{4.5}$$

where $S(\hat{u})$ is the diffusion signal, S_0 is the non-diffusion-weighted image, FRT is the Funk-Radon transform, and $\triangledown_b^2$ is the Laplace-Beltrami operator.

The resulting graph covers the entire white matter region by using an FA mask (0.2 threshold) and connects the previously parcellated cortical regions. We now compute a measure of the amount of flow passing between regions by finding paths between these regions as follows. The classic maximum flow problem has a variety of algorithms (Schrijver, 2002; Ford Jr et al., 1963) to compute flow from the flow network. The Edmonds-Karp algorithm (Edmonds and Karp, 1972) is a widely used method to solve this problem. It works by finding a path between two regions in the graph, using breadth-first search, and subtracting the minimum flow along this path from the graph. The minimum flow or minimum edge constitutes the bottleneck along the path and represents the maximum flow along the path. This process is repeated with the updated graph until no more paths can be found between the two regions. The sum of the minimum flow edges from the paths is the maximum flow between the two regions.

In this paper, we adapt the Edmonds-Karp method (Edmonds and Karp, 1972) to our diffusion data by restricting the paths to follow tractography fibers between the two regions. In addition, we use the mean flow along these paths instead of the minimum flow. Our flow graph can have millions of nodes and edges, leading to an enormous number of possible paths between any two regions, but many of these may not be biologically plausible. Unlike prior methods, we restrict the paths to follow fibers from tractography and restrict our search to n paths per connection instead of all possible paths to estimate the flow in a practical amount of time. If we rely on the minimum flow along a path between two regions, it may falsely restrict the capacity of a path because of noise in our diffusion data (causing a single edge

74

along the path to be abnormally small). To remedy this, we compute the mean flow or mean edge weight along the path as an estimate of the passing flow. We summarize the analysis pipeline in Fig. 4.5.

Tractography Fiber Projection We compute connections between two regions in the flow graph by following paths closest to tractography fibers connecting the two regions, and further select those fibers from our Hough tractography method that intersect the two regions of interest. These fibers are then projected into the flow graph by finding the closest nodes and edges along each fiber. This set of nodes and edges represents a projected path between the two regions in the flow graph.

Flow Between Regions Paths between two regions of interest are then scored based on the their mean edge value or flow. We then select the path with the maximum mean flow and subtract its maximum mean flow from its edges in the flow graph. This modified graph is used to re-score the projected paths between the two regions and the process is repeated for a pre-specified number, n, of iterations. In our experiments, $n = 10$ was sufficient to capture the flow between all regions in our image. The procedure for our method is described as pseudo-code in Algorithm 1.

The above process is repeated for each of the 68 cortical areas to generate a 68×68 flow connectivity matrix (Bullmore and Sporns, 2009). In addition, we computed standard fiber connectivity matrices that count the number of fibers connecting each region. Fig. 4.6 compares the flow-based connectivity matrix and fiber count connectivity matrix in the same subject.

4.2.2.3 Connectivity Measures

We interpreted our connectivity matrices with eight network measures described in (Rubinov and Sporns, 2010), computed with the Brain Connectivity Toolbox. We chose global efficiency, transitivity, path length, mean degree, density, modularity, small world, and assor-

Algorithm 1 A summary of our flow connectivity method

1: **for** each cortical region pair (2 sets of voxels) **do**
2:     Compute ODF graph with FA mask (0.2 threshold)
3:     Select Tractography fibers that intersect both regions
4:     **for** each fiber **do**
5:         Find the set of closest nodes in the graph
6:         Connect these nodes by edges
7:         This set (nodes/edges) is the projected path
8:     **end for**
9:     **for** 1 to n **do**
10:         **for** each projected path **do**
11:             Find its mean edge cost (flow) in the ODF graph
12:         **end for**
13:         Select the path with the maximum mean flow
14:         Add this maximum mean flow to the connectivity matrix
15:         Remove the maximum mean flow from the ODF graph along the maximum mean flow path edges
16:     **end for**
17: **end for**

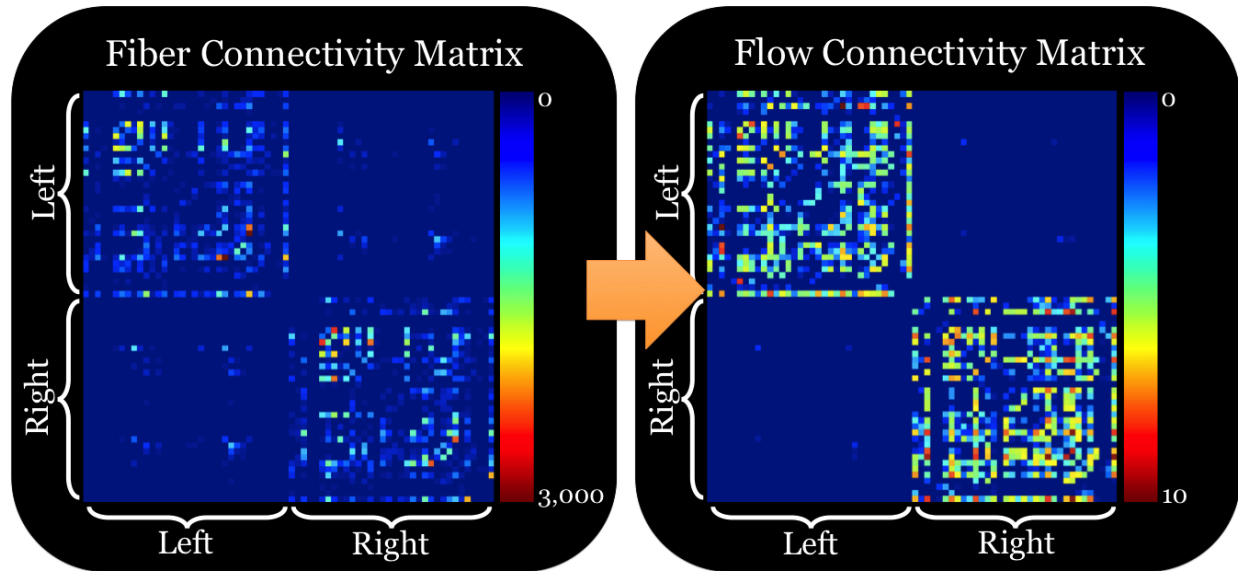Figure 4.6: We show the fiber connectivity matrix and flow connectivity matrix. The flow connectivity matrix may be thought of as a particular weighting of the fiber connectivity matrix from which it is derived.

tativity. We computed all 8 measures from both weighted and binary undirected networks.

As is standard, ten different thresholds were applied to each connectivity matrix, to preserved

76

Table 4.2: Most significant results in our network measure comparison between network measures derived from fiber connectivity matrices and flow connectivity matrices and their ability to discriminate between disease states. The p-values are from two-sample t-tests and the threshold represents the proportion of weights preserved in the connectivity matrices before network measures are computed. % represents the threshold used on the connectivity matrix. Ten separate thresholds were applied to each connectivity matrix that preserved a proportion of the weights ranging from .1 to 1, at intervals of .1.

	Fiber Connectivity			Flow Connectivity		
Test	Measure	p-value (z-score)	%	Measure	p-value (z-score)	%
Control vs. AD	Weighted Small World	$4.28 \times 10^{-6}(4.5)$	.90	Mean Degree	$1.29 \times 10^{-6}(4.7)$	.20
Control vs. eMCI	Mean Degree	$2.92 \times 10^{-12}(6.9)$	.20	Mean Degree	$8.40 \times 10^{-14}(7.4)$	.20
Control vs. late-MCI	Mean Degree	$1.17 \times 10^{-3}(3.0)$	.20	Mean Degree	$1.56 \times 10^{-5}(4.2)$	.20
eMCI vs. late-MCI	Mean Degree	$2.50 \times 10^{-8}(5.5)$	.20	Weighted Small World	$1.16 \times 10^{-2}(2.3)$	.20

a fixed fraction of the weights ranging from .1 to 1, in intervals of .1.

4.2.3 Results

For each subject, we computed both a flow connectivity matrix and a fiber connectivity matrix. For each of these matrices, we computed a total of fourteen (global efficiency, transitivity, path length, density, modularity, and small world were found both in weighted and binary matrices while mean degree and assortativity were found in only binary) different network measures at ten different weight preservation thresholds. The resulting measures were then used in two-sample t-tests comparing controls vs. eMCI, controls vs. LMCI, controls vs. AD, and eMCI vs. LMCI, allowing the network sparsity (threshold) to vary. 2-tailed tests were used, as it is not clear in advance, for these newly defined networks, how some network properties would change with disease (e.g., small world properties). We then analyzed the lowest p-value from each test to compare the discriminative ability of each network. Table 4.2 shows the 'most significant' network measures (i.e., the ones with greatest effect size) from fiber and flow connectivity matrices for each test along with its p-value, and the threshold used to create the connectivity matrices.

4.3 Brain Connectivity and Network Measures in Alzheimer's Disease Identification

4.3.1 Introduction

Magnetic resonance imaging measures of volume, CSF biomarkers, ApoE genotype, age, sex, body mass index, and diagnostic tests have widely been used as features to classify Alzheimer's disease (AD) (Kohannim et al., 2010; Klöppel et al., 2008). We present here results showing the utility of features derived only from brain connectivity in diffusion weighted images of the brain. The features we chose came from standard tractography fiber connectivity between regions and a novel flow based connectivity method. We aimed to understand the information contained in the raw connectivity matrices versus network measures derived from them in differentiating disease states related to Alzheimer's disease. To do this we employed support vector machines (SVMs), a machine learning algorithm for classification, to learn from training data and then classify a separate test set.

SVMs used features from spatial atrophy and white matter alterations was used to classify amnestic MCI in (Cui et al., 2012). They used a technique to rank features and select the top features to use in a 10-fold cross-validation. Our method differs in that we use only measures of connectivity and use principal components analysis to reduce our feature set. Laplacian Regularized Least Squares was used to classify Alzheimer's disease in (Zhang and Shen, 2011) where they tried to incorporate features from mild cognitive impairment into the AD classifier. Cortical thickness features were evaluated in (Eskildsen et al., 2012) using classification though they focused on conversion from MCI to AD.

Our feature extraction and classification was applied to publicly available structural magnetic resonance imaging (MRI) and diffusion MRI from the Alzheimer's Disease Neuroimaging Initiative (ADNI) (Mueller et al., 2005). We studied the imaging data from 110 subjects: 28 normal healthy controls (C), 11 with late-stage mild cognitive impairment (L-MCI), 56 early MCI (eMCI), and 15 AD.

We explain how we extracted connectivity between 68 automatically parcellated regions of interest in the form of connectivity matrices. From these connectivity matrices we computed 28 different network measures. This set of features was reduced in dimensionality using principal components analysis (PCA) and fed into a 10-fold cross-validation design using SVMs to classify C vs. AD, C vs. eMCI, C vs. L-MCI, and eMCI vs. L-MCI. Our results show a significant different in the accuracy of different features sets to distinguish between the two groups.

4.3.2 Methods

4.3.2.1 Data

Our data was from 110 subjects scanned as part of ADNI-2, a continuation of the ADNI project in which diffusion imaging (among other scans) was added to the standard MRI protocol. The dataset as of 11/01/12 included 28 cognitively normal controls (C), 56 early- and 11 late-stage MCI subjects (eMCI, LMCI), and 15 with Alzheimer's disease (AD).

