

Multimodal Fusion for Assessing Functional Segregation and Integration in the Human Brain

by

Burak Yoldemir

B.Sc., Bogazici University, 2009

M.Sc., Bogazici University, 2011

A THESIS SUBMITTED IN PARTIAL FULFILLMENT OF
THE REQUIREMENTS FOR THE DEGREE OF

DOCTOR OF PHILOSOPHY

in

The Faculty of Graduate and Postdoctoral Studies

(Electrical and Computer Engineering)

THE UNIVERSITY OF BRITISH COLUMBIA

(Vancouver)

February 2016

© Burak Yoldemir 2016

Abstract

Mental and neurological diseases account for a major portion of the global disease burden. Neuroimaging has greatly contributed to the characterization and understanding of such disorders by enabling noninvasive investigation of the human brain. Among neuroimaging technologies, magnetic resonance imaging (MRI) stands out as a relatively widespread and safe imaging modality that can be sensitized to capture different aspects of the brain. Historically, MRI studies have investigated anatomy or function of the brain in isolation, which created an apparent dichotomy. In this thesis, we aim to bridge this divide using novel multimodal techniques. In particular, we present techniques to reconcile information regarding anatomical and functional connectivity (AC and FC) in the brain estimated from diffusion MRI (dMRI) and functional MRI (fMRI) data, respectively. Our first contribution is to show that the consistency between AC and FC is understated when standard analysis methods are used. We illustrate how the estimation of AC can be improved to increase the AC-FC consistency, which facilitates a more meaningful fusion of these two types of information. Specifically, we propose to improve AC estimation by the use of a dictionary based super-resolution approach to increase the spatial resolution in dMRI, reconstructing the white matter tracts using global tractography instead of conventional streamline tractography, and quantifying AC using fiber count as the metric. Our second contribution is to develop novel multimodal approaches for investigating functional segregation and integration in the human brain. We show that task fMRI data can be fused with dMRI and resting state fMRI data to mitigate the effects of noise and deconfound the effects of spontaneous fluctuations in brain activity on activation detection. Further, we show that sensitivity in unraveling the modular structure of the brain can be increased by fusing dMRI and fMRI data. Our results collectively suggest that combining dMRI and fMRI data outperforms classical unimodal analyses in understanding the brain's organization, bringing us one step closer to understanding the most complex organ in the human body.

Preface

The research presented herein was approved by the UBC Clinical Research Ethics Board, certificate number: H15-02343. This thesis is primarily based on the following papers, resulting from collaboration between multiple researchers.

Studies described in Chapter 2 have been published in:

- [P1] B. Yoldemir, B. Ng, R. Abugharbieh. Effects of tractography approach on consistency between anatomical and functional connectivity estimates. In Proceedings of the IEEE International Symposium on Biomedical Imaging (ISBI'14), pp. 250-253, Beijing, China, April 2014. Oral presentation - Acceptance rate: $\sim 20\%$.
- [P2] B. Yoldemir, M. Bajammal, R. Abugharbieh. Dictionary based super-resolution for diffusion MRI. In Proceedings of the MICCAI Workshop on Computational Diffusion MRI, pp. 194-204, Boston, MA, September 2014.
- [P3] M. Bajammal, B. Yoldemir, R. Abugharbieh. Comparison of structural connectivity metrics for multimodal brain image analysis. In Proceedings of the IEEE International Symposium on Biomedical Imaging (ISBI'15), pp. 934-937, Brooklyn, NY, USA, April 2015.

Studies described in Chapter 3 have been published in:

- [P4] B. Yoldemir, B. Ng, R. Abugharbieh. Deconfounding the effects of resting state activity on task activation detection in fMRI. In Proceedings of the MICCAI Workshop on Multimodal Brain Image Analysis, LNCS, vol. 7509, pp. 51-60, Nice, France, October 2012.
- [P5] B. Yoldemir, B. Ng, T. Woodward, R. Abugharbieh. Fiber connectivity integrated brain activation detection. In Proceedings of the 23rd Biennial International Conference on Information Processing in Medical Imaging (IPMI'13), LNCS, vol. 7917, pp. 135-146, Asilomar, CA, June 2013. Acceptance rate: 32.1%.

Studies described in Chapter 4 have been published in:

- [P6] B. Yoldemir, B. Ng, R. Abugharbieh. Overlapping replicator dynamics for functional subnetwork identification. In Proceedings of the International Conference on Medical Image Computing and Computer Assisted Intervention (MICCAI'13), LNCS, vol. 8150, pp. 682-689, Nagoya, Japan, September 2013. Acceptance rate: 32.8%.
- [P7] B. Yoldemir, B. Ng, R. Abugharbieh. Coupled stable overlapping replicator dynamics for multimodal brain subnetwork identification. In Proceedings of the 24th Biennial International Conference on Information Processing in Medical Imaging (IPMI'15), LNCS, vol. 9123, pp. 770-781, Isle of Skye, Scotland, June 2015. Acceptance rate: 32.3%.
- [P8] B. Yoldemir, B. Ng, R. Abugharbieh. Stable overlapping replicator dynamics for brain community detection. IEEE Transactions on Medical Imaging, vol. 35, no. 2, pp. 529-538, February 2016. Impact factor: 3.39, h5-index: 63.

[P1, P4, P6-8]: I was the primary author of these papers and main contributor to the design, implementation and testing of the proposed methods under the supervision of Dr. Rafeef Abugharbieh. Dr. Bernard Ng helped with his valuable input in improving the methodology and design of the experiments. The papers were edited by all co-authors.

[P2]: I contributed the application idea conception, preprocessing of the data and parcellation, and the validation scheme. Mohammad Bajammal contributed the algorithmic idea conception, implementation of the method and validation scheme, as well as generation of the results. I contributed the majority of the effort for manuscript writing. The paper was edited by all co-authors.

[P3]: I contributed the paper idea, preprocessing of the data and parcellation, and implementation of the tractography method and three of the four anatomical connectivity metrics. Mohammad Bajammal contributed the implementation of the fourth anatomical connectivity metric and the validation code as well as generated the results. I and Mohammad Bajammal contributed equally to writing the paper. The paper was edited by all co-authors.

Preface

[P5]: I was the primary author of this paper and main contributor to the design, implementation and testing of the proposed method under the supervision of Dr. Rafeef Abugharbieh. Dr. Bernard Ng assisted in improving the methodology and design of the experiments. Dr. Todd Woodward provided data and clinical input. The paper was edited by all co-authors.

Table of Contents

Abstract	ii
Preface	iii
Table of Contents	vi
List of Tables	ix
List of Figures	x
List of Abbreviations	xii
Acknowledgements	xiv
Dedication	xv
1 Introduction	1
1.1 Motivation	1
1.2 Problem Statement	3
1.2.1 Objectives	3
1.2.2 Research Questions Addressed	4
1.3 Brain Imaging for Macroscale Connectomics	5
1.3.1 Diffusion Magnetic Resonance Imaging	5
1.3.2 Functional Magnetic Resonance Imaging	6
1.4 Anatomical and Functional Brain Connectivity	7
1.4.1 Estimating Anatomical Connectivity	7
1.4.2 Estimating Functional Connectivity	12
1.4.3 Relationship between Anatomical and Functional Brain Connectivity	13
1.5 Functional Segregation in the Brain: Activation Detection	15
1.5.1 Unimodal Approaches	15
1.5.2 Multimodal Approaches	16

Table of Contents

1.6	Functional Integration in the Brain: Subnetwork Identification	17
1.6.1	Unimodal Approaches	17
1.6.2	Multimodal Approaches	19
1.7	Thesis Contributions	20
1.7.1	Anatomical and Functional Brain Connectivity	20
1.7.2	Multimodal Brain Activation Detection	21
1.7.3	Multimodal Brain Subnetwork Identification	21
2	Anatomical and Functional Brain Connectivity	22
2.1	Super-resolution for dMRI	24
2.1.1	Acquisition Model	25
2.1.2	Dictionary Construction	25
2.1.3	Super-Resolved Volume Generation	26
2.1.4	Experiments	27
2.1.5	Discussion	30
2.2	Estimating Anatomical Connectivity	33
2.2.1	Modelling of Local Diffusion Properties	34
2.2.2	Tractography	34
2.2.3	Experiments	35
2.2.4	Discussion	37
2.3	Quantifying Anatomical Connectivity	38
2.3.1	Anatomical Connectivity Metrics	39
2.3.2	Experiments	40
2.3.3	Discussion	42
2.4	Summary	43
3	Multimodal Brain Activation Detection	44
3.1	Modelling Intrinsic Functional Connectivity	45
3.1.1	Group Activation Detection	46
3.1.2	RS Subnetwork Detection	47
3.1.3	RS Activity Estimation and Removal	48
3.1.4	Experiments	50
3.1.5	Discussion	52
3.2	Modelling Anatomical Connectivity	53
3.2.1	Random Walker for Activation Estimation	54
3.2.2	Anatomical Connectivity Prior Estimation	55
3.2.3	Activation Likelihood Estimation	56
3.2.4	Group Activation Inference	57

Table of Contents

3.2.5	Re-estimation of Group Activation Using Partial Con-	
	nectome	58
3.2.6	Experiments	58
3.2.7	Discussion	63
3.3	Summary	63
4	Multimodal Brain Subnetwork Identification	64
4.1	Replicator Dynamics	65
4.2	Coupled Replicator Dynamics	66
4.3	Coupled Stable Replicator Dynamics	68
4.3.1	Graph Incrementation	68
4.3.2	Stability Selection	68
4.4	Coupled Overlapping Stable Replicator Dynamics	69
4.5	Parameter Selection and Implementation	70
4.6	Materials	71
4.6.1	Synthetic Data	71
4.6.2	Real Data	72
4.7	Results	72
4.7.1	Synthetic Data Results	72
4.7.2	Real Data Results	73
4.8	Discussion	77
5	Conclusions	78
5.1	Future Work	79
	Bibliography	81