Subjects were scanned on 3-Tesla GE Medical Systems scanners, which collected both T1-weighted 3D anatomical spoiled gradient echo (SPGR) sequences (256×256 matrix; voxel size = $1.2 \times 1.0 \times 1.0$ mm^3; TI=400 ms; TR = 6.98 ms; TE = 2.85 ms; flip angle = $11°$), and diffusion weighted images (DWI; 256×256 matrix; voxel size: $2.7 \times 2.7 \times 2.7$ mm^3; scan time = 9 min). Per subject, the DWI consisted of 41 diffusion images with $b = 1000$ s/mm^2 and 5 T2-weighted b_0 images. This protocol was chosen after an effort to study trade-offs between spatial and angular resolution in a tolerable scan time (Jahanshad et al., 2011).

Image Preprocessing We processed the T1-weighted images to parcellate them into 68 cortical regions. We first automatically removed extra-cerebral tissues from the anatomical images using ROBEX (Iglesias et al., 2011b), a method that learned from manual segmentations of hundreds of healthy young adults. Skull-stripped brains were inhomogeneity corrected using the MNI N3 tool (Sled et al., 1998) and aligned to the Colin27 template

(Holmes et al., 1998) with FSL Flirt (Jenkinson et al., 2002). The resulting images were segmented into 34 cortical regions (in each hemisphere) using FreeSurfer (Fischl et al., 2004). These segmentations were then dilated with an isotropic box kernel of $5 \times 5 \times 5$ voxels to make sure they intersected with the white matter for subsequent connectivity analysis.

We corrected head motion and eddy current distortion in each subject by aligning the DWI images to the average b_0 image with FSL's eddy correct tool. The brain extraction tool (BET) (Smith, 2002) was then used to skull-strip the brains. We EPI corrected these images with an elastic mutual information registration algorithm (Leow et al., 2007) that aligned the DWI images to the T1 scans. Preprocessing steps are further detailed in (Nir et al., 2012).

We used a global probabilistic tractography method based on the Hough transform (Aganj et al., 2011). While ADNIs scans are not high angular resolution, due to the need for a fast scan, this method takes advantage of all the diffusion information provided at each voxel, parametrized by the orientation distribution function (ODF). The Hough method generates curves in the fiber space and scores them based on fractional anisotropy (FA) and the ODF at each point along the curve. FA was computed from the single-tensor model of diffusion (Basser and Pierpaoli, 1996). ODFs at each voxel were computed with a normalized, dimensionless estimator derived from Q-ball imaging (QBI) (Aganj et al., 2010). This model is more accurate and outperforms the previous QBI definition (Tuch, 2004), offering better detection of multiple fiber orientations (Aganj et al., 2010; Fritzsche et al., 2010) and additional information for the scoring function.

To generate close to $50,000$ fibers per subject, we used an accelerated form of this tractography method (Prasad et al., 2013f). Our optimizations included an ODF lookup table and randomized search of the parameter space to generate fibers in less than 1/60 of the original time.

4.3.2.2 Connectivity Features

We used features directly from connectivity matrices of the brain that categorize connection between different regions of interest on the cortex. From these matrices we computed a set of network measures that quantify different network characteristics. We chose different subsets of these features in our experiments to understand the information they captured from the brain.

Connectivity Matrix We compute connectivity matrices using two methods, the first is the relative frequency of fibers connecting regions, and the second is a novel method that computes flow between regions.

Our first method takes fibers computed using the accelerated Hough tractography method and computes the number that intersect pairs of regions from the 68 cortical areas around the brain. We use these frequencies to populate a 68×68 connectivity matrix (with no normalization).

The second method we used a flow based measure of connectivity between all region pairs (Prasad et al., 2013b). In short, we first created a lattice network by connecting all lattice points (voxel centers) to all their neighbors. Edge weights were based on the orientation density function (ODF) value in the direction of the edge. In contrast to counting fibers between ROIs, we computed the maximum flow or capacity between each ROI pair by following connecting tractography fibers projected onto the flow network edges. We utilized a modified maximum-flow algorithm that is robust to noise in the diffusion data and guided by biologically viable pathways and structure of the brain. The resulting flow is used to create a separate 68×68 connectivity matrix. Fig. 4.7 gives an example connectivity matrix using this method.

Network Measures We the two types of connectivity matrices with 28 network measures described in (Rubinov and Sporns, 2010), computed with the Brain Connectivity Toolbox. We computed the following from both weighted and binary connectivity matrices: global ef-

Connectivity Matrix
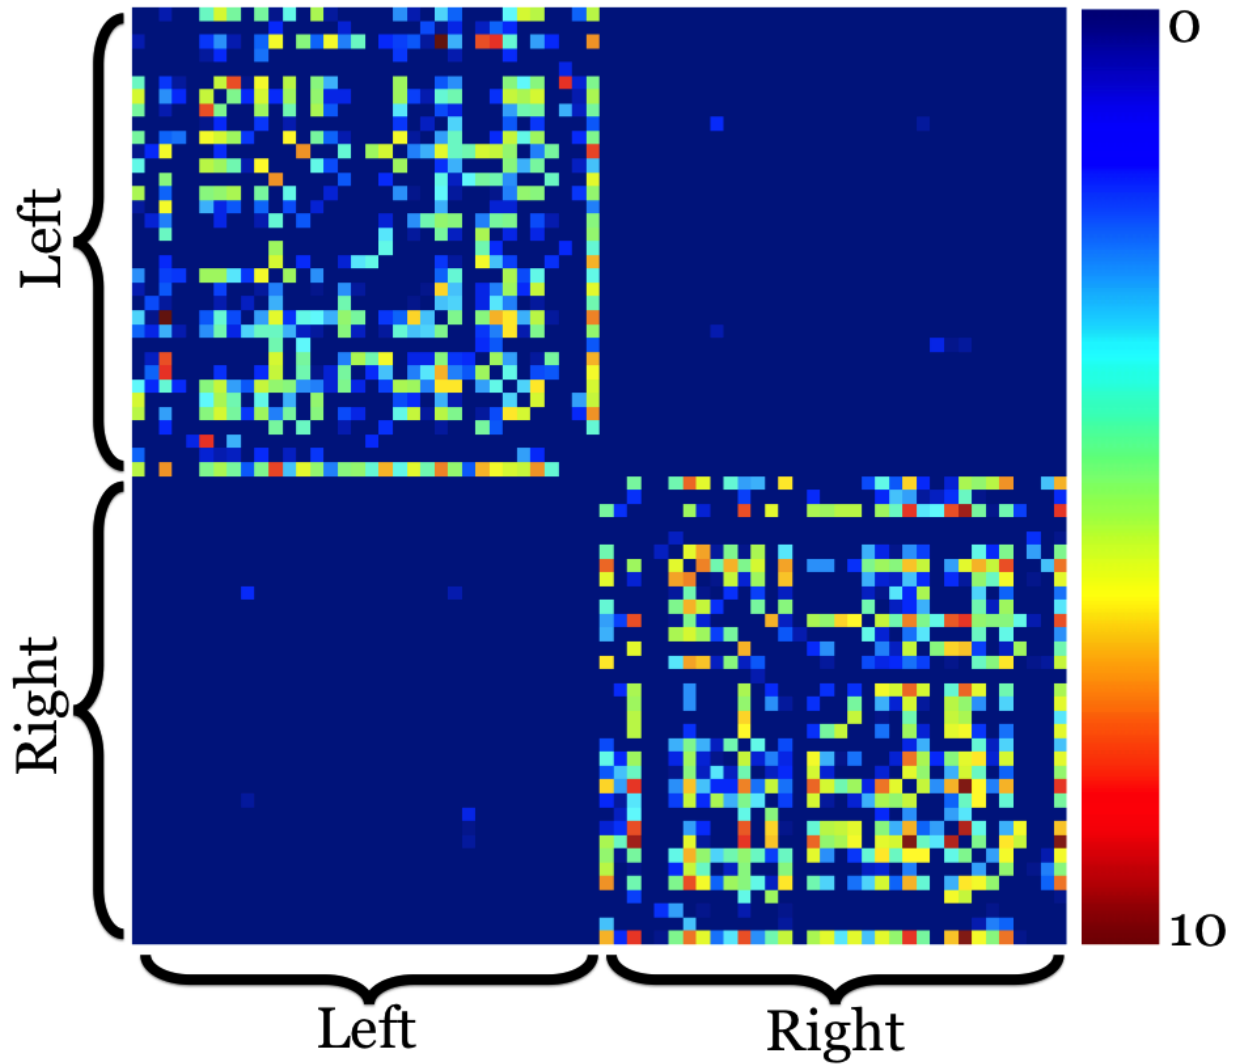
Figure 4.7: An example 68×68 connectivity matrix. This particular matrix was computed using the flow connectivity method that computed the amount of flow between pairs of ROIs segmented in the cortex.

ficiency, transitivity, path length, density, modularity, small world, radius, diameter, number of vertices, number of edges, participation, edge betweenness centrality, nodal betweeness centrality, local efficiency, optimal community structure, eigenvector, subgraph centrality, and eccentricity. In addition, we computed assortativity, nodal flow coefficient, average flow coefficient, total flow across central node, degree, matching index, edge neighborhood

overlap, node pairs degree, and connected component sizes from only binary matrices and strengths from only weighted matrices. As is standard, ten different thresholds were applied to each connectivity matrix, to preserved a fixed fraction of the weights ranging from .1 to 1, in intervals of .1.

If a feature was multi-dimensional we took the mean value in addition to its raw values and vectorized the entire set of features for each subject.

4.3.2.3 Dimensionality Reduction

The combined features from the raw flow and fiber based connectivity matrices with the 28 different network measures for binary and weighted matrices in 10 different thresholds totaled 872,550 features per subject. To synthesize these features into a manageable number for our machine learning classification algorithm we employed dimensionality reduction with principal components analysis (PCA).

Principal Components Analysis (PCA) (Pearson, 1901) is a method that computes the mutually orthogonal directions of maximum variance in a collection of d-dimensional data (in our case $d = 872,550$). These d directions form an orthogonal basis that spans the data space. PCA is usually used for dimensionality reduction, by projecting the data on a subset $\tilde{d} \le d$ of the basis vectors. The subset of vectors are termed the *principal components*, and they account for a large percentage of the variance. They thus represent each d-dimensional data point as a linear combination of the $\tilde{d}$ principal components and, hence, by the $\tilde{d}$-dimensional vector of coefficients of the linear combination. Among its many uses, PCA can be used to compress data, to classify data, to visualize data trends in a lower-dimensional space.

We did principal components analysis of our 110×872550 data matrix representing all the features from our 110 subjects. Most of the variance was accounted for by just a few of the components. In our experiments we chose $\tilde{d}$ to be 10, meaning we used 10 principal components to represent each of our subjects as a 10-dimensional vector. These 10 principal

components represented greater than 99% of the variance in our data.

4.3.2.4   Classification

Support vector machines (SVMs) (Cortes and Vapnik, 1995) is supervised learning method that we used to classify our connectivity features to differentiate between disease states. The method works by learning the best hyperplane that separates the two classes being compared. The hyperplane is chosen to maximize the distance between the hyperplane and the closest data point. Formally, if we represent each subject as a $\tilde{d}$-dimension vector x_i and its category as $y_i = \pm 1$, we can find the hyperplane

$$y_i(<w, x_i> +b) \geq 1 \tag{4.6}$$

such that w and b minimize $||w||$ for all x_i and y_i. In many instances a hyperplane cannot be found to completely separate the two classes of data and slack variables are added to create soft margins to separate most of the data. In addition, SVMs can perform nonlinear classification by applying a nonlinear transformation to the original feature data and computing the hyperplane in this high-dimensional space, which may be nonlinear in the original feature space. This is done efficiently using the "kernel trick", as described in (Boser et al., 1992). We use Gaussian radial basis functions for nonlinear SVM in the classification results presented here.