List of Tables

2.1 Comparison of fiber length bias for AC metrics.	42
-------------------------------------------------------------	----

List of Figures

1.1	Projected average number of years in full health lost due to disability and premature death of the 2010 to 2020 birth cohort in Canada.	2
1.2	Projected indirect economic costs of selected neurological conditions due to working-age death and disability.	3
1.3	Sample workflow of a dMRI experiment.	10
1.4	Depiction of the false negative problem in AC estimation.	11
1.5	Sample workflow of an fMRI experiment.	14
2.1	Qualitative comparison between the tract-density maps and fiber tracts reconstructed from the original and super-resolved dMRI data.	29
2.2	AC-FC correlation for 10 subjects with AC estimated from the data at its original resolution (2 mm isotropic), and high-resolution data (1 mm isotropic) obtained using trilinear interpolation, spline interpolation, CLASR and the proposed method.	31
2.3	Number of inter-parcel tracts reconstructed from the data at its original resolution (2 mm isotropic), and high-resolution data (1 mm isotropic) obtained using trilinear interpolation, spline interpolation, CLASR and the proposed method.	32
2.4	Fiber tracts reconstructed for a representative subject using streamline tractography on diffusion tensors, streamline tractography on ODFs, and global tractography on ODFs with an ROI placed in the corpus callosum.	37
2.5	Correlation between AC and FC using four commonly used AC metrics. FC was estimated from RS fMRI data, thus depicts intrinsic connections.	41
2.6	Correlation between AC and FC using four commonly used AC metrics. FC was estimated from task fMRI data acquired during 7 different tasks.	41

List of Figures

3.1	Graphical depiction of proposed RS activity removal approach.	47
3.2	Dominant principal component extracted from time courses of intrinsically-connected brain areas.	49
3.3	Comparison of activation detection results with and without the proposed RS activity removal method.	51
3.4	Detected activation patterns with and without the proposed RS activity removal method.	52
3.5	Overview of the proposed multimodal task activation detection approach.	54
3.6	Quantitative results of the proposed multimodal activation detection technique.	61
3.7	Qualitative results of the proposed multimodal activation detection method.	62
4.1	Subnetwork identification accuracy on synthetic data.	73
4.2	Subnetwork identification results on real data.	75
4.3	Visuospatial subnetwork identified by CSORD at the group level.	75
4.4	Right executive control subnetwork identified by CSORD at the group level.	76
4.5	A group subnetwork identified by CSORD that comprises the visual corticostriatal, CTC, and STC loops.	76

List of Abbreviations

AC	Anatomical connectivity
BOLD	Blood oxygen level dependent
CIS	Connected Iterative Scan
CRD	Coupled Replicator Dynamics
CSORD	Coupled Stable Overlapping Replicator Dynamics
CSRD	Coupled Stable Replicator Dynamics
CSA	Constant solid angle
CTC	Cerebello-thalamo-cortical
DALY	Disability-adjusted life years
dMRI	Diffusion magnetic resonance imaging
DSI	Diffusion spectrum imaging
DTI	Diffusion tensor imaging
DW	Diffusion weighted
EM	Expectation-maximization
FA	Fractional anisotropy
FC	Functional connectivity
FDR	False discovery rate
fMRI	Functional magnetic resonance imaging
FP	False positive
GBD	Global Burden of Diseases
GΓG	Gamma-Gaussian-Gamma
GLM	General linear model
GMM	Gaussian mixture model
HARDI	High angular resolution diffusion imaging
HCP	Human Connectome Project
HRF	Hemodynamic response function
ICA	Independent component analysis
ICC	Intraclass correlation
LSC	Left somatomotor cortex
MNI	Montreal Neurological Institute
MRF	Markov random field

List of Abbreviations

NC	Normalized cut
ODF	Orientation distribution function
OSLOM	Order Statistics Local Optimization Method
PCA	Principal component analysis
PCC	Posterior cingulate cortex
PDD	Principal diffusion direction
PDF	Probability density function
RD	Replicator Dynamics
ROI	Region of interest
RW	Random walker
QBI	Q-ball imaging
RS	Resting state
RSC	Right somatomotor cortex
SNR	Signal-to-noise ratio
STC	Striato-thalamo-cortical
SVD	Singular value decomposition
TDI	Tract-density imaging
TP	True positive
WHO	World Health Organization

Acknowledgements

I would like to thank my advisor, Dr. Rafeef Abugharbieh, for supervising me throughout my doctoral studies. I owe particular thanks to Dr. Bernard Ng for his guidance and dedication to research. Furthermore, I would like to thank my friends inside and outside of BiSICL for their support.

Last but not least, I would like to thank my better half, Başak Öztaş Yoldemir, for always being there for me.

Dedication

To Başak, my amazing wife

Chapter 1

Introduction

The human brain is the control center of the nervous system and regulates almost every aspect of human behavior. It is the most complex organ in the human body with hundreds of billions of neurons, which communicate with each other via hundreds of trillions of connections. As such, the human brain has been described as a three-pound universe [1]. Anatomically, it is an elaborate network of processing centers connected through neural axons. Functionally, the brain comprises specialized regions that work together to enable complex cognitive processes. It is thus important to not only characterize anatomy and function in the brain, but also interrelations, or connectivity.

1.1 Motivation

According to the Mapping Connections report published by the Government of Canada and Neurological Health Charities Canada in September 2014 [2], Canadians born in the current decade (2010 to 2020) who will suffer from neurological disorders are projected to lose tens of years of life living in full health. The average number of years in full health lost due to disability or premature death for five neurological conditions is given in Fig. 1.1. In addition to their physical, emotional and cognitive impacts, such disorders have immense economic costs. The Mapping Connections study suggests that the indirect economic cost of neurological disorders to the Canadian society is billions of dollars every year due to working-age death and disability. Moreover, these costs are expected to steadily increase with the life expectancy of the population increasing (Fig. 1.2). Worldwide impacts of neurological conditions are becoming more widely recognised. According to the latest Global Burden of Diseases (GBD) study published in December 2012 [3], mental and neurological diseases account for a major portion of the global disease burden. To quantify the number of years lost as a result of premature death and disability, the GBD study defines a metric dubbed disability-adjusted life years (DALYs), where one DALY equals one lost year of healthy life. In 2010, mental and neurological diseases resulted

1.1. Motivation

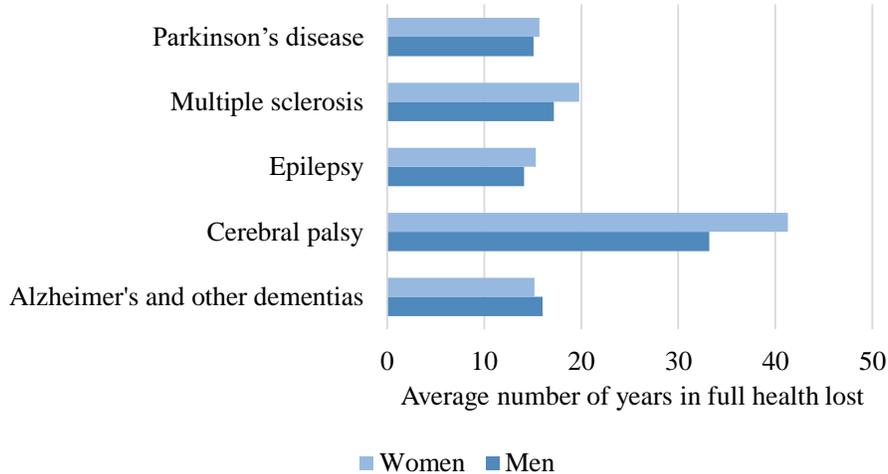


Figure 1.1: Projected average number of years in full health lost due to disability and premature death of the 2010 to 2020 birth cohort in Canada [2]. Difference between the life expectancy and the health-adjusted life expectancy for each of the selected neurological conditions is reported.

in 201 million DALYs with an increase of 40% since 1990 [3].

During the recent years, neuroimaging has greatly contributed to the characterization and understanding of mental and neurological diseases. Several imaging modalities have been explored, each delineating the brain from a different view. Traditionally, these modalities are individually considered when making inferences about the brain. The main premise of this thesis is that fusing data from different imaging modalities should prove beneficial by leveraging their respective strengths. Specifically, our focus is on combining information regarding brain activity, and the structural substrate that enables brain functions to arise.