4.3.2.5   Experiments

We designed experiments to test the utility of different subsets of features to identify between sets of two disease states from our data. Our metric was the accuracy, sensitivity, and specificity from 10-fold cross-validation comparing controls vs. AD, controls vs. early-MCI, controls vs. late-MCI, and early-MCI vs. late-MCI. In each of these classification problems we used nine different sets of features: the fiber connectivity matrix (Fiber(M)), the flow connectivity matrix (Flow(M)), the fiber network measures (Fiber(N)), the flow network measures

(Flow(M)), combinations of these sets as Fiber(N+M), Flow(N+M), Fiber(N)+Flow(N), Fiber(M)+Flow(M), and Fiber(N+M)+Flow(N+M). Each of these sets of features was reduced to ten features using PCA and then fed into the SVM algorithm using a 10-fold cross-validation design. A summary of our experimental design is in Fig. 4.8.
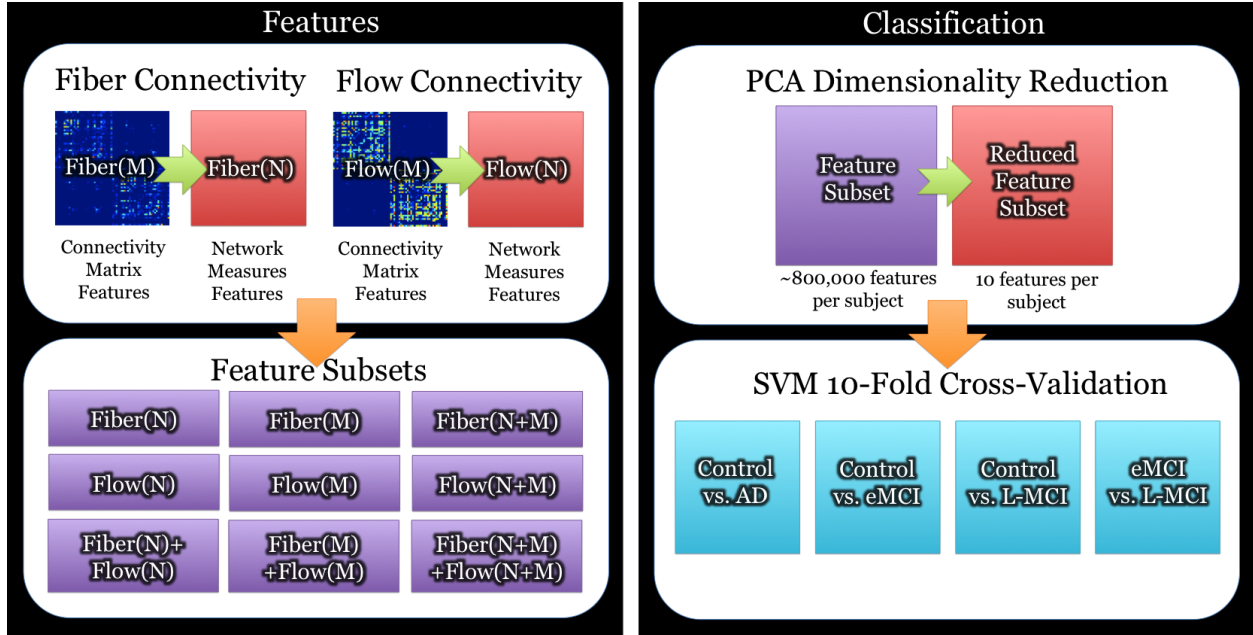

Figure 4.8: A summary of how we select features from brain connectivity and classify them using support vector machines (SVMs).

4.3.3 Results

Table 1 shows the results of each of our four classification problems using the nine different subsets of features. It shows the accuracy, sensitivity, and specificity as percentages.

Table 4.3: Results from support vector machine (SVM) 10-fold cross-validated classification with different sets of features. The features are drawn from sets of network measures (N) and connectivity matrices (M) using a fiber and flow based method. Accuracy (AC), sensitivity (SE), and specificity (SP) are shown as percentages. *Fiber* and *Flow* signify the fiber and flow based connectivity methods. N represents the subset of features composed of network measures derived from the connectivity matrices. M represents the subset of features composed of the raw connectivity matrices values. For each of these subsets of features, we used principal components analysis (PCA) to reduce the dimension of features to a vector of 10 values per subject. The highest accuracy feature combinations are in bold.

	Control vs. AD			Control vs. Early-MCI			Control vs. Late-MCI			Early-MCI vs. Late-MCI		
	AC	SE	SP	AC	SE	SP	AC	SE	SP	AC	SE	SP
Fiber(N)	83.72	85.71	80.00	65.47	64.28	66.07	**79.48**	92.85	45.45	70.14	67.85	81.81
Fiber(M)	**86.04**	85.71	86.66	63.09	60.71	64.28	74.35	75.00	72.72	70.14	71.42	63.63
Fiber(N+M)	83.72	85.71	80.00	65.47	64.28	66.07	**79.48**	92.85	45.45	70.14	67.85	81.81
Flow(N)	**86.04**	85.71	86.66	66.66	67.85	66.07	74.35	78.57	63.63	67.16	75.00	27.27
Flow(M)	79.06	78.57	80.00	**67.85**	67.85	67.85	61.53	67.85	45.45	58.20	66.07	18.18
Flow(N+M)	**86.04**	85.71	86.66	66.66	67.85	66.07	74.35	78.57	63.63	70.14	82.14	9.09
Fiber(N)+Flow(N)	83.72	82.14	86.66	65.47	67.85	64.28	74.35	78.57	63.63	68.65	67.85	72.72
Fiber(M)+Flow(M)	**86.04**	85.71	86.66	63.09	60.71	64.28	74.35	75.00	72.72	**73.13**	80.35	36.36
Fiber(N+M)+Flow(N+M)	83.72	82.14	86.66	65.47	67.85	64.28	74.35	78.57	63.63	68.65	67.85	72.72

4.4 Optimizing Nodes in Brain Connectivity Analyses using Markov Chain Monte Carlo Methods for Alzheimer's Disease Classification

4.4.1 Introduction

Connectomics (Hagmann et al., 2008) – or the study of brain connectivity has become popular in recent years, due to improvements in diffusion imaging and resting-state functional MRI, which reveal neural pathways and functional synchronization between pairs of brain regions. Connectivity in the brain is often characterized by determining a set of connections among a set of N regions in the cortex. In the case of structural connectivity, tractography can be applied to diffusion-weighted MRI data to extract fibers throughout the brain, and the density or number of connections between all pairs of cortical regions can be represented as an $N \times N$ connectivity matrix for each subject in the study (Sporns et al., 2004). This representation of connectivity has been used to further our understanding of aging (Meunier

et al., 2009), brain development, left/right hemisphere differences in connectivity, various diseases, psychiatric disorders, and even genetic variants that affect the brain's neural pathways.

Previously, researchers have proposed several choices for the cortical parcellation particularly for brain connectivity analysis. Tzourio-Mazoyer et al. define the parcellation according to the regions of interests that identify specific functional areas in the brain (Tzourio-Mazoyer et al., 2002), whereas Zalesky et al. (Zalesky et al., 2011) propose an exhaustive approach that treats each voxel as its own ROI, resulting in tens of thousands of ROIs on the cortex. An intermediate approach uses random parcellations of the cortex as nodes (Wig et al., 2011), but this may still fail to capture the true node borders. The meaning of connectivity may depend on the scale of the parcellation, and the most anatomically accurate parcellation might cluster or bundle fibers with similar geometries or trajectories.

While the parcellations based on functional or structural organizations of the brain are useful and widely applied in network connectivity studies, in this paper, we ask, what kinds of cortical parcellations lead to better detection of effects of disease on brain connectivity? The answer depends upon the statistical method used to detect significant effects of disease, and the disease itself, as different diseases may affect some connections more than others. It makes sense to adopt a data-driven approach to estimate the optimal parcellation corresponding to a given disease effect. In this paper, the optimization criteria for the nodes are imposed by the choice of a parcellation that is more sensitive to the classification rate in discriminating disease. The type of networks that we study here are similar to the inference-based networks (He et al., 2007) that rely on the effective connectivity derived from the statistical dependencies of various nodes (regions) on each other, but whose nodes are optimally chosen to improve the classification performance of brain disease. Furthermore, we allow the cortical regions to assume different shapes and sizes. Our design allows optimization of the node set for any connectivity application (it would also be applicable to functional connectivity analyses, or to any connectivity maps constructed from diverse criteria); although in this paper, we use structural connectivity measures to illustrate the

concepts.

As an application, we focus our analysis on the optimal node selection or segmentation for classifying Alzheimer's disease (AD).

4.4.2   Methods

The input data for our method consists of T1-weighted images (See Sec. 4.4.3.2 for more details). We skull-strip the images, and after preprocessing, and correcting for intensity inhomogeneities, register the brains to the Colin27 (Holmes et al., 1998) template. The resulting images are segmented into 34 cortical regions (in each hemisphere) using FreeSurfer (Fischl et al., 2004). These segmentations are then dilated with an isotropic box kernel of $5 \times 5 \times 5$ voxels to make sure they intersect with the white matter for the subsequent connectivity analysis.

4.4.2.1   Sampling Framework

We formulate our problem to assess the ability of the chosen cortical segmentation, s - which dictates the nodes in connectivity matrices that are used in an object function, C, (in our case a classifier to differentiate diseased from normal brains) based on our image data I, which is grouped into a set of patches of the cortex that can be used for the set partition.

$$C(s(p(I, n))), \tag{4.7}$$

where C takes on a value of goodness of fit for a particular set partition s, and p is a randomized patch function that created n random patches of the cortex as the building blocks of the connectivity architecture.

4.4.2.2 Connectivity Representation

The connectivity of the brain may be characterized by a set of nodes and edges (N, E). The nodes represent different cortical regions and the edges denote connections between regions, which can be of varying strength. In our case, we first randomly segmented the cortex into $m = 500$ mutually exclusive patches. We then computed the patches by selecting m random directions, M, and chose a label for each point, β, on the cortex based on the direction that satisfies

$$\operatorname*{argmax}_{\alpha \in M} \left[\left(\left| \frac{\alpha \cdot (\beta - w)}{\|\alpha\| \|(\beta - w)\|} \right| \right) = \cos(\theta) \right], \tag{4.8}$$

where w is the center of the Colin27 template and θ is the angle between the two directions. This process is repeated for each subject for the same random directions so there is a correspondence across the cohort. Figure 4.9 shows an example set of random patches on the cortex.

Random Patches of the Cortex

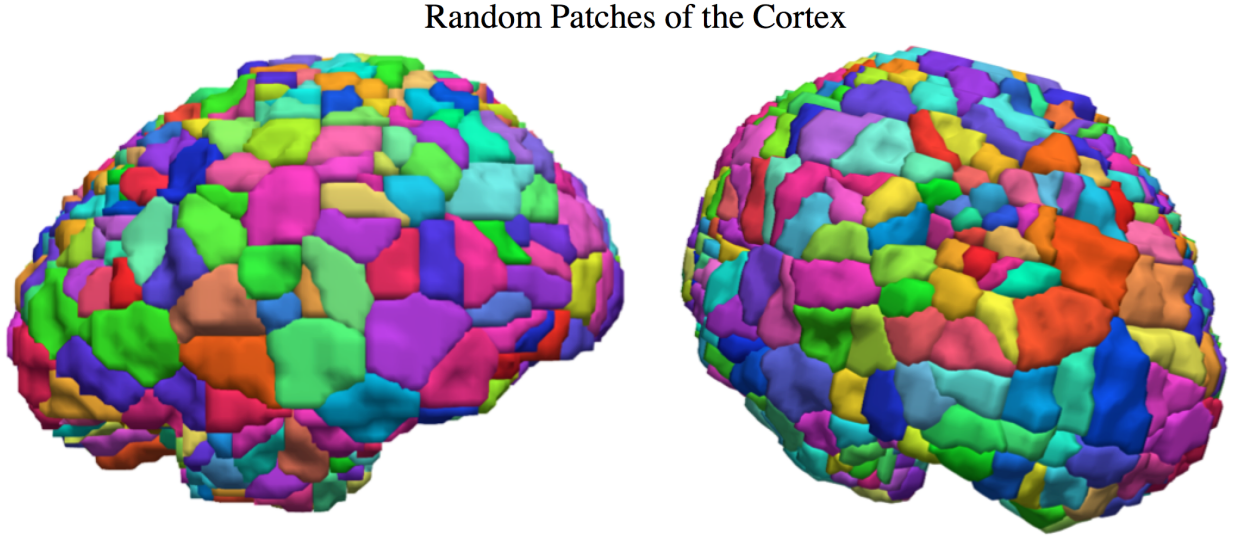Figure 4.9: An example (two frontal views shown) of a parcellation consisting of $m = 500$ random patches on the cortex for a single subject.

For a particular segmentation of the cortex, we also performed tractography to evaluate the edges or connectivity of the nodes by their relationship to the underlying fiber connectivity structure. This is done by finding all combinations of two nodes and counting the

number of fibers that intersect the two areas. We do not normalize this number across all combinations, but keep the number of tractography fibers in each subject constant. The resulting set of nodes and edges is represented by a $N \times N$ connectivity matrix. We compute transitivity, density, node count, edge count, degree, strengths, eigenvector, global efficiency, assortativity, flow coefficient, radius, diameter, betweenness, local efficiency, connected components, subgraph centrality, and eccentricity organized into a feature vector representing the connectivity of each subject.

4.4.2.3 Classification

We use support vector machines (SVMs) (Cortes and Vapnik, 1995) to classify the connectivity features comparing the groups of healthy controls and disease. A SVM is a supervised machine learning algorithm that classifies two-class data by training to find the best hyperplane between the two classes. In addition, we use the nonlinear form of SVMs by transforming features using a radial basis function (RBF). Our classification design tests the information provided by the connectivity configuration in a 10-fold cross validation setting. In each fold, we randomly select an equal number of subjects for each class to avoid bias that can develop from unequal training class sizes (Chew et al., 2001). Finally, we evaluate the results based on accuracy, sensitivity, and specificity. We let $\psi(s(p(I, n))) = C$ represent the classification accuracy for a segmentation/partition s and set of patches p.

4.4.2.4 Markov Monte Carlo Chain Sampling

We represent the space of possible segmentations of the cortex as all possible partitions of the set of m random patches. A partition, s, of a set M composed of m patches is a set of nonempty subsets such that each element of M is in only one of the subsets. This definition means a partition can specify one to m groups/nodes. The total number of partitions of a set of m-elements is the Bell number (Bell, 1934), which may be computed using the recurrence

relations

$$B_m = \sum_{k=0}^{m-1} B_k \binom{m-1}{k}, \tag{4.9}$$

where $\binom{a}{b}$ is the binomial coefficient. When $m = 500$, $B_m = 1.61 \times 10^{843}$ making an exhaustive search of the space infeasible. Instead we opt for a method to generate samples from the posterior distribution and derive our results from this approximation.

Markov Chain Monte Carlo (Andrieu et al., 2003) is a randomized sampling technique to generate samples from a target distribution over time. The Monte Carlo principle approximates a distribution by a set of N independent and identically distributed samples

$$\lim_{N \to \infty} \left[P_N(x) = \frac{1}{N} \sum_{i=1}^{N} 1(x^{(i)} = x) \right] = P(x). \tag{4.10}$$

We used the Metropolis-Hastings algorithm, which is one of the most widely used versions of MCMC. It works by generating a Markov Chain, which represents a sample from the target distribution. Algorithm 1 gives a summary of the method.

Algorithm 2 Metropolis-Hastings Markov Chain Monte Carlo

1: Sample partition $x^{(0)}$ $q(s)$ as the first element of the Markov Chain
2: **for** $i = 0$ to $N - 1$ **do**
3:    Sample s' from the proposal distribution $q(s'|s^{(i)})$
4:    Sample u from the uniform distribution $U(0, 1)$
5:    Compute the acceptance probability as

$$A(x', x^{(i)}) = \min \left\{ 1, \frac{C(s'(p(I, n))q(s^{(i)}|s')}{C(s^{(i)}(p(I, n)))q(s'|s^{(i)})} = \frac{C(s'(p(I, n)))}{C(s^{(i)}(p(I, n)))} = \frac{\psi(s'(p(I, n)))}{\psi(s^{(i)}(p(I, n)))} \right\} \tag{4.11}$$

6:    Choose $s^{(i+1)}$ as

$$s^{(i+1)} = \begin{cases} s' & \text{if } u < A(s', s^{(i)}) \\ s^{(i)} & \text{if } u \geq A(s', s^{(i)}) \end{cases} \tag{4.12}$$

7: **end for**

We select candidate partitions that follow the proposal distribution, q, such that each partition has equal probability of being selected. Our proposal distribution first selects the number of groups in a sample partition based on the probability density function (PDF),

$1/S(m, k)$, over $k = 1, ..., m$, where $S(m, k)$ is the Sterling number of the second kind as

$$S(m, k) = \frac{1}{k!} \sum_{i=1}^{k} (-1)^i \binom{k}{i} (k - i)^m. \tag{4.13}$$

We generated samples from this PDF using a modified version of inverse transform sampling. One we have selected the number of groups, g, we select the elements in each group by randomly assigning a label to each of the m patches we randomly delineated. This is represented as a vector

$$s = [l_1, l_2, l_3, ..., l_m]^T, \tag{4.14}$$

where each l_i is an integer sampled from the uniform distribution $U(1, g)$. The set of labels in s will assign the patches into groups, creating g regions and a segmentation of the cortex. We then apply the specified permutation, s, to our $m \times m$ connectivity matrix, B, using a $m \times g$ matrix, D, that sums together rows and columns of B that belong to the same region or group of patch nodes. It is constructed so that each row contains 1 at the index of elements in a group and 0 otherwise. Each row represents a different group, so that there are g rows total to create a new connectivity matrix

$$B_s = DBD^T. \tag{4.15}$$

If we let $\gamma(s) = g$ specify the number of unique labels or groups in a partition, the resulting transition probability does not depend on the previous state and is characterized as

$$q(s|y) = q(s) = \frac{S(m, \gamma(s))}{B_m} \frac{1}{S(m, \gamma(s))} = \frac{1}{B_m}, \tag{4.16}$$

which is the product of the probability of choosing a particular number of groups with the probability of choosing a partition in the set of partitions of the same number of groups. This means that each partition has the same probability of being chosen and simplifies the acceptance probability of Algorithm 1.

In MCMC we evaluate the a proposed state as $C(s'(p(I, n))) = \psi(s'(p(I, n)))$ in Algorithm 1. We summarize the steps we use to evaluate a proposed segmentation in Algorithm 2.

Algorithm 3 Evaluating Proposed Segmentation

1: Generate the node partition matrix G_{S_i} from proposed label vector s_i
2: **for** $j = 1$ to N (number of subjects) **do**
3:     Compute partitioned connectivity matrix as $B_{s_i}^j = G_{S_i} B^j G_{S_i}^T$
4:     Compute graph measures array, m_i^j from $B_{s_i}^j$
5: **end for**
6: Use the set of m_i^j ($j = 1, ..., N$) in 10-fold cross validated SVM classification
7: Return the accuracy of the classification as $\psi(s'(p(I, n)))$

4.4.3 Results

4.4.3.1 Data

Our data was collected from 43 subjects scanned as part of the ADNI-2 (Trojanowski et al., 2010) project in which diffusion imaging was added to the standard MRI protocol. The dataset included 28 normal controls (C) and 15 Alzheimer's disease (AD) patients.

Subjects were scanned on 3-Tesla GE Medical Systems scanners, which collected both T1-weighted 3D anatomical spoiled gradient echo (SPGR) image volumes, and diffusion-weighted images (DWI; 256×256 matrix; voxel size: $2.7 \times 2.7 \times 2.7$ mm^3; scan time = 9 min). For each subject, the DWI consisted of 41 diffusion-weighted images with $b = 1000$ s/mm^2 and 5 T2-weighted b_0 images. The ADNI data would not be considered high angular resolution, but the protocol was optimized to avoid long scan times.

4.4.3.2 Image Preprocessing

We processed the T1-weighted images as follows. We first removed extra-cerebral tissues from the anatomical images, inhomogeneity corrected and aligned to the Colin27 template with FSL Flirt (Jenkinson et al., 2002), and then processed using Freesurfer for cortical

Table 4.4: This table compares controls vs. AD classification measures between the random patches, Freesurfer, and our method methods for selecting nodes in the connectivity matrix.

Node Selection Method	Accuracy	Sensitivity	Specificity
Random Patches	76.74	71.43	86.67
Freesurfer	72.09	75.00	66.67
Our Method	81.40	78.57	86.67

extraction and labeling (Fischl et al., 2004). For the diffusion weighted images (DWI), we corrected for head motion and eddy current distortion in each subject by aligning the images to the average b_0 image with FSL's eddy correct tool. The images were EPI corrected with an elastic registration algorithm using mutual information that aligned the DWI images to the T1-weighted scans. We used a global probabilistic tractography method based on the Hough transform (Aganj et al., 2011) to generate close to 50,000 fibers per subject.

4.4.3.3 Results

We compare the classification of controls vs. AD results of each node selection method in Table 4.4 using accuracy, sensitivity, and specificity. In the sample generated by our method, the highest performing partition had an accuracy of 81.40%, the highest of the three methods. We generated 3000 samples using our method. This set of samples had a maximum accuracy of 81.40%, minimum 23.26%, mean 51.87%, and variance .75%.

CHAPTER 5

Conclusion

5.1   Brain Segmentation

The use of deformable organisms to segment the brain in whole head human MR images can help in cases where other algorithms include areas of tissue around the eyes and cases where they include parts of the skin. This method could help in difficult cases where those boundaries are hard to delineate. The behaviors within the deformable organisms framework can be customized to perform different specific brain segmentation tasks. If a brain study was sensitive to a particular lobe in the brain then the organisms could be customized so that area had better accuracy during the skull-stripping process.

It also allows the segmentation process to be sensitive to many different types of image information. In our case we process the images using K-Means classification, a threshold, and the gradient. Each organism has its own set of goals and new organisms can build off of what has already been accomplished in the segmentation task. The deformable organisms framework also easily allows the addition of components using different types of images and the use of additional deformations to improve segmentation performance.

The metrics in Table 2.3 suggest that the performance of our deformable organisms approach is comparable to that of other widely used methods. It has the lowest Hausdorff distance average between its automatic results and the gold standard delineations. It did have a higher false positive error, meaning it may include slightly more voxels inside the boundary of the brain.