Among available imaging modalities, we chose to employ magnetic resonance imaging (MRI) since it is relatively widespread and safe. Global popularity of MRI makes it advantageous in terms of uptake, i.e. the likelihood of the developed techniques being adopted is high. To investigate brain function, we use functional MRI (fMRI) in this thesis. fMRI has shown promise in a number of clinical applications including localization of specific brain functions to guide neurosurgery [4], characterization of mechanisms of disease [5], prediction of treatment response [6], and characterization of disease

1.2. Problem Statement

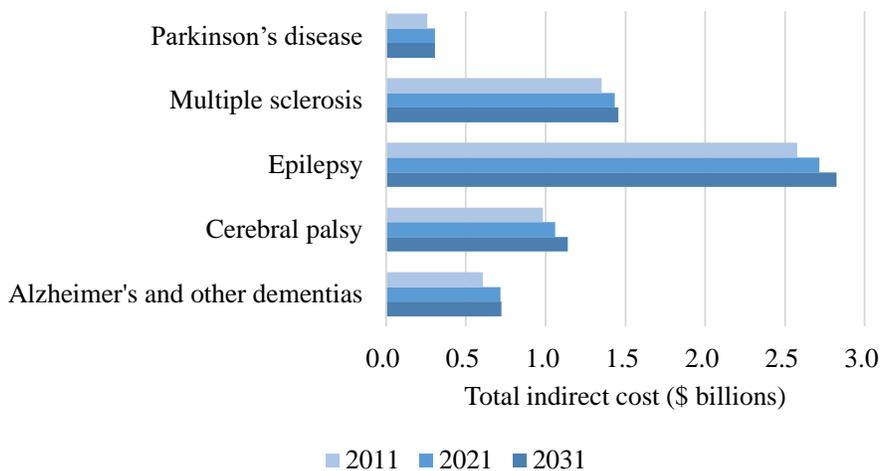


Figure 1.2: Projected indirect economic costs of selected neurological conditions due to working-age (age 15 to 74 years) death and disability [2]. Data are expressed in Canadian dollars.

risk [7]. Although fMRI provides us with an effective tool to probe brain's functional organization, integrating this information with anatomical connectivity, which essentially facilitates the functional interactions [8], might provide richer insights about inner dynamics of the brain. Hence in this thesis, we propose complementing analyses of brain function with anatomical connectivity information learned through diffusion MRI (dMRI), which is currently the only modality that allows us to infer fiber pathways in the brain *in vivo* and noninvasively [9].

1.2 Problem Statement

1.2.1 Objectives

The overarching objective of this thesis is to increase the sensitivity in the assessment of functional brain organization. Accurate mapping of functional brain organization is immensely important since it may ultimately enable identifying reliable biomarkers for various neurological diseases. However, there are several challenges to this problem that are commonly encountered in practice, such as poor data quality, low spatial and temporal resolution, and low number of subjects. We focus on multimodal frameworks to

mitigate such challenges. Our rationale is that combining information from different neuroimaging modalities should prove beneficial by leveraging their respective strengths and increasing the amount of information available.

1.2.2 Research Questions Addressed

Anatomical and functional connectivity (AC and FC) analyses provide two complementary views of brain circuitry. Motivated by how anatomical wiring of the brain facilitates and shapes functional interactions between brain regions, we propose multimodal fusion methods to integrate information regarding the two types of brain connectivity. To facilitate meaningful data fusion, we first address the important problem of brain connectivity estimation and quantification. Specifically, we focus on various steps of AC estimation from dMRI data and present ways of increasing the consistency between estimates of AC and FC. Next, we propose novel multimodal integration methods to investigate the brain’s organization. It is known that functional organization in the brain is governed by a subtle balance between functional segregation and integration [10]. The segregation principle states that some functional processes specifically engage localized and specialized brain regions, whereas the integration principle emphasizes the importance of connections across distributed brain regions in the emergence of brain function [11]. These two views of functional organization are generally addressed by investigating:

- which brain regions are activated upon stimulus, i.e. activation detection,
- which brain regions work together at rest or in response to stimulus, i.e. subnetwork (also known as community or module) identification.

We propose novel methods combining fMRI and dMRI data to address these two key problems.

In the next sections, we first provide an overview of dMRI and fMRI and describe how they enable *in vivo* estimation of AC and FC. We then provide an outline of the two key research questions regarding functional brain organization, describe unimodal and multimodal techniques employed to tackle these two problems, and conclude with a summary of the thesis contributions.

1.3 Brain Imaging for Macroscale Connectomics

Connectomics is the study of the human “connectome”, a term which was introduced in 2005 to refer to the complete map of neural connections in the brain [12]. Recently, the term connectome is also used in a broader sense to refer to brain regions, their anatomical connections, and their functional interactions [13]. The connectome can be analyzed at multiple scales from the microscale, which focuses on individual neurons and their synaptic connections, to the macroscale, which focuses on large brain regions with cognitive and behavioral associations. Currently, macroscale connectomics is by far the most common since noninvasive tools for *in vivo* imaging of the human connectome are only available at the macroscale. Among the modalities used for macroscale connectomics, dMRI and fMRI are especially popular [13].

1.3.1 Diffusion Magnetic Resonance Imaging

dMRI is currently the only imaging modality that allows the assessment of white matter connectivity by examining the translational displacement of water molecules [14]. Water molecules are in constant thermal motion constrained by physical boundaries. In fibrous tissues such as axons in the white matter, water diffuses more rapidly in the direction aligned with the fibrous structure and more slowly perpendicular to it. Measurements of this anisotropic diffusion hence reveal micro-structural properties of the underlying tissue.

In practice, several so-called diffusion weighted (DW) images which are sensitized to different diffusion directions are acquired, where each voxel intensity reflects the estimate of the rate of local water diffusion for the corresponding diffusion direction. These estimates are then combined in a reconstruction step to estimate the 3D diffusion probability density function (PDF) at each voxel [15]. Depending on how reconstruction is done, the imaging technique takes on different names such as diffusion tensor imaging (DTI), diffusion spectrum imaging (DSI) and Q-ball imaging (QBI) [16]. Among these, DTI is the simplest and the most common imaging technique, where the local diffusion at each voxel is characterized by a 3×3 symmetric, positive semi-definite, second-order tensor called the diffusion tensor. The inherent assumption here is that the underlying diffusion process is Gaussian, which is the main limiting factor of DTI. Common to all mentioned PDF reconstruction techniques, tractography is one of the most powerful tools dMRI offers. In its simplest sense, tractography is a way of delineating

the fiber tracts in the brain by integrating pathways of maximum diffusion coherence [16]. Hence, tractography allows us to infer AC patterns of the brain *in vivo* and noninvasively.

The most common artifacts in DW images are subject's head motion and eddy currents arising from rapid switching of gradient pulses during data acquisition. Eddy current artifacts are manifested in three different forms: contraction or dilation of the image, overall shift and shear [17]. These artifacts can be corrected with an affine registration to the $b=0$ image, i.e. the volume with no diffusion weighting. Head motion can be corrected with a rigid body registration to the $b=0$ image. These two artifacts are typically accounted for using a single registration step [18].

1.3.2 Functional Magnetic Resonance Imaging

fMRI is an *in vivo* and noninvasive imaging modality for studying brain activity, most popularly through blood oxygen level dependent (BOLD) contrast [19]. Due to hemodynamic response, i.e. rapid increase in blood supply to active brain regions, changes in oxygen concentration are indirect implications of brain activity. Exploiting this mechanism, BOLD contrast imaging is based on MR images made sensitive to changes in the oxygenation state of hemoglobin molecule [19]. Hemoglobin behaves as a diamagnetic substance when it is saturated with oxygen (oxyhemoglobin) and as a paramagnetic substance when some oxygen molecules have been removed (deoxyhemoglobin). These changes in magnetic properties are manifested in the fMRI observations, areas with high concentration of oxyhemoglobin giving a higher signal than areas with low concentration [20]. Broadly, there are two different types of fMRI experiments. In the first one, the subject inside the scanner performs a series of cognitive tasks while BOLD images are being collected. This experiment is typically followed by a statistical analysis of the data in which signal intensity changes are compared to the task paradigm [21]. In the second type, the subject is asked not to be engaged in any cognitive task without falling asleep, where spontaneous fluctuations in the brain are recorded [22]. These two approaches are referred to as task fMRI and resting state fMRI (RS fMRI), and they are regarded as two different modalities in this thesis.

The inherently low signal-to-noise ratio (SNR) of BOLD signal (typically between 0.2 and 0.5 [23, 24]) together with confounds such as scanner drifts and subject's head motion pose major challenges to the interpretation of fMRI data. As such, fMRI data typically undergo several steps of preprocessing before any analysis of the data. Standard preprocessing steps

include slice timing correction [25], head motion correction [26], and temporal autocorrelation correction [27, 28]. Slice timing correction accounts for the fact that BOLD signals at different slices of the brain are recorded at slightly different time points. To correct for this, individual slices are temporally realigned to a reference slice based on their relative timing. Similar to dMRI, head motion in fMRI acquisitions is corrected by a rigid body registration to a reference volume (typically the first volume). Temporal autocorrelations arise from scanner drifts, cardiac noise, respiratory noise, and head motions [28]. Since the residual of the voxel time courses is often assumed to be identical independently distributed (i.i.d.) in typical analysis methods, pre-whitening is performed to satisfy this assumption. In addition to these steps, spatial normalization to a common space, such as the Montreal Neurological Institute (MNI) space [29], can be performed to ensure spatial correspondence between different subjects [26]. There exist several toolboxes such as SPM [30], FSL [31] and DPARSF [32] for streamlined preprocessing of fMRI data.

1.4 Anatomical and Functional Brain Connectivity

1.4.1 Estimating Anatomical Connectivity

Diffusion Modelling and Tractography

Tractography using dMRI data relies on the fundamental assumption that water diffusion reflects the axonal orientations. The first attempt to reconstruct white matter bundles throughout the brain involved the use of diffusion tensors, where a 3×3 symmetric, positive semi-definite tensor field is used to model 3D Gaussian distribution of diffusing water molecules [14, 33]. It was shown that the direction of least hindrance to diffusion, i.e. the principal diffusion direction (PDD), is correlated with the major eigenvector of the diffusion tensor at each voxel [34]. This led to several “streamline” tractography approaches where the main idea is to integrate the PDD at each voxel, using it as a surrogate for axons’ orientation [35–37]. Streamline tractography performed on diffusion tensors currently remains to be the norm in estimating AC, mainly due to the intuitive nature of this approach: Diffusion tensors can be visualized as ellipsoids oriented along their PDDs, and streamline tractography boils down to connecting the ellipsoids “pointing” at each other. However, it is now clear that PDD may not always be a good representative of the underlying fiber bundles’ orientation, especially

at regions with complex fiber geometries, e.g. where fiber bundles meet or cross. For instance, when there are two crossing fiber bundles with similar radii at a given location, the diffusion tensor has a planar shape. In such a case, we cannot draw conclusions about the underlying microstructure since different bundle configurations may lead to the same planar shape.