Table 2.4 shows that learning errors improves the segmentation results; all metrics ex-

amined were improved, especially the Jaccard coefficient. This additional training step may be useful if a large data set needs to be segmented, making it reasonable to segment some images manually for error correction. The method could be trained on a small subset of the manual and automatically segmented data, in a first pass, and the error corrections learned could be useful to segment the rest of the dataset.

Deformable organisms provide an adaptable framework to perform segmentation. They can encode a high-level plan into deformable models, to help them work together to accomplish segmentation tasks. The different control layers may be adapted to fit any type of segmentation tasks.

5.2   White Matter Representation & Registration

In this work, we combined whole-brain HARDI tractography with a standard white matter tract atlas. We were able to automatically group traced fibers in the brain irrespectively of the chosen tractography method. By reconstructing maximum density paths for each individual, we can compactly represent known anatomical fiber tracts in new datasets. This allows us to obtain group statistics for cross-subject comparisons of scalar attributes along these paths (e.g., such as FA, GFA, or mean diffusivity). This method lends itself to population studies as it recovers homologous anatomical tracts across subjects without requiring manual labeling of tracts. It also avoids computationally intensive high-dimensional clustering methods to match corresponding clusters across subjects, which is anyhow not always successful. The threshold and median filter are additional processing steps that do not require manual modification, once established for a specific dataset.

Maximum density paths may be used to compute population statistics for clinical or genetic analyses of connectivity. Features such as curvature along the paths may also be examined. The paths may be registered together across large numbers of subjects to investigate statistical factors associated with differences in tract geometry and fiber characteristics across populations.

The cingulum, anterior limb of the internal capsule, superior longitudinal fasciculus, and anterior corona radiata, were some of the white matter regions with significant differences between the two sexes. The analysis of the Hough-derived fibers analysis showed more significant regions than that of the streamline-derived fibers i.e. the Hough results showed in both hemispheres the cingulum and the superior longitudinal fasciculus, hinting that the greater resolution afforded by HARDI images can help find more significant regions.

We implemented and tested a new pipeline for tract-based analysis of DTI, which uses an atlas to help cluster tracts for point-wise curve-based analysis. Sex differences in the FA of 50 regions of interest delineated in an ROI atlas, suggesting promise of the method for detecting other factors that affect tracts, such as disease and risk genes.

5.3   Connectivity

Our results showed that in some cases a larger number of fibers provided more significant results. Weighted small world in controls vs. AD and mean degree on the MCI comparisons seemed very sensitive to fiber density. The variability in p-values at the 50,000 fiber count may be the result of a decrease in the number of subjects for the test. In a few of the subjects there were less than 50,000 seed points with only 86 of the 110 represented at that fiber count level.

The advantage of more fibers could be higher for catching subtle differences in controls vs. MCI stages or differences between MCI stages. In the case of controls vs. AD, the network may be so different that the network is not influenced by fiber counts as much.

Mean degree stood out among the network measures in differentiating the disease groups; further studies may be useful evaluating other degree-related measures, including subnetworks derived by filtering the networks by nodal degree (such as K-cores).

From Table 4.2, we observe that the network measures from our flow connectivity matrices are able to produce more significant differences (i.e., higher effect sizes) in distinguishing the disease groups when compared to network measures from fiber connectivity matrices in each

test, except for eMCI vs. late-MCI. The selection of mean degree as a discriminative feature was also a reasonable choice, as the progression of Alzheimers disease results in a gradual loss of connectivity between regions. The flow connectivity matrices may be thought of as a weighting of the original fiber connectivity matrices because the procedure of computing flow between a pair of regions, expects an already-existing connection of tractography fibers. The flow measure may help differentiate the disease states relative to simply counting fibers, as it gives a sense of the diffusion along those fibers. Additionally, we conjecture that the flow network measures result in more significant differences (greater effect sizes) due to the additional global flow information retained in the network. To be precise, the flow measure is not measuring the total diffusion as the ODFs are normalized to unit mass. It will reflect myelination as greater myelination would create dominant peaks in the ODFs.

For controls vs. AD Table 1 shows Fiber(M) and Flow(N) having the highest classification accuracy, though when combined with more features the accuracy does go down in some instances. Flow(M) was able to distinguish in controls vs. eMCI the best and Fiber(N) was the best in controls vs. L-MCI. In the eMCI vs. L-MCI experiment a combination of Fiber based features produced the best results.

Different sets of features may uncover detail in the connectivity structure of the brain and be more important for different disease states.

Our method was able to find a segmentation of the cortex that improved the accuracy of the classification. We were able to understand and quantify the distribution of the classification accuracy based on different partitions. In contrast to the approach by Zalesky et al. (Zalesky et al., 2010), that focuses on node selection based on both node resolution and parcellation for connectivity analysis, our method optimizes an objective function in the space of parcellations and nodes. We could use different methods to select nodes; allow a diffeomorphic flow of the cortical parcellation, and use the iterative flow to optimize the classification performance, and use the Jump-Diffusion process for generating the model estimates in MCMC sampling (Srivastava et al., 2002). We can use the current parcellation, and use the merge-split operation to improve the performance. Our method does not restrict the

number, or the shape of the cortical segmentations and the resulting node groups, allowing us to optimize over a much greater space.

5.4   Future Work

The evaluation of our deformable organisms algorithm showed the areas of the brain where the deformable organisms approach needs more tuning and where it provides good results. The ability for each organism to be customized could allow for development of a skull-stripping plan customized for each area of the brain and would be a future way to extend our method.

The low p-value for the test comparing mean FA along the maximum density paths between the Hough based and TrackVis approaches could show that although the overall shape of the maximum density curves is similar, as shown in Figure 3.5, the distribution of FA along the curves is significantly different. This could demonstrate that each method clusters effectively with sufficient internal coherence to capture the differences in the two tractography methods. In addition, the Hough method tended to pick tracts with consistently higher FA. This may be useful future approach in statistical studies, where high FA regions typically offer greater SNR, making it easier to detect effects of specific clinical or genetic factors on fiber microstructure.

In addition, our method enables population analysis of diffusion-weighted images without relying exclusively on global registration of the images into the same space. The method we introduced uses the registration of the ROI atlas to the subject space to roughly initialize seed points which are used to generate maximum density paths derived from fibers from tractography. The MDPs are then put into correspondence through curve registration, allowing us to focus the registration specifically on the white matter region we are interested in without involving the rest of the image. It could be used in subsequent studied to drastically cut down on the computational resources required for an analysis.

Our optimization of the Hough transform tractography method made the large fiber

counts feasible in a reasonable amount of computation time. Future work will focus on more formally bounding the errors in the lookup table and random search compared to the original implementation.

Future work will assess how these types of connectivity networks differ in overall content and organization. Another goal will be to test how much the flow connectivity method depends on the tractography algorithm, e.g., Hough versus streamline methods, and ODF versus tensor-based models. Finally it will be useful to see how this method compares to traditional maximum flow algorithms by using synthetic or simplified data.

Learning algorithms such as SVM, Adaboost, or random forest classification can be sensitive to the feature set used. We chose PCA here to both quickly reduce the features to those that explained the variation in the data, but other schemes can be used and their effects would be useful to categorize in this dataset. In addition, subjects are continuously being added to the ADNI dataset and more training data would give us a stronger understanding of these relationships.

Future work could explore using different proposal distributions or using training data to learn possible distributions that are closer to the underlying data. Other connectivity measures such as functional connectivity could be used with this framework to determine the changes in the nodes of connectivity matrix based on multimodal information. In our experiments we chose the classification rate as our objective function, but the framework allows for other measures or objective mentions to be used in its place. Potential measures could be genetic effects, differences in aging, or changes in development for normal pediatric populations.

Bibliography

Abe, O., Aoki, S., Hayashi, N., Yamada, H., Kunimatsu, A., Mori, H., Yoshikawa, T., Okubo, T., and Ohtomo, K. (2002). Normal aging in the central nervous system: quantitative MR diffusion-tensor analysis. *Neurobiology of Aging*, 23(3):433–441. 43

Aganj, I., Lenglet, C., Jahanshad, N., Yacoub, E., Harel, N., Thompson, P., and Sapiro, G. (2011). A Hough transform global probabilistic approach to multiple-subject diffusion MRI tractography. *Medical Image Analysis*, 15(4):414–425. 2, 7, 37, 45, 61, 73, 80, 94

Aganj, I., Lenglet, C., Sapiro, G., Yacoub, E., Ugurbil, K., and Harel, N. (2010). Reconstruction of the orientation distribution function in single-and multiple-shell Q-ball imaging within constant solid angle. *Magnetic Resonance in Medicine*, 64(2):554–566. 1, 36, 45, 62, 73, 74, 80

Andersson, J., Smith, S., and Jenkinson, M. (2008). FNIRT-FMRIB's non-linear image registration tool. *Human Brain Mapping*, pages 15–19. 1

Andrews-Hanna, J., Snyder, A., Vincent, J., Lustig, C., Head, D., Raichle, M., and Buckner, R. (2007). Disruption of large-scale brain systems in advanced aging. *Neuron*, 56(5):924. 5

Andrieu, C., De Freitas, N., Doucet, A., and Jordan, M. (2003). An introduction to MCMC for machine learning. *Machine Learning*, 50(1):5–43. 91

Ardekani, B., Guckemus, S., Bachman, A., Hoptman, M., Wojtaszek, M., and Nierenberg, J. (2005). Quantitative comparison of algorithms for inter-subject registration of 3D volumetric brain MRI scans. *Journal of Neuroscience Methods*, 142(1):67–76. 3

Ashburner, J. (2007). A fast diffeomorphic image registration algorithm. *NeuroImage*, 38(1):95–113. 1

Basser, P., Mattiello, J., and Le Bihan, D. (1994). Estimation of the effective self-diffusion tensor from the NMR spin echo. *Journal of Magnetic Resonance-Series B*, 103(3):247–254. 1

Basser, P., Pajevic, S., Pierpaoli, C., Duda, J., and Aldroubi, A. (2000). In vivo fiber tractography using DT-MRI data. *Magnetic Resonance in Medicine*, 44(4):625–632. 2

Basser, P. and Pierpaoli, C. (1996). Microstructural and physiological features of tissues elucidated by quantitative-diffusion-tensor MRI. *Journal of Magnetic Resonance: Series B*, 111(3):209–219. 5, 43, 45, 62, 73, 80

Bell, E. (1934). Exponential numbers. *American Mathematical Monthly*, pages 411–419. 90

Boesen, K., Rehm, K., Schaper, K., Stoltzner, S., Woods, R., L
"uders, E., and Rottenberg, D. (2004). Quantitative comparison of four brain extraction algorithms. *NeuroImage*, 22(3):1255–1261. 22

Boser, B., Guyon, I., and Vapnik, V. (1992). A training algorithm for optimal margin classifiers. In *Proceedings of the Fifth Annual workshop on Computational Learning Theory*, pages 144–152. ACM. 84

Bovenkamp, E., Dijkstra, J., Bosch, J., and Reiber, J. (2004). Multi-agent segmentation of IVUS images. *Pattern Recognition*, 37(4):647–663. 13

Bozzali, M., Falini, A., Franceschi, M., Cercignani, M., Zuffi, M., Scotti, G., Comi, G., and Filippi, M. (2002). White matter damage in Alzheimer's disease assessed in vivo using diffusion tensor magnetic resonance imaging. *Journal of Neurology, Neurosurgery & Psychiatry*, 72(6):742–746. 59

Brun, A., Knutsson, H., Park, H., Shenton, M., and Westin, C. (2004). Clustering fiber traces using normalized cuts. *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2004*, 7:368–375. 43

Bullmore, E. and Sporns, O. (2009). Complex brain networks: graph theoretical analysis of structural and functional systems. *Nature Reviews Neuroscience*, 10(3):186–198. 2, 8, 75

Chew, H., Bogner, R., and Lim, C. (2001). Dual ν-support vector machine with error rate and training size biasing. In *IEEE International Conference on Acoustics, Speech, and Signal Processing*, volume 2, pages 1269–1272. 90

Cook, P., Zhang, H., Avants, B., Yushkevich, P., Alexander, D., Gee, J., Ciccarelli, O., and Thompson, A. (2005). An automated approach to connectivity-based partitioning of brain structures. *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2005*, pages 164–171. 60

Corouge, I., Fletcher, P., Joshi, S., Gilmore, J., and Gerig, G. (2005). Fiber tract-oriented statistics for quantitative diffusion tensor MRI analysis. *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2005*, pages 131–139. 41

Corouge, I., Fletcher, P., Joshi, S., Gouttard, S., Gerig, G., et al. (2006). Fiber tract-oriented statistics for quantitative diffusion tensor MRI analysis. *Medical Image Analysis*, 10(5):786–798. 5, 44

Cortes, C. and Vapnik, V. (1995). Support-vector networks. *Machine Learning*, 20(3):273–297. 84, 90

Cui, Y., Wen, W., Lipnicki, D., Beg, M., Jin, J., Luo, S., Zhu, W., Kochan, N., Reppermund, S., Zhuang, L., et al. (2012). Automated detection of amnestic mild cognitive impairment in community-dwelling elderly adults: A combined spatial atrophy and white matter alteration approach. *NeuroImage*, 59(2):1209–1217. 78

Cuingnet, R., Gerardin, E., Tessieras, J., Auzias, G., Lehéricy, S., Habert, M., Chupin, M., Benali, H., and Colliot, O. (2011). Automatic classification of patients with Alzheimer's disease from structural MRI: A comparison of ten methods using the ADNI database. *NeuroImage*, 56(2):766–781. 2

Daianu, M., Jahanshad, N., Nir, T., Toga, A., C.R. Jack, J., Weiner, M., Thompson, P., and the ADNI (2012). Analyzing the structural k-core of brain connectivity networks in normal aging and alzheimer's disease. In *Medical Image Computing and Computer Assisted Intervention (MICCAI) - NIBAD*. 66

Dale, A., Fischl, B., and Sereno, M. (1999). Cortical Surface-Based Analysis* 1:: I. Segmentation and Surface Reconstruction. *NeuroImage*, 9(2):179–194. 11

De Zubicaray, G., Chiang, M., McMahon, K., Shattuck, D., Toga, A., Martin, N., Wright, M., and Thompson, P. (2008). Meeting the challenges of neuroimaging genetics. *Brain Imaging and Behavior*, 2(4):258–263. 36

Dean, J. and Ghemawat, S. (2008). MapReduce: simplified data processing on large clusters. *Communications of the ACM*, 51(1):107–113. 64

Dennis, E., Jahanshad, N., Rudie, J., Brown, J., Johnson, K., McMahon, K., de Zubicaray, G., Montgomery, G., Martin, N., and Wright, M. (2011). Altered structural brain connectivity in healthy carriers of the autism risk gene, CNTNAP2. *Brain Connectivity*, 1(6):447–459. 6

Desbrun, M., Meyer, M., Schroder, P., and Barr, A. (1999). Implicit fairing of irregular meshes using diffusion and curvature flow. In *Proceedings of the 26th Annual Conference on Computer Graphics and Interactive Techniques*, pages 317–324. ACM Press/Addison-Wesley Publishing Co. 23

Dijkstra, E. (1959). A note on two problems in connexion with graphs. *Numerische Mathematik*, 1(1):269–271. 36, 40, 47

Edmonds, J. and Karp, R. (1972). Theoretical improvements in algorithmic efficiency for network flow problems. *Journal of the ACM (JACM)*, 19(2):248–264. 74

Eskildsen, S., Coupé, P., García-Lorenzo, D., Fonov, V., Pruessner, J., Collins, L., et al. (2012). Improving prediction of Alzheimer's disease using patterns of cortical thinning

and homogenizing images according to disease stage. In *MICCAI 2012 Workshop on Novel Biomarkers for Alzheimer's Disease and Related Disorders*, pages 79–90. 78

Fair, D., Dosenbach, N., Church, J., Cohen, A., Brahmbhatt, S., Miezin, F., Barch, D., Raichle, M., Petersen, S., and Schlaggar, B. (2007). Development of distinct control networks through segregation and integration. *Proceedings of the National Academy of Sciences*, 104(33):13507–13512. 6

Fennema-Notestine, C., Ozyurt, I., Clark, C., Morris, S., Bischoff-Grethe, A., Bondi, M., Jernigan, T., Fischl, B., Segonne, F., Shattuck, D., et al. (2006). Quantitative evaluation of automated skull-stripping methods applied to contemporary and legacy images: effects of diagnosis, bias correction, and slice location. *Human Brain Mapping*, 27(2):99–113. 12

Fillard, P., Gilmore, J., Piven, J., Lin, W., and Gerig, G. (2003). Quantitative analysis of white matter fiber properties along geodesic paths. *Medical Image Computing and Computer-Assisted Intervention-MICCAI 2003*, pages 16–23. 41

Fischl, B., Salat, D., Busa, E., Albert, M., Dieterich, M., Haselgrove, C., van der Kouwe, A., Killiany, R., Kennedy, D., Klaveness, S., et al. (2002). Whole brain segmentation: automated labeling of neuroanatomical structures in the human brain. *Neuron*, 33(3):341–355. 1

Fischl, B., Sereno, M., Dale, A., et al. (1999). Cortical surface-based analysis. *NeuroImage*, 9(2):195–207. 3

Fischl, B., Van Der Kouwe, A., Destrieux, C., Halgren, E., Ségonne, F., Salat, D., Busa, E., Seidman, L., Goldstein, J., Kennedy, D., et al. (2004). Automatically parcellating the human cerebral cortex. *Cerebral Cortex*, 14(1):11–22. 61, 72, 80, 88, 94

Ford Jr, L., Fulkerson, D., and Ziffer, A. (1963). Flows in networks. *Physics Today*, 16:54. 74

Freund, Y. and Schapire, R. (1995). A decision-theoretic generalization of on-line learning and an application to boosting. In *Computational Learning Theory*, pages 23–37. Springer. 30

Fritzsche, K., Laun, F., Meinzer, H., and Stieltjes, B. (2010). Opportunities and pitfalls in the quantification of fiber integrity: What can we gain from Q-ball imaging? *NeuroImage*, 51(1):242–251. 45, 62, 73, 80

Gaffney, S. and Smyth, P. (2005). Joint probabilistic curve clustering and alignment. *Advances in neural information processing systems*, 17:473–480. 35

Gering, D., Nabavi, A., Kikinis, R., Hata, N., O'Donnell, L., Grimson, W., Jolesz, F., Black, P., and Wells III, W. (2001). An integrated visualization system for surgical planning and guidance using image fusion and an open MR. *Journal of Magnetic Resonance Imaging*, 13:967–975. 11, 21

Gharachorloo, N., Gupta, S., Sproull, R., and Sutherland, I. (1989). A characterization of ten rasterization techniques. *ACM SIGGRAPH Computer Graphics*, 23(3):355–368. 24

Hagmann, P., Cammoun, L., Gigandet, X., Meuli, R., Honey, C., Wedeen, V., and Sporns, O. (2008). Mapping the structural core of human cerebral cortex. *PLoS biology*, 6(7):e159. 86

Hagmann, P., Kurant, M., Gigandet, X., Thiran, P., Wedeen, V., Meuli, R., and Thiran, J. (2007). Mapping human whole-brain structural networks with diffusion MRI. *PLoS One*, 2(7):e597. 6

Hamarneh, G. and McIntosh, C. (2005). Physics-based deformable organisms for medical image analysis. In *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, volume 5747, pages 326–335. 21

He, Y., Chen, Z., and Evans, A. (2007). Small-world anatomical networks in the human brain revealed by cortical thickness from MRI. *Cerebral Cortex*, 17(10):2407. 87

Holmes, C., Hoge, R., Collins, L., Woods, R., Toga, A., and Evans, A. (1998). Enhancement of MR images using registration for signal averaging. *Journal of Computer Assisted Tomography*, 22(2):324–333. 38, 50, 61, 72, 80, 88

Holzapfel, M., Barnea-Goraly, N., Eckert, M., Kesler, S., and Reiss, A. (2006). Selective alterations of white matter associated with visuospatial and sensorimotor dysfunction in turner syndrome. *The Journal of Neuroscience*, 26(26):7007–7013. 5

Huttenlocher, D., Klanderman, G., and Rucklidge, W. (1993). Comparing images using the Hausdorff distance. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pages 850–863. 31

Ibanez, L., Schroeder, W., Ng, L., and Cates, J. (2005). The ITK software guide. *Kitware Inc.* 21

Iglesias, J., Jiang, J., Liu, C., and Tu, Z. (2011a). Classification of Alzheimer's disease using a self-smoothing operator. *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2011*, pages 58–65. 2

Iglesias, J., Liu, C., Thompson, P., and Tu, Z. (2011b). Robust brain extraction across datasets and comparison with publicly available methods. *Medical Imaging, IEEE Transactions on*, 30(9):1617–1634. 61, 72, 79

Jahanshad, N., Aganj, I., Lenglet, C., Joshi, A., Jin, Y., Barysheva, M., McMahon, K., de Zubicaray, G., Martin, N., Wright, M., et al. (2011). Sex differences in the human connectome: 4-Tesla high angular resolution diffusion imaging (HARDI) tractography in 234 young adult twins. In *Biomedical Imaging: From Nano to Macro, 2011 IEEE International Symposium on*, pages 939–943. IEEE. 6, 34, 61, 72, 79

Jahanshad, N., Lee, A., Barysheva, M., McMahon, K., de Zubicaray, G., Martin, N., Wright, M., Toga, A., and Thompson, P. (2010). Genetic influences on brain asymmetry: A DTI study of 374 twins and siblings. *NeuroImage*, 52(2):455–469. 38, 43

Jenkinson, M., Bannister, P., Brady, M., and Smith, S. (2002). Improved optimization for the robust and accurate linear registration and motion correction of brain images. *NeuroImage*, 17(2):825–841. 61, 72, 80, 93

Jenkinson, M. and Smith, S. (2001). A global optimisation method for robust affine registration of brain images. *Medical Image Analysis*, 5(2):143–156. 28, 38, 46

Jones, D., Simmons, A., Williams, S., and Horsfield, M. (1999). Non-invasive assessment of axonal fiber connectivity in the human brain via diffusion tensor MRI. *Magnetic Resonance in Medicine*, 42(1):37–41. 5

Joshi, A., Lepore, N., Joshi, S., Lee, A., Barysheva, M., Stein, J., McMahon, K., Johnson, K., de Zubicaray, G., Martin, N., et al. (2011). The contribution of genes to cortical thickness and volume. *NeuroReport*, 22(3):101. 32

Joshi, A., Shattuck, D., and Leahy, R. (2012). A method for automated cortical surface registration and labeling. *Biomedical Image Registration*, pages 180–189. 1

Joshi, S., Klassen, E., Srivastava, A., and Jermyn, I. (2007a). A novel representation for Riemannian analysis of elastic curves in Rn. In *2007 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–7. 48