It is possible to characterize the diffusion process without imposing a diffusion model, such as the Gaussian model assumed in the case of diffusion tensors. It is apparent that this approach is theoretically less prone to biases imposed by an explicit diffusion model [38]. The relationship between the 3D PDF of diffusion, also known as the spin propagator, and measured dMRI data is given by the theory of q -space [39]. Given the spin propagator, the marginal probability of diffusion in any given direction, i.e. orientation distribution function (ODF), provides a reduced representation of the PDF and is of interest for mapping white matter architecture. ODF can be found using a linear radial projection of the PDF [40] or by doing the projection in a cone of constant solid angle (CSA) [41]. It has been shown that using the projection in a cone of CSA, denoted as CSA-ODF, results in a sharper representation of the diffusion process leading to clearer recovery of certain fiber bundles [41].

Using ODFs, which are more sensitive in capturing intricate intravoxel fiber geometries than diffusion tensors, may not directly translate into higher sensitivity in AC estimation since the tractography method used also needs to make use of this extra information. For example, if only the dominant peak of the ODF is used in a greedy manner in reconstructing white matter bundles, there may be no performance improvement at regions where a large fiber bundle crosses a smaller bundle. This is due to the partial volume averaging effect of diffusion tensors, i.e. the diffusion tensor would point along the more dominant bundle (the direction of the dominant ODF peak) at such a location. On the other hand, additionally using the secondary ODF peak may enable identifying the fiber crossing. However, this approach is also prone to errors since it necessitates estimating the number of bundles in a voxel to determine the number of ODF peaks to consider. A notable method that bypasses this potential source of error is global tractography [42], where the decisions are made on a global level as opposed to the greedy nature of streamline tractography. Global tractography reconstructs all fiber tracts simultaneously, which enables fiber trajectories to be jointly considered in determining the most possible fiber configuration. In this approach, short fiber segments are bridged together to form the set of fiber tracts that best explains the measured dMRI data. This global approach is particularly beneficial in resolving cases where local diffusion characteristics are ambiguous

due to crossing fibers. Streamline tractography tends to terminate tracking at such locations, whereas global tractography exploits the geometry of the surrounding fibers by considering the whole-brain fiber configuration in aggregate to compensate for the lack of reliable local information.

Both streamline and global tractography provide us with point estimates, i.e. the variability of the reconstructed tracts based on underlying dMRI data is not quantified. Several so-called probabilistic tractography methods have been proposed to ameliorate this limitation. Probabilistic tractography techniques are typically based on either modelling noise parameters by some probability distribution [43, 44], or bootstrap-based methods that capture the uncertainty in the data [45]. However, these methods are generally very computationally expensive, limiting their use when whole-brain connectivity is to be estimated. On a regular workstation, probabilistic tractography of one brain at a typical resolution may take almost one day [46]. The intensive computation involved thus inhibits the adoption of probabilistic tractography in clinical tasks. Steps in a typical dMRI experiment are shown in Fig. 1.3.

A major problem with AC estimation using tractography is the false negative connections due to premature termination of reconstructed tracts at fiber crossings. Given the recent evidence that 60% to 90% of the white matter voxels contain crossing fibers [47], this limitation has serious implications in terms of the quality of AC estimates. An example tractography result depicting this problem is given in Fig. 1.4, where the reconstructed tracts passing through corpus callosum are shown in red, and those passing through internal capsule are shown in green. Dotted red and green arrows indicate the expected path of the tracts which could not be reconstructed. Such false negatives significantly hinder reliable AC estimation, especially when estimating interhemispheric connections.

Quantifying Anatomical Connectivity

Performing tractography on dMRI data results in a set of reconstructed fiber tracts. Unfortunately, this reconstruction is inherently non-quantitative in terms of the connectivity strength. Although probabilistic tractography appears to be more quantitative, it is important to realize that it provides estimates of our confidence on the fiber trajectories, and not connection strength. As an intuitive example, consider a case where brain regions A and B are connected through a relatively small fiber bundle that has no immediate neighboring bundles. Also consider regions C and D that are connected through a major fiber bundle that crosses another fiber bundle

1.4. Anatomical and Functional Brain Connectivity

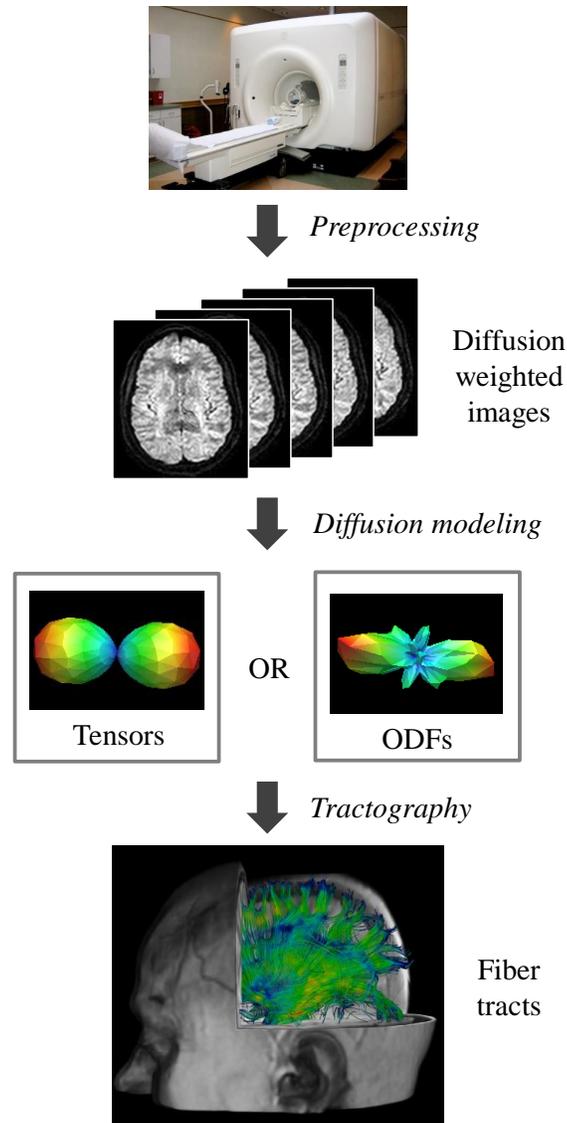


Figure 1.3: Sample workflow of a dMRI experiment. During a dMRI experiment, DW images sensitized to different directions are acquired, where intensity of each voxel reflects the rate of diffusion at that location. Information from multiple directions are combined to reconstruct the diffusion tensors or ODFs. Fiber tracts are then delineated by following tensors or ODFs with coherent orientations. MR scanner image: Big MRI by Liz West, CC BY 2.0, <https://www.flickr.com/photos/calliope/223220955>.

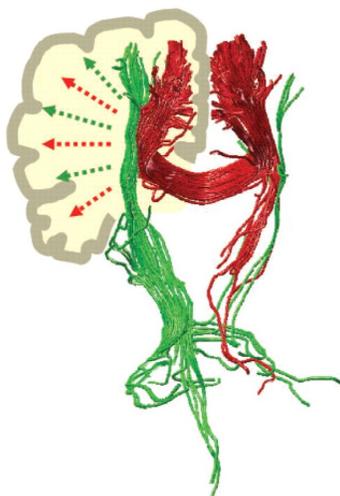


Figure 1.4: Depiction of the false negative problem in AC estimation using streamline tractography on diffusion tensors. Tracts extending from corpus callosum (red) and internal capsule (green) are shown. Dotted arrows indicate the expected path of the tracts extending from corpus callosum and internal capsule, which tractography failed to reconstruct. This figure is reproduced from [48] with the permission of the publisher.

between C and D. In this hypothetical case, probabilistic tractography might assign higher confidence to the connection between A and B, even though C and D are connected through a major fiber bundle that provides a higher capacity for interaction. We need quantitative measures of this capacity (which is loosely termed as AC strength here) for many problems such as devising biomarkers of tissue structure and analyzing the relationship between anatomical and functional connectivity in the brain.

Quantifying AC in the brain is most commonly based on quantifying one or more aspects of the fiber tracts reconstructed using streamline or global tractography, though computationally expensive approaches based on probabilistic tractography techniques were also explored [49]. The choice of which property to measure and of how to map it into an AC metric are key aspects affecting the AC estimates. Arguably the most common AC metric is the number of reconstructed fiber tracts between pairs of brain regions, commonly referred to as fiber count. A variant of this approach involves a normalization of the fiber count by the total volume of the region pairs they

connect to account for the variable size of the brain regions [50]. Besides metrics based on fiber count, use of the total length of reconstructed fiber tracts has also been suggested as a measure of AC [51] aiming to correct for the fact that longer tracts may have larger accumulated error, leading to lower fiber counts. Another metric, the average fractional anisotropy (FA) along streamlines connecting regions, has also been proposed as a proxy for anatomical connectivity strength [52]. FA is a local measure that describes how anisotropic the underlying diffusion process is, and takes on values between 0 (perfectly isotropic diffusion) and 1 (an infinite cylinder).

It is important to acknowledge that all of the aforementioned AC metrics are confounded by several factors, limiting their interpretability. First, tractography can only delineate bundles of fibers in the brain, and not individual fibers. The term *fiber count* can thus be misleading. Indeed, using the term *streamline count* has been recently proposed as an alternative [53]. Nonetheless, we use the term fiber count for easier interpretation and to conform to the jargon used in existing literature, with the understanding that it is the streamlines that are actually being counted. Moreover, we note that the number of fibers is dependent on the number of seeds used for tracking the fibers, the tractography method used, and several features of the pathway such as curvature, length and width [53]. Additionally, we highlight that FA not only depends on the reliability of local diffusion information, but also on a large number of modulating factors such as axonal ordering, axonal density, amount of myelination, and increase in extracellular or intracellular water [53]. However, such confounding factors did not impede the adoption of a variety of AC metrics, driven by a practical need for quantifying the degree of connection between brain regions. To this end, reconciling the presence of confounding factors with the practical need for quantitative connectivity estimation calls for a detailed analysis to determine which AC metric has the highest potential of being of practical use in multimodal brain image analysis efforts.

1.4.2 Estimating Functional Connectivity

FC is conventionally defined as the temporal dependence of neuronal activity patterns of remote brain regions [54], and it describes how brain region pairs co-activate (positive synchrony, negative synchrony or no synchrony). Such co-activation patterns exist both during rest and task performance [22]. FC studies have been gaining traction in the clinical realm with applications such as identifying group differences between populations, devising biomarkers for obtaining diagnostic and prognostic information, and guid-

ance of invasive and noninvasive treatments [55].

Several metrics have been proposed to quantify the dependence between time courses of the BOLD signal. Popular metrics include Pearson’s correlation [54], partial correlation [56], coherence [57] and partial coherence [57], with Pearson’s correlation being the most widely used metric by far. Pearson’s correlation measures marginal (direct and indirect) dependencies, whereas partial correlation quantifies conditional (only direct) dependencies. Coherence and partial coherence are the spectral counterparts of correlation metrics. Though not commonly used, there exist more complex metrics such as those based on higher-order statistics and temporal lag [58]. In simulation based studies, it has been shown that correlation metrics perform significantly better than more complex metrics [58]. On real data, test-retest analyses confirmed that correlation is a more stable metric than coherence [59]. We have thus used Pearson’s correlation in this thesis to quantify FC.

It is important to acknowledge that correlation between voxel time courses is not informative of causality, or whether the connection is direct. For example, if brain region A is feeding into regions B and C , FC analysis based on correlation would conclude that B and C are connected. Although this behaviour is consistent with the conceptual definition of FC given above (neural synchrony, or co-activation), it is important to understand the precise nature of connectivity measured by the metric of choice. Steps in a typical fMRI experiment are shown in Fig. 1.5.

A major problem with FC estimation is the many false positive connections arising from confounds such as head movements, cardiac and respiratory pulsations, and scanner drifts. Among these, head motion artifacts are likely the most severe problem in fMRI studies [60]. Head motion results in artificial brain activations if it is not properly corrected for, since data analyses assume that each voxel represents a unique location in the brain. Such effects are typically visible as “ring activations” around the edges of the brain. Also, the voxels at brain boundaries may appear to be functionally connected due to the factitious peak in brain activity they commonly exhibit.

1.4.3 Relationship between Anatomical and Functional Brain Connectivity

Motivated by how anatomical wiring of the brain facilitates and shapes functional interactions between brain regions, a growing interest in understanding the relationship between AC and FC has emerged in recent years. To this end, AC-FC comparison has been done on a single axial slice [61],

1.4. Anatomical and Functional Brain Connectivity

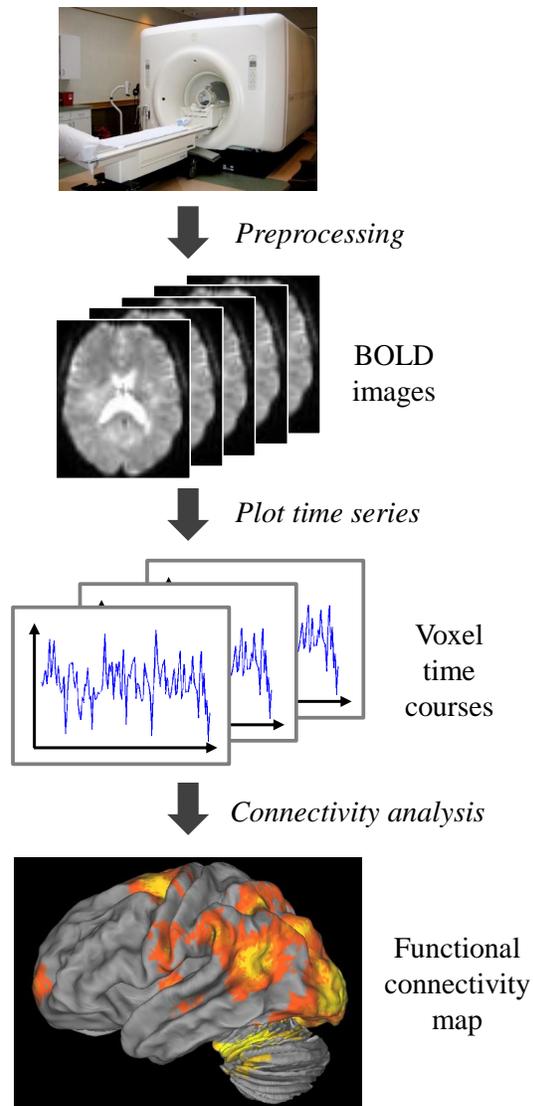


Figure 1.5: Sample workflow of an fMRI experiment. During an fMRI experiment, the subject is either not engaged in any cognitive task (resting state fMRI), or performs a series of cognitive tasks (task fMRI) while BOLD images are being acquired. The preprocessed voxel time courses are compared with each other to infer the functional connectivity structure. MR scanner image: Big MRI by Liz West, CC BY 2.0, <https://www.flickr.com/photos/calliope/223220955>.

within a specific subnetwork [8, 62], and for the whole brain [50, 63, 64]. Collectively, these studies suggest that there is a robust and strong correlation between estimates of AC and FC. Although these earlier works comparing AC and FC suggest high consistency between their estimates, it has been more recently shown that the AC-FC correlation is rather low with typical data acquisitions [65]. A fundamental factor governing the level of consistency is the way in which FC and AC are estimated. Although neuroimaging community largely settled on the use of Pearson’s correlation between voxel time courses for estimating FC [58, 59], investigations of AC estimation and quantification have received less attention. False negatives are common in tractography mainly due to the abundance of crossing fibers in the human brain. There are two potential solutions to ameliorate this problem. The first one involves increasing the spatial resolution of dMRI data such that each voxel covers a smaller region in the brain, reducing the probability that a given voxel will comprise multiple fiber bundles. The second solution is to employ a tractography algorithm that handles fiber crossings better than the widely used streamline tractography. Furthermore, it is important to determine which AC metric has the highest potential of being of practical use in multimodal brain image analysis efforts.

1.5 Functional Segregation in the Brain: Activation Detection

1.5.1 Unimodal Approaches

The aim of activation detection is to map brain regions to function by investigating the regions that are activated in response to stimuli. The standard approach for inferring brain activation from fMRI data is to model the fMRI observations at each voxel independently as a linear combination of expected temporal responses using the general linear model (GLM) [66] to estimate the voxel-wise activation effects. Statistics, such as t-values, corresponding to the p-values of observing the estimated activation effects when there is no activation are combined in an activation statistics map. This map is then thresholded, resulting in a set of supra-threshold voxels that are declared as active. However, since inference is done for all voxels separately, there is a need for adjusting the threshold for multiple comparisons [67]. A major shortcoming of the GLM approach is that, spatial correlations among voxels are ignored by independently testing each voxel. To account for the tendency of the nearby voxels to work in tandem, activation statistics maps

are typically smoothed with a Gaussian kernel as a post-processing step [68]. Bayesian schemes incorporating spatial priors to smooth the data in a data-driven way have also been proposed [69, 70]. These adaptive approaches have been shown to preserve the spatial features of the activation patterns that would be lost with a fixed Gaussian kernel [70]. Evidently, both adaptive and nonadaptive filtering approaches result in reduced spatial resolution, decreasing the localization power of the analyses. Alternatively, Markov random fields (MRF) based approaches which view activation detection as a class labelling problem, i.e. active and nonactive, have also been explored [71–73]. These methods encourage spatial continuity by penalizing the discrepancies of labels among neighbouring voxels, as opposed to refining activation statistics maps. Although such methods help suppress false spatially-isolated activations, they completely ignore long-range functional interactions. To address this limitation, incorporation of functional connectivity information into activation detection has been put forth [74], but the approach taken estimates both activation effects and functional connectivity from the same dataset, hence the information gain might be limited. Given the typically strong noise in fMRI data, exploring other sources of information to regularize activation detection may be beneficial.

1.5.2 Multimodal Approaches

A major challenge in brain activation detection is the existence of spontaneous fluctuations in the brain in addition to task-evoked signal [75]. Such spontaneous modulations in brain activity are continually present at all times [76]. Thus, part of the noise observed in task fMRI data ascribes to ongoing RS activity. As opposed to the traditional view that RS activity is just noise that needs to be averaged out, strong synchrony in RS modulations between specific brain areas has been observed in numerous studies [77]. In fact, many of the detected RS subnetworks exhibit high resemblance to subnetworks seen in task experiments [78]. Further supporting this finding is a recent work [79] that demonstrated enhanced sensitivity in task activation detection by incorporating an RS connectivity prior. Taken together, these findings suggest that there is spatial structure in RS activity which is likely similar at rest and during task performance [75].

Since the time at which RS activity peaks and troughs is difficult to predict, decoupling task-evoked activity from RS fluctuations is non-trivial. To the best of our knowledge, the only previous work that attempted to tackle this challenging problem of RS activity removal from task fMRI data was by Fox *et al.* [75]. Specifically, the authors showed that for a right

handed motor task, subtracting out fMRI signals in the right somatomotor cortex (RSC) from the left somatomotor cortex (LSC) after weighting by a regression coefficient derived from the RS data significantly reduced inter-trial variability in fMRI response, suggesting that RS fluctuations and task-evoked signal are superimposed roughly in a linear fashion in the human brain. The basis of this approach is twofold. First, the presence of coherent RS activity between LSC and RSC is well established [22]. Second, right handed motor tasks typically activate only LSC, thus signals in RSC would largely correspond to RS activity. Subtraction of signals in RSC from LSC would thereby remove the RS components within the task fMRI time courses of LSC. However, not all tasks evoke only lateralized activation. Thus, simply subtracting signals in one side of the brain from the other side is not always suitable for removing RS modulations from task fMRI data.

1.6 Functional Integration in the Brain: Subnetwork Identification

1.6.1 Unimodal Approaches

Since the interactions between specialized regions in the brain are known to play a key role in mediating brain function [80], fMRI has been widely used for the modelling of voxel interactions. One of the main approaches for studying this integrative property of the brain is to group voxels into functional subnetworks based on the temporal dependencies among their fMRI observations. The most widely used techniques for this include the seed-based approach and independent component analysis (ICA) [81]. In the seed-based approach, functional subnetworks are estimated based on the temporal correlations between a seed region and all other regions in the brain [81]. Since pre-specifying a seed might be difficult in some cases, techniques such as k-means [82], hierarchical clustering [83], and Gaussian mixture model (GMM) [84] have been proposed to automatically identify seed regions. Despite its simplicity and intuitiveness, the seed-based approach has the drawback that results are restricted by the choice of seeds [85]. In contrast, ICA provides a data-driven means for extracting subnetworks by decomposing fMRI observations into maximally independent spatial components and associated time courses [86]. At the intra-subject level, voxels assigned weights larger than a threshold within the same spatial component are grouped as a subnetwork. However, the threshold is generally chosen in an ad hoc manner, which precludes statistical interpretation [86]. In prac-

tice, ICA is typically performed at the group level by pooling information from multiple subjects, in which a statistically meaningful threshold can be set [87].

An important point commonly overlooked in the literature is the fact that certain brain regions are known to interact with multiple subnetworks [88]. It is thus important to develop methods that allow spatial overlap among identified subnetworks. Recently, a number of techniques have been put forth for identifying overlapping brain subnetworks. In [89], the technique involves first performing maximal clique detection and treating each maximal clique as a node of a new graph. Subnetworks are then extracted by applying modularity optimization [90] on this new graph. Since the maximal cliques might share nodes, overlaps between subnetworks are enabled. The main shortcoming of this technique is that it requires graph edges to be binarized since it cannot handle weighted graphs, but brain connectivity lies on a continuum. This need for edge binarization could thus result in important information being discarded. A technique adopted from social network analysis called connected iterative scan (CIS) has also been explored for brain subnetwork identification [91]. Taking each node as a candidate subnetwork, CIS determines if other nodes belong to this candidate subnetwork based on their influence on a graph density metric. A limitation of CIS is that it is sensitive to a weighting factor that controls subnetwork size. Another technique for finding overlapping brain subnetworks is temporal ICA (tICA), which is more suited for finding overlaps than the more widely-used spatial ICA (sICA), since maximizing spatial independence tends to result in little overlaps [88]. However, the proposed implementation uses an ad hoc threshold for extracting interacting nodes from the spatial maps, which limits the statistical interpretability of the results. Grouping the edges between nodes instead of the nodes themselves has also been proposed for detecting subnetwork overlaps [92]. In that technique, edges are grouped using e.g. hierarchical clustering with the set of nodes associated with any edge of a given edge cluster declared as a subnetwork [93]. The main limitation of this technique is scalability since in weighted graphs where all nodes are connected by an edge, the number of pairwise relationships between edges grows quartically with the number of nodes, e.g. 100 nodes would have more than 10^8 relationships between edges. Several other overlapping community detection techniques have been proposed in the general context, a review of which can be found in [94]. However, to the best of our knowledge, none of these overlapping subnetwork identification methods permit multimodal integration.

Another limitation in most community detection techniques is their sole

focus on optimizing a criterion, such as modularity, without considering the variability of the subnetwork assignment under data perturbation or what subnetworks one would get from random graphs to assess statistical significance. In the area of brain activation detection, rigorous statistics has long been employed to control for false positives [67]. It is known that even random networks can have high modularity [95], thus statistically controlling for false node inclusion in subnetwork identification is crucial. A notable exception is the order statistics local optimization method (OSLOM) [96]. OSLOM uses a statistical criterion based on a null model derived from random graphs for assessing the significance of the identified subnetworks, and is shown to outperform many state-of-the-art community detection techniques [94].

1.6.2 Multimodal Approaches

Incorporating anatomical information extracted from dMRI data into the investigation of functional brain dynamics has attracted growing interest since functional synchronization between spatially distinct brain regions is enabled through neural fiber pathways [97, 98]. Most of the early work focused on direct comparisons of anatomical and functional connectivity learned separately from dMRI and fMRI data as explained in Section 1.4.3. The main findings were that brain areas with high AC typically display high FC but the converse does not always hold [98]. However, recent studies have shown that the correlation between FC and AC estimates is rather low with typical data acquisition schemes [65]. As explained in Section 1.4, FC estimates tend to contain many false positive connections whereas AC estimates often suffer from false negatives. These discrepancies may appear as setbacks at first sight, but these inter-modality differences in fact render multimodal integration useful since incorporating AC would reduce the effects of noise-induced FC and vice versa.

The inherent relationship between anatomy and function in the brain revealed by the aforementioned studies implicates the potential benefits of fMRI-dMRI fusion such as improving the sensitivity of data analyses using complementary information or improving our understanding of how brain function is facilitated by anatomy in general. In line with this observation, multimodal integration for joint anatomical and functional connectivity inference have been recently explored to better delineate the architecture of the human brain [99, 100]. It has been shown that multimodal connectivity inference offers higher inter-subject consistency of the detected connection patterns [99] and higher discrimination power between the control and clin-

ical populations [100] compared to analyzing each modality in isolation.

In addition to multimodal connectivity inference studies, multimodal subnetwork extraction problem has also recently started to attract attention. For this problem, a multi-view spectral clustering algorithm has been recently adopted in [52], where FC and AC are considered as two views of the brain. The Laplacian of the similarity matrix of one view is projected onto the eigenspace of the other view in an alternating manner, with spectral clustering subsequently applied on the converged eigenspace to identify subnetworks. A related approach has been used in [101], where the common eigenspace between the similarity matrices of FC and AC is found via joint diagonalization, on which spectral clustering is performed.

1.7 Thesis Contributions

We propose in this thesis novel directions for multimodal brain image analysis for assessing functional segregation and integration in the human brain. We demonstrate the strengths of multimodal approaches over traditional unimodal analyses on carefully designed synthetic datasets as well as real fMRI and dMRI datasets. Our contributions in advancing the state-of-the-art are listed below. We note that the original mathematical notations used in publications that led to this thesis are kept to ease cross-referencing. Hence, there might be some overlaps in notations between different sections. To resolve any ambiguities that might arise due to such overlaps, we have explicitly redefined the notations in each subsection.

1.7.1 Anatomical and Functional Brain Connectivity

- We showed that the consistency between estimates of AC and FC can be increased using streamline tractography on ODFs as opposed to conventional diffusion tensors when estimating AC. We also demonstrated a further improvement when global tractography on ODFs is deployed [P1].
- We showed that part of the inconsistency between estimates of AC and FC can be attributed to the relatively low spatial resolution of dMRI data, often leading to partial volume effects. We proposed to use a super-resolution approach based on dictionary learning for alleviating this problem and showed that our method outperforms traditional interpolation strategies [P2].

- We showed that the way AC is quantified has a significant effect on the agreement between estimates of AC and FC. We demonstrated that fiber count and volume-normalized fiber count are the two metrics that correlate best with FC, and that fiber count has lower bias in favor of shorter fibers than volume-normalized fiber count [P3]. Fiber count is thus used to quantify AC in this thesis.

1.7.2 Multimodal Brain Activation Detection

- We proposed a novel approach that exploits the intrinsic task-rest relationships in brain activity to reduce the confounding effects of RS activity on task activation detection. We showed that removal of the estimated RS modulations from task fMRI data significantly improves activation detection [P4].
- We proposed incorporating AC into brain activation detection to mitigate the effects of noise. We illustrated that incorporating AC significantly increases sensitivity in detecting task activation for healthy subjects. We also showed that this multimodal approach enables the detection of significant group differences between schizophrenia patients and controls that are missed with standard methods [P5].

1.7.3 Multimodal Brain Subnetwork Identification

- We proposed a multimodal integration technique for identifying subnetworks of brain regions that exhibit high inter-connectivity both functionally and anatomically. Our proposed technique allows for subnetwork overlaps, and provides statistical control on false node inclusion in the identified subnetworks. We demonstrated that our technique achieves significantly higher subnetwork identification accuracy than state-of-the-art techniques on synthetic data. On real data, we showed that our method attains improved test-retest reliability on multiple network measures. Further, we showed that the tasks subjects are involved in can be more accurately classified based on the subnetworks employed during different tasks with this multimodal approach compared to unimodal approaches [P6-P8].

Chapter 2

Anatomical and Functional Brain Connectivity

Since anatomical wiring of the brain shapes functional interactions, it is intuitive to expect a strong correlation between AC and FC in the brain. Indeed, a number of studies suggested a strong positive correlation between AC and FC estimates [63, 64, 98]. However, it has been more recently shown that the AC-FC correlation is as low as ~ 0.12 using the most common metrics to quantify AC and FC (fiber count for AC, Pearson’s correlation between fMRI time courses for FC) on typical dMRI (~ 32 gradient directions) and RS fMRI (~ 7 min) data [65]. We hypothesize that this result is in part due to error propagation during the several steps of AC and FC estimation. As explained in Section 1.4.2, the neuroimaging community has largely settled on the estimation of FC using Pearson’s correlation between fMRI time courses following a series of standard preprocessing steps. However, AC estimation using dMRI data is far from being standardized. We thus focused our attention on various steps leading to quantitative AC estimates.

Firstly, accuracy of the reconstructed fiber tracts is often hampered by the inherently low resolution of dMRI data. Currently achievable spatial dMRI resolution is around $2 \times 2 \times 2$ mm³, while the actual neuronal fiber diameter is on the order of $1 \mu\text{m}$ [102]. A voxel can thus comprise several distinct fiber bundles with differing orientations, leading to partial volume averaging [103]. At such locations, diffusion information typically becomes ambiguous, and tractography is often falsely terminated. Therefore, increasing the spatial resolution in dMRI holds great promise towards more accurate delineation of fiber tracts. To this end, we propose to use a super-resolution approach based on dictionary learning for alleviating this problem (Section 2.1). Unlike the majority of existing super-resolution algorithms, our proposed solution does not entail acquiring multiple scans from the same subject which renders it practical in clinical settings and applicable to legacy data. Moreover, this approach can be used in conjunction with any diffusion model.

Secondly, given a set of DW images, the diffusion model and tractogra-

phy algorithm employed directly affect the quality of the reconstructed fiber tracts. Suboptimal choices made at this step may result in erroneous AC estimates regardless of the quality of the underlying data. dMRI data have overwhelmingly been modelled using tensors to date, which approximates the local diffusion process by a single Gaussian probability density function. Due to this simplistic approximation, diffusion tensors can only capture the principal diffusion direction, which makes them prone to errors induced by crossing fibers. With recent advances in dMRI acquisition, a larger number of gradient directions can now be acquired within a practical amount of time, which enables more sophisticated methods that can model multiple fiber directions, such as ODF reconstruction [15], to be readily employed. To better exploit the high angular resolution data, ample efforts have been placed on improving ODF estimation strategies during recent years. However, less attention has been paid to tractography methods, which are equally important for AC estimation. Streamline tractography performed on diffusion tensors or ODFs remains to be the norm. Typically, streamline tractography follows either the principal diffusion directions of diffusion tensors or largest ODF peaks in tracking fibers one at a time. Approaches that consider all fibers in aggregate have recently been explored [42]. Whether these global tractography techniques can help increase the consistency between AC and FC estimates has to date not been investigated. To this end, we analyze the impact of AC estimation strategies on the consistency between AC and FC estimates. Specifically, we show increased AC-FC consistency with streamline tractography on ODFs compared to using conventional diffusion tensors. We also demonstrate a further improvement when global tractography on ODFs is deployed (Section 2.2).

Lastly, the choice of which streamline property to measure and of how to map it into an anatomical connectivity metric are key aspects affecting the AC estimates once the streamlines are reconstructed. To determine which AC metric has the highest potential of being of practical use in multimodal brain image analysis efforts, we compare four commonly used AC metrics in terms of their impact on the relationship between estimates of AC and FC. We demonstrate that fiber count and volume-normalized fiber count are the two metrics that correlate best with FC, and that fiber count has lower bias in favor of shorter distances than volume-normalized fiber count (Section 2.3). Fiber count is thus used to quantify AC in this thesis.

2.1 Super-resolution for dMRI

Due to partial volume averaging, local dMRI observations typically become ambiguous in voxels comprising several distinct fiber bundles, and tractography is often falsely terminated. Therefore, increasing the spatial resolution in dMRI holds great promise towards more accurate delineation of fiber tracts. There are practical limitations in increasing the resolution of the acquired data directly, such as reduced SNR and prolonged scanning time [104]. Such limitations motivate the search for post-processing solutions for increasing the spatial resolution, such as super-resolution techniques.

Super-resolution techniques have been previously adopted to increase the spatial detail in dMRI. In the literature, the term super-resolution is used for two distinct classes of methods which follow different paradigms. The first class of methods are based on performing multiple low-resolution acquisitions, followed by the fusion of the information in these images to generate high-resolution images. To this end, fusing images spatially shifted at subvoxel level [105], as well as fusing multiple anisotropic images with high resolution only along one axis [102, 106] have been explored. In a fairly similar spirit, combining DW images acquired at two different resolutions to infer high-resolution diffusion parameters using a Bayesian model has also been proposed [107]. The inherent drawback of these approaches is the dependence on a specific acquisition protocol, limiting their usability in general settings. The second class of methods do not require multiple acquisitions, and these are typically based on examples or priors about the correspondence between low and high resolution images. Falling in this category, an approach to reconstruct diffusion tensors at a resolution higher than the underlying DW images using a single dMRI acquisition has been recently proposed [108]. Even though this method eliminates the need for multiple acquisitions, it is only geared towards estimating diffusion tensors, and cannot be easily extended to higher order diffusion models such as ODFs. To the best of our knowledge, the only previous work that tackled the problem of super-resolving dMRI data from a single acquisition independent of the diffusion model was by Coupé *et al.* [109]. Specifically, the authors showed that super-resolving $b=0$ (non-diffusion-weighted) image using a locally adaptive patch-based strategy, and using this high-resolution $b=0$ image to drive the reconstruction of DW images outperforms upsampling of dMRI data using classical interpolation methods. Beyond these two classes, a new perspec-

Section 2.1 is adapted from: B. Yoldemir, M. Bajammal, R. Abugharbieh. Dictionary based super-resolution for diffusion MRI. In Proceedings of the MICCAI Workshop on Computational Diffusion MRI, pp. 194-204, Boston, MA, September 2014.

tive to gain spatial resolution in dMRI has been proposed which is termed as super-resolution tract-density imaging (TDI) [110]. This approach is fundamentally different than the aforementioned super-resolution methods in the sense that the aim is to generate high resolution tract density maps through counting the number of tracts present in each element of a subvoxel grid, rather than super-resolving the DW volumes prior to tractography.

We start by presenting our assumed data acquisition model (Section 2.1.1). Given a set of acquired DW volumes, we form a training set that includes the original volumes and a set of corresponding downsampled volumes at double the voxel size. We then construct two over-complete dictionaries from the original and downsampled set of volumes (Section 2.1.2). For a previously unseen input DW volume, we obtain the super-resolution data in two steps. First, we sparsely code the volume against the dictionary learned from the downsampled volumes in the training set. We finally apply the generated sparse code to the dictionary learned from the original resolution set to obtain the super-resolution data (Section 2.1.3).

2.1.1 Acquisition Model

Let v_L be an acquired volume and v_H be the corresponding unobserved higher resolution volume. We assume that the relationship between these two volumes is modelled by [111]

$$v_L = \mathbf{S}\mathbf{B}v_H + n, \tag{2.1}$$

where \mathbf{S} is a downsampling operator, \mathbf{B} is a blurring operator and n is additive white Gaussian noise. We aim to invert this acquisition model to approximate the unobserved higher resolution volume through super-resolution. The maximum-likelihood solution to this problem involves the minimization of $\|\mathbf{S}\mathbf{B}\hat{v}_H - v_L\|_2$, where \hat{v}_H is the estimated high resolution volume. However, the inversion of $\mathbf{S}\mathbf{B}$ is ill-posed [111], hence infinitely many maximum-likelihood solutions exist. We thus cast the problem in a dictionary learning framework instead, as explained next.

2.1.2 Dictionary Construction

We model each 3D patch in dMRI volumes as a sparse linear combination of atoms from a learned dictionary \mathbf{D} . In the proposed approach, we use two dictionaries to capture the correspondence between low and high resolution dMRI volumes. These two dictionaries are learned from the original training dataset and its downsampled version, respectively.

Let \mathbf{v}_O be the set of original training volumes concatenated across scans and \mathbf{v}_D be the corresponding set of downsampled volumes. We extract all overlapping patches in these two sets of volumes, denoted by \mathbf{p}_O and \mathbf{p}_D , respectively. Using \mathbf{p}_O and \mathbf{p}_D , we construct two over-complete dictionaries as follows:

$$\min_{\mathbf{D}_O, \mathbf{D}_D, \mathbf{y}} \sum \|\mathbf{p}_D - \mathbf{D}_D \mathbf{y}\|_2^2 + \sum \|\mathbf{p}_O - \mathbf{D}_O \mathbf{y}\|_2^2 + \psi(\mathbf{y}), \quad (2.2)$$

where $\mathbf{y} = \{\mathbf{y}_{(i,j,k)}\}$ is the set of sparse coding vectors for each image location (i, j, k) , and \mathbf{D}_D and \mathbf{D}_O are the generated over-complete dictionaries of the downsampled and original volumes, respectively [111]. $\psi(\mathbf{y})$ is a regularization term which we set to be $\psi(\mathbf{y}) = \|\mathbf{y}\|_1$, inducing sparsity on the generated coding vector [112]. We note that the same set of coding vectors \mathbf{y} is used for both dictionaries. In other words, the learned atoms of the two dictionaries represent matched pairs. We set the number of atoms in each dictionary to 1000 and the patch size to $3 \times 3 \times 3$ voxels, which were empirically chosen to strike a balance between representation accuracy and overfitting. Note that these two parameters can be optimized via cross-validation, however we did not perform such analysis since our aim is to highlight that AC-FC agreement is affected by the relatively low spatial resolution of dMRI data rather than achieving the highest correlation between the estimates of these two types of connectivity.

2.1.3 Super-Resolved Volume Generation

Let \mathbf{p}_I be the set of low resolution overlapping patches obtained from a previously unseen input volume that we wish to super-resolve. We code \mathbf{p}_I with respect to \mathbf{D}_D as

$$\min_{\mathbf{y}_I} \|\mathbf{p}_I - \mathbf{D}_D \mathbf{y}_I\|_2^2 + \psi(\mathbf{y}_I), \quad (2.3)$$

where \mathbf{y}_I is the set of coding vectors for \mathbf{p}_I , with $\psi(\mathbf{y}_I)$ again being the l_1 norm of \mathbf{y}_I , enforcing sparsity on the coefficients. Once the input volume is sparsely coded using \mathbf{D}_D , we generate a new set of super-resolved patches, \mathbf{p}_S , by applying the sparse coding vector \mathbf{y}_I to \mathbf{D}_O previously constructed from the training data:

$$\mathbf{p}_S = \mathbf{D}_O \mathbf{y}_I. \quad (2.4)$$

We note that this process results in a patch being generated for each voxel. We then reconstruct the super-resolved volume by averaging neighbouring overlapping patches. We used K-singular value decomposition (K-SVD) [113] to construct the dictionaries and orthogonal matching pursuit

[114] to sparsely code the 3D patches. Theoretically, \mathbf{p}_O , \mathbf{p}_D and \mathbf{p}_I can be extracted at once from the volumes of all gradient directions in the DW images. However, we opt to apply super-resolution for each gradient direction separately. This helps circumvent computational limitations that might arise, especially with the increasingly large number of gradient directions acquired in practice.

2.1.4 Experiments

Materials

We validated our method on the publicly available multimodal Kirby 21 dataset.¹ Along with other imaging modalities, this dataset comprises dMRI and RS fMRI scans of 21 subjects with no history of neurological disease (11 men, 10 women, 32 ± 9.4 years old). We summarize the key acquisition parameters below. Further details on data acquisition can be found in [115]. In our experiments, we used 10 subjects for dictionary training, and 10 other subjects for testing.

RS fMRI data of 7-minute duration were collected with a TR of 2 s and a voxel size of 3 mm (isotropic). Preprocessing steps included motion correction, bandpass filtering at 0.01 and 0.1 Hz, and removal of white matter and cerebrospinal fluid confounds. A common approach to brain image analysis involves examining regions of interest (ROIs) as internally coherent units. The representative signal of a given ROI is defined either as the average time course of the voxels in the ROI or the first principal component from a principal component analysis (PCA) performed on the time courses of individual voxels [116]. The main reasons for working on voxel groups instead of voxels include reducing the data dimensionality for computational tractability and controlling for type-I error by reducing the number of statistical tests [117]. Since each voxel is unlikely to function in isolation [118], ROI-based approach offers an efficient alternative to voxel-based analyses. We thus adopted an ROI-based approach instead of a voxel-based approach in this thesis. We divided the brain into 150 parcels by applying Ward clustering [119] on the voxel time courses, which were temporally concatenated across subjects. Parcel time courses were then found by averaging the voxel time courses within each parcel.

The dMRI data had 32 diffusion-weighted images with a b-value of 700 s/mm^2 in addition to a single $b=0$ image, with a voxel size of $0.83 \times 0.83 \times 2.2 \text{ mm}^3$. Since anisotropic voxels were previously shown to be suboptimal for

¹This dataset is available online at: <http://www.nitrc.org/projects/multimodal>

fiber tractography [120], we resampled each volume to 2 mm isotropic resolution prior to any analysis. We also applied a Rician-adapted denoising filter [121] to eliminate nonstationary noise commonly observed in DW images, since our acquisition model assumes Gaussian noise. We then warped our functionally derived group parcellation map to the $b=0$ volume of each subject using FSL [31] to facilitate the computation of fiber count.

Results

We first present a qualitative comparison between the fiber tracts reconstructed from the original (2 mm isotropic) and super-resolved (1 mm isotropic) dMRI data. For ease of interpretation, we chose to employ deterministic streamline tractography with the diffusion tensor model, which is by far the most popular tractography approach to date. However, we highlight that our super-resolution approach can be used with any diffusion model and any tractography method. Tractography was carried out using Dipy [122], with 750,000 seed points for both the original and super-resolution data. We generated the tract-density maps by calculating the total number of fiber tracts present in each voxel. Fig. 2.1(a),(c) and (b),(d) show sample tract-density maps with the original and super-resolved dMRI data, respectively. As observed from these figures, the tract-density maps generated from the super-resolution data clearly convey more spatial information. Fig. 2.1(e),(f) and (g),(h) show the corticospinal tracts extracted using an ROI placed on the brain stem for two representative subjects. It can be observed that fiber tracts reconstructed from the super-resolution data can capture the fan-shape configuration of the corticospinal tract more fully.

To quantify the improvement in tractography with the suggested approach, we analyzed the consistency between measures of intra-subject AC and FC. We estimated AC using the fiber counts between brain region pairs, and FC using Pearson’s correlation between parcel time courses. For each subject, AC and FC are vectors of size $d(d - 1)/2 \times 1$ comprising the corresponding connectivity estimates between each region pair, where d is the number of brain regions. We then calculated Pearson’s correlation between intra-subject AC and FC to quantify the consistency between the two connectivity estimates. This linear model was used instead of higher-order nonlinear models in estimating AC-FC consistency for easier interpretability, and based on the empirical finding that AC and FC in the human brain are roughly linearly related [64, 97]. Using this correlation measure, we compared the proposed super-resolution approach with trilinear and spline interpolation in addition to an alternative super-resolution method; collab-

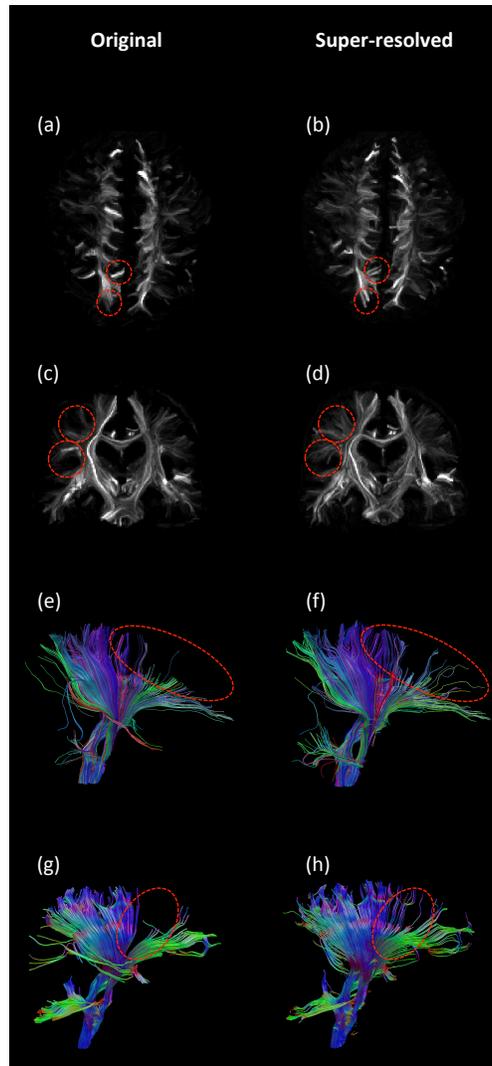


Figure 2.1: Qualitative comparison between the tract-density maps and fiber tracts reconstructed from the original (left) and super-resolved (right) dMRI data. Original dataset has 2 mm isotropic resolution which is super-resolved to 1 mm isotropic resolution. Each row corresponds to a different test subject. Tract-density maps of super-resolved data ((b) and (d)) show markedly improved spatial detail compared to those of original data ((a) and (c)). Corticospinal tracts reconstructed from super-resolved data ((f) and (h)) can capture the fan-shape configuration more accurately than those generated from original data ((e) and (g)).

orative and locally adaptive super-resolution (CLASR) [109]. To the best of our knowledge, CLASR is the only existing single image super-resolution method for dMRI which is independent of the diffusion model employed. Fig. 2.2 shows the AC-FC correlation for each subject tested. Taking the average AC-FC correlation across the group when using the original data as a baseline, the improvement was 5.7% with spline interpolation, 13.6% with CLASR, and 27.1% with our proposed method. On the other hand, there was a 6.3% decrease in the correlation when trilinear interpolation was used. The difference in the performance of our method and every other method tested was found to be statistically significant at $p < 0.01$ based on the Wilcoxon signed-rank test, showing its potential for enhanced anatomical connectivity assessment. Our results thus suggest that low spatial resolution of dMRI data can partially account for the low AC