Joshi, S., Klassen, E., Srivastava, A., and Jermyn, I. (2007b). Removing shape-preserving transformations in square-root elastic (SRE) framework for shape analysis of curves. In *Energy Minimization Methods in Computer Vision and Pattern Recognition*, pages 387–398. 48

Klein, A., Andersson, J., Ardekani, B., Ashburner, J., Avants, B., Chiang, M., Christensen, G., Collins, D., Gee, J., Hellier, P., et al. (2009). Evaluation of 14 nonlinear deformation algorithms applied to human brain MRI registration. *NeuroImage*, 46(3):786. 1, 11, 31

Klöppel, S., Stonnington, C., Chu, C., Draganski, B., Scahill, R., Rohrer, J., Fox, N., Jack,

C., Ashburner, J., and Frackowiak, R. (2008). Automatic classification of MR scans in Alzheimer's disease. *Brain*, 131(3):681–689. 7, 78

Kohannim, O., Hua, X., Hibar, D., Lee, S., Chou, Y., Toga, A., Jack, C., Weiner, M., and Thompson, P. (2010). Boosting power for clinical trials using classifiers based on multiple biomarkers. *Neurobiology of Aging*, 31(8):1429–1442. 7, 78

Langton, C. and Shimohara, T. (1997). *Artificial Life V: Proceedings of the Fifth International Workshop on the Synthesis and Simulation of Living Systems*, volume 5. MIT Press. 3

Lee, H., Chung, M., Kang, H., Kim, B., and Lee, D. (2011). Discriminative persistent homology of brain networks. In *Biomedical Imaging: From Nano to Macro, 2011 IEEE International Symposium on*, pages 841–844. IEEE. 66

Leow, A., Yanovsky, I., Chiang, M., Lee, A., Klunder, A., Lu, A., Becker, J., Davis, S., Toga, A., and Thompson, P. (2007). Statistical properties of Jacobian maps and the realization of unbiased large-deformation nonlinear image registration. *Medical Imaging, IEEE Transactions on*, 26(6):822–832. 61, 72, 80

Lorensen, W. and Cline, H. (1987). Marching cubes: A high resolution 3D surface construction algorithm. *ACM Siggraph Computer Graphics*, 21(4):163–169. 23

Maddah, M., Zollei, L., Grimson, W., and Wells, W. (2008). Modeling of anatomical information in clustering of white matter fiber trajectories using dirichlet distribution. In *Computer Vision and Pattern Recognition Workshops, 2008. CVPRW'08. IEEE Computer Society Conference on*, pages 1–7. IEEE. 43

McInerney, T., Hamarneh, G., Shenton, M., and Terzopoulos, D. (2002). Deformable organisms for automatic medical image analysis. *Medical Image Analysis*, 6(3):251–266. 12, 21

McInerney, T. and Terzopoulos, D. (1996). Deformable models in medical image analysis. In *Mathematical Methods in Biomedical Image Analysis, 1996., Proceedings of the Workshop on*, pages 171–180. IEEE. 3, 21

McIntosh, C. and Hamarneh, G. (2006a). I-DO: A Deformable Organisms framework for ITK. *Medical Image Analysis.* 21

McIntosh, C. and Hamarneh, G. (2006b). Spinal crawlers: Deformable organisms for spinal cord segmentation and analysis. *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2006*, pages 808–815. 21

McIntosh, C. and Hamarneh, G. (2006c). Vessel crawlers: 3D physically-based deformable organisms for vasculature segmentation and analysis. In *Computer Vision and Pattern Recognition, 2006 IEEE Computer Society Conference on*, volume 1, pages 1084–1091. 21

McIntosh, C., Hamarneh, G., and Mori, G. (2007). Human limb delineation and joint position recovery using localized boundary models. In *Motion and Video Computing, 2007. WMVC'07. IEEE Workshop on*, pages 31–31. IEEE. 21

Mettler, F. (1964). Connections of the Cerebral Cortex. *Archives of Neurology*, 10(6):637. 34

Meunier, D., Achard, S., Morcom, A., and Bullmore, E. (2009). Age-related changes in modular organization of human brain functional networks. *NeuroImage*, 44(3):715–723. 86

Mori, S., Crain, B., Chacko, V., and Van Zijl, P. (1999). Three-dimensional tracking of axonal projections in the brain by magnetic resonance imaging. *Annals of Neurology*, 45(2):265–269. 36

Morra, J., Tu, Z., Apostolova, L., Green, A., Toga, A., and Thompson, P. (2008). Automatic subcortical segmentation using a contextual model. *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2008*, pages 194–201. 3

Mueller, S., Weiner, M., Thal, L., Petersen, R., Jack, C., Jagust, W., Trojanowski, J., Toga, A., and Beckett, L. (2005). Ways toward an early diagnosis in Alzheimer's disease: The Alzheimer's disease neuroimaging initiative (ADNI). *Alzheimer's and Dementia*, 1(1):55–66. 59, 70, 78

Nimsky, C., Ganslandt, O., Hastreiter, P., Wang, R., Benner, T., Sorensen, A., and Fahlbusch, R. (2005). Preoperative and intraoperative diffusion tensor imaging-based fiber tracking in glioma surgery. *Neurosurgery*, 56(1):130–138. 34

Nir, T., Jahanshad, N., Toga, A., Jack, C., Weiner, M., and Thompson, P. (2012). Connectivity network breakdown predicts imminent volumetric atrophy in early mild cognitive impairment. *Multimodal Brain Image Analysis*, pages 41–50. 61, 80

O'Donnell, L., Haker, S., and Westin, C. (2002). New approaches to estimation of white matter connectivity in diffusion tensor MRI: Elliptic PDEs and geodesics in a tensor-warped space. *Medical Image Computing and Computer-Assisted Intervention—MICCAI 2002*, pages 459–466. 70

O'Donnell, L., Kubicki, M., Shenton, M., Dreusicke, M., Grimson, W., and Westin, C. (2006). A method for clustering white matter fiber tracts. *American Journal of Neuroradiology*, 27(5):1032–1036. 35, 43

O'Donnell, L. and Westin, C. (2005). White matter tract clustering and correspondence in populations. *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2005*, 8:140–147. 43

O'Donnell, L. and Westin, C. (2007). Automatic tractography segmentation using a high-dimensional white matter atlas. *Medical Imaging, IEEE Transactions on*, 26(11):1562–1575. 35

Ohtake, Y., Belyaev, A., and Bogaevski, I. (2000). Polyhedral surface smoothing with simultaneous mesh regularization. In *Geometric Modeling and Processing 2000. Theory and Applications. Proceedings*, pages 229–237. IEEE. 23

Pearson, K. (1901). On lines and planes of closest fit to systems of points in space. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, 2(11):559–572. 83

Peled, S., Gudbjartsson, H., Westin, C., Kikinis, R., and Jolesz, F. (1998). Magnetic resonance imaging shows orientation and asymmetry of white matter fiber tracts. *Brain Research*, 780(1):27–33. 43

Pineda, J. (1988). A parallel algorithm for polygon rasterization. *ACM SIGGRAPH Computer Graphics*, 22(4):17–20. 24

Powell, H., Parker, G., Alexander, D., Symms, M., Boulby, P., Wheeler-Kingshott, C., Barker, G., Noppeney, U., Koepp, M., and Duncan, J. (2006). Hemispheric asymmetries in language-related pathways: a combined functional MRI and tractography study. *NeuroImage*, 32(1):388–399. 34

Prasad, G., Jahanshad, N., Aganj, I., Lenglet, C., Sapiro, G., Toga, A., and Thompson, P. (2011a). Atlas-based fiber clustering for multi-subject analysis of high angular resolution diffusion imaging tractography. In *Biomedical Imaging: From Nano to Macro, 2011 IEEE International Symposium on*, pages 276–280. 2, 5

Prasad, G., Joshi, A., Feng, A., Barysheva, M., Mcmahon, K., De Zubicaray, G., Martin, N., Wright, M., Toga, A., Terzopoulos, D., and Thompson, P. (2011b). Deformable organisms and error learning for brain segmentation. In Pennec, X., Joshi, S., and Nielsen, M., editors, *Proceedings of the Third International Workshop on Mathematical Foundations of Computational Anatomy - Geometrical and Statistical Methods for Modelling Biological Shape Variability*, Toronto, Canada. 2, 3

Prasad, G., Joshi, A., Thompson, P., Toga, A., Shattuck, D., and Terzopoulos, D. (2011c). Skull-stripping with deformable organisms. In *Biomedical Imaging: From Nano to Macro, 2011 IEEE International Symposium on*, pages 1662–1665. 2, 3, 21

Prasad, G., Joshi, S., Jahanshad, N., Reina, J., Aganj, I., Lenglet, C., Sapiro, G., McMahon, K., de Zubicaray, G., Martin, N., et al. (2012). Genetic analysis of fibers in white matter pathways from HARDI images. In *Organization for Human Brain Mapping (OHBM 2012)*, Beijing, China. 2, 5

Prasad, G., Joshi, S., Jahanshad, N., Villalon, J., Aganj, I., Lenglet, C., Sapiro, G., McMahon, K., de Zubicaray, G., Martin, N., et al. (2011d). White matter tract analysis in 454 adults using maximum density paths. In *Proceedings of the MICCAI Workshop on Computational Diffusion MRI (CDMRI 2011)*, pages 1–12, Toronto, Canada. 2, 5

Prasad, G., Joshi, S., Nir, T., Toga, A., and Thompson, P. (2013a). Brain connectivity based on maximum flow in Alzheimer's disease: The EMFATIC method. In *Organization for Human Brain Mapping (OHBM 2013)*, Seattle, USA, submitted. 2, 7

Prasad, G., Joshi, S., Nir, T., Toga, A., and Thompson, P. (2013b). Flow-based network measures of brain connectivity in Alzheimer's disease. In *Biomedical Imaging: From Nano to Macro, 2013 IEEE International Symposium on.* 2, 7, 81

Prasad, G., Joshi, S., Nir, T., Toga, A., and Thompson, P. (2013c). Machine learning for connectivity-based Alzheimer's disease classification. In *Organization for Human Brain Mapping (OHBM 2013)*, Seattle, USA, submitted. 2, 7

Prasad, G., Joshi, S., Nir, T., Toga, A., and Thompson, P. (2013d). Optimizing nodes in brain connectivity analyses using markov chain monte carlo methods for Alzheimer's disease classification. In *Information Processing in Medical Imaging 2013*, Asilomar, USA, submitted. 2, 8

Prasad, G., Joshi, S., Nir, T., Toga, A., and Thompson, P. (2013e). Refining brain connectivity networks to optimally identify brain disease: The EPIC algorithm. In *Organization for Human Brain Mapping (OHBM 2013)*, Seattle, USA, submitted. 2, 8

Prasad, G., Nir, T., Toga, A., and Thompson, P. (2013f). Fiber density and connectivity in

Alzheimer's disease. In *Organization for Human Brain Mapping (OHBM 2013)*, Seattle, USA, submitted. 2, 7, 80

Prasad, G., Nir, T., Toga, A., and Thompson, P. (2013g). Tractography density and network measures in Alzheimer's disease. In *Biomedical Imaging: From Nano to Macro, 2013 IEEE International Symposium on.* 2, 7

Resnick, S., Pham, D., Kraut, M., Zonderman, A., and Davatzikos, C. (2003). Longitudinal magnetic resonance imaging studies of older adults: a shrinking brain. *Journal of Neuroscience*, 23(8):3295–3301. 11, 21

Roche, A., Malandain, G., Pennec, X., and Ayache, N. (1998). The correlation ratio as a new similarity measure for multimodal image registration. *Medical Image Computing and Computer-Assisted Intervention–MICCAI 1998*, page 1115. 28

Rose, S., Chen, F., Chalk, J., Zelaya, F., Strugnell, W., Benson, M., Semple, J., and Doddrell, D. (2000). Loss of connectivity in Alzheimer's disease: an evaluation of white matter tract integrity with colour coded MR diffusion tensor imaging. *Journal of Neurology, Neurosurgery & Psychiatry*, 69(4):528–530. 5, 59

Rubinov, M. and Sporns, O. (2010). Complex network measures of brain connectivity: uses and interpretations. *Neuroimage*, 52(3):1059–1069. 7, 65, 75, 81

Schmitt, J., Lenroot, R., Wallace, G., Ordaz, S., Taylor, K., Kabani, N., Greenstein, D., Lerch, J., Kendler, K., Neale, M., et al. (2008). Identification of genetically mediated cortical networks: a multivariate study of pediatric twins and siblings. *Cerebral Cortex*, 18(8):1737–1747. 6

Schrijver, A. (2002). On the history of the transportation and maximum flow problems. *Mathematical Programming*, 91(3):437–445. 74

Segonne, F., Dale, A., Busa, E., Glessner, M., Salat, D., Hahn, H., Fischl, B., et al. (2004).

A hybrid approach to the skull stripping problem in MRI. *NeuroImage*, 22(3):1060–1075. 3, 12, 20, 22

Shattuck, D. and Leahy, R. (2002). BrainSuite: an automated cortical surface identification tool. *Medical Image Analysis*, 6(2):129–142. 3, 22

Shattuck, D., Mirza, M., Adisetiyo, V., Hojatkashani, C., Salamon, G., Narr, K., Poldrack, R., Bilder, R., and Toga, A. (2008). Construction of a 3D probabilistic atlas of human cortical structures. *NeuroImage*, 39(3):1064–1080. 19, 23, 32

Shattuck, D., Prasad, G., Mirza, M., Narr, K., and Toga, A. (2009). Online resource for validation of brain segmentation methods. *NeuroImage*, 45(2):431–439. 19, 24

Shattuck, D., Sandor-Leahy, S., Schaper, K., Rottenberg, D., and Leahy, R. (2001). Magnetic resonance image tissue classification using a partial volume model. *NeuroImage*, 13(5):856–876. 12, 20

Shimony, J., Snyder, A., Lori, N., and Conturo, T. (2002). Automated fuzzy clustering of neuronal pathways in diffusion tensor tracking. In *Proc. Intl. Soc. Mag. Reson. Med*, volume 10. 43

Shin, Y., Kim, D., Hyon, T., Park, H., Moon, W., Chung, E., Lee, J., Kim, I., Kim, S., and Kwon, J. (2005). Sex differences in the human corpus callosum: diffusion tensor imaging study. *Neuroreport*, 16(8):795–798. 5

Skudlarski, P., Jagannathan, K., Calhoun, V., Hampson, M., Skudlarska, B., Pearlson, G., et al. (2008). Measuring brain connectivity: diffusion tensor imaging validates resting state temporal correlations. *NeuroImage*, 43(3):554–561. 6

Sled, J., Zijdenbos, A., and Evans, A. (1998). A nonparametric method for automatic correction of intensity nonuniformity in MRI data. *Medical Imaging, IEEE Transactions on*, 17(1):87–97. 61, 72, 79

Smith, S. (2002). Fast robust automated brain extraction. *Human Brain Mapping*, 17(3):143–155. 3, 12, 20, 22, 38, 50, 61, 72, 80

Smith, S., Jenkinson, M., Johansen-Berg, H., Rueckert, D., Nichols, T., Mackay, C., Watkins, K., Ciccarelli, O., Cader, M., Matthews, P., et al. (2006). Tract-based spatial statistics: voxelwise analysis of multi-subject diffusion data. *NeuroImage*, 31(4):1487–1505. 5, 44

Sporns, O., Chialvo, D., Kaiser, M., and Hilgetag, C. (2004). Organization, development and function of complex brain networks. *Trends in Cognitive Sciences*, 8(9):418–425. 86

Sporns, O., Tononi, G., and Kötter, R. (2005). The human connectome: a structural description of the human brain. *PLoS Computational Biology*, 1(4):e42. 2

Srivastava, A., Grenander, U., Jensen, G., and Miller, M. (2002). Jump–diffusion markov processes on orthogonal groups for object pose estimation. *Journal of Statistical Planning and Inference*, 103(1):15–37. 98

Stam, C., Jones, B., Nolte, G., Breakspear, M., and Scheltens, P. (2007). Small-world networks and functional connectivity in Alzheimer's disease. *Cerebral Cortex*, 17(1):92–99. 6

Steels, L. (1993). The artificial life roots of artificial intelligence. *Artificial Life*, 1(1_2):75–110. 21

Stein, J., Medland, S., Vasquez, A., Hibar, D., Senstad, R., Winkler, A., Toro, R., Appel, K., Bartecek, R., Bergmann, Ø., et al. (2012). Identification of common variants associated with human hippocampal and intracranial volumes. *Nature Genetics*, 44(5):552–561. 2

Storey, J. and Tibshirani, R. (2003). Statistical significance for genomewide studies. *Proceedings of the National Academy of Sciences of the United States of America*, 100(16):9440. 51

Takahashi, S., Yonezawa, H., Takahashi, J., Kudo, M., Inoue, T., and Tohgi, H. (2002). Selective reduction of diffusion anisotropy in white matter of Alzheimer's disease brains

measured by 3.0 Tesla magnetic resonance imaging. *Neuroscience Letters*, 332(1):45–48. 43

Terzopoulos, D., Platt, J., Barr, A., and Fleischer, K. (1987). Elastically deformable models. *ACM Siggraph Computer Graphics*, 21(4):205–214. 21

Thompson, P., Hayashi, K., De Zubicaray, G., Janke, A., Rose, S., Semple, J., Herman, D., Hong, M., Dittmer, S., Doddrell, D., et al. (2003). Dynamics of gray matter loss in Alzheimer's disease. *Journal of Neuroscience*, 23(3):994–1005. 11, 21

Thompson, P., Hayashi, K., Sowell, E., Gogtay, N., Giedd, J., Rapoport, J., De Zubicaray, G., Janke, A., Rose, S., Semple, J., et al. (2004). Mapping cortical change in Alzheimer's disease, brain development, and schizophrenia. *NeuroImage*, 23:S2–S18. 21

Trojanowski, J., Vandeerstichele, H., Korecka, M., Clark, C., Aisen, P., Petersen, R., Blennow, K., Soares, H., Simon, A., Lewczuk, P., et al. (2010). Update on the biomarker core of the Alzheimer's disease neuroimaging initiative subjects. *Alzheimer's & dementia: The Journal of the Alzheimer's Association*, 6(3):230. 72, 93

Tu, Z., Narr, K., Dollár, P., Dinov, I., Thompson, P., and Toga, A. (2008). Brain anatomical structure segmentation by hybrid discriminative/generative models. *Medical Imaging, IEEE Transactions on*, 27(4):495–508. 1

Tuch, D. (2004). Q-ball imaging. *Magnetic Resonance in Medicine*, 52(6):1358–1372. 1, 37, 43, 45, 62, 73, 80

Tuch, D., Belliveau, J., and Wedeen, V. (2000). A path integral approach to white matter tractography. In *Proceedings of the 8th Annual Meeting of ISMRM, Denver*, page 791. 6

Tuch, D., Reese, T., Wiegell, M., Makris, N., Belliveau, J., and Wedeen, V. (2002). High angular resolution diffusion imaging reveals intravoxel white matter fiber heterogeneity. *Magnetic Resonance in Medicine*, 48(4):577–582. 2

Tustison, N., Avants, B., Cook, P., Kim, J., Whyte, J., Gee, J., and Stone, J. (2012). Logical circularity in voxel-based analysis: Normalization strategy may induce statistical bias. *Human Brain Mapping*. 5

Tyan, S. (1981). Median filtering: Deterministic properties. *Two-Dimensional Digital Signal Processing II*, 43:197–217. 39

Tzourio-Mazoyer, N., Landeau, B., Papathanassiou, D., Crivello, F., Etard, O., Delcroix, N., Mazoyer, B., Joliot, M., et al. (2002). Automated anatomical labeling of activations in SPM using a macroscopic anatomical parcellation of the MNI MRI single-subject brain. *NeuroImage*, 15(1):273–289. 8, 87

Visser, E., Nijhuis, E., Buitelaar, J., and Zwiers, M. (2011). Partition-based mass clustering of tractography streamlines. *NeuroImage*, 54(1):303–312. 35, 43

Wakana, S., Jiang, H., Nagae-Poetscher, L., van Zijl, P., and Mori, S. (2004). Fiber tract–based atlas of human white matter anatomy. *Radiology*, 230(1):77–87. 38

Wang, H., Das, S., Suh, J., Altinay, M., Pluta, J., Craige, C., Avants, B., and Yushkevich, P. (2011). A learning-based wrapper method to correct systematic errors in automatic image segmentation: Consistently improved performance in hippocampus, cortex and brain segmentation. *NeuroImage*. 22, 30, 35, 51

Wang, R., Benner, T., Sorensen, A., and Wedeen, V. (2007). Diffusion toolkit: A software package for diffusion imaging data processing and tractography. *Intl Soc Mag Reson Med 15*, 15:3720. 36, 37

Wassermann, D., Kanterakis, E., Gur, R., Deriche, R., and Verma, R. (2010). Diffusion-based population statistics using tract probability maps. *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2010*, 13:631–639. 43

Wig, G., Schlaggar, B., and Petersen, S. (2011). Concepts and principles in the analysis of brain networks. *Annals of the New York Academy of Sciences*, 1224(1):126–146. 8, 87

Woods, R., Dapretto, M., Sicotte, N., Toga, A., and Mazziotta, J. (1999). Creation and use of a Talairach-compatible atlas for accurate, automated, nonlinear intersubject registration, and analysis of functional imaging data. *Human Brain Mapping*, 8(2-3):73–79. 11, 21

Zalesky, A. and Fornito, A. (2009). A DTI-derived measure of cortico-cortical connectivity. *Medical Imaging, IEEE Transactions on*, 28(7):1023–1036. 70

Zalesky, A., Fornito, A., Egan, G., Pantelis, C., and Bullmore, E. (2011). The relationship between regional and inter-regional functional connectivity deficits in schizophrenia. *Human Brain Mapping*. 2, 87

Zalesky, A., Fornito, A., Harding, I., Cocchi, L., Yücel, M., Pantelis, C., and Bullmore, E. (2010). Whole-brain anatomical networks: does the choice of nodes matter? *NeuroImage*, 50(3):970. 98

Zhang, D. and Shen, D. (2011). Semi-supervised multimodal classification of Alzheimer's disease. In *Biomedical Imaging: From Nano to Macro, 2011 IEEE International Symposium on*, pages 1628–1631. IEEE. 78

Zhang, Y., Brady, M., and Smith, S. (2001). Segmentation of brain MR images through a hidden markov random field model and the expectation-maximization algorithm. *Medical Imaging, IEEE Transactions on*, 20(1):45–57. 3

Zöllei, L., Stevens, A., Huber, K., Kakunoori, S., and Fischl, B. (2010). Improved tractography alignment using combined volumetric and surface registration. *NeuroImage*, 51(1):206–213. 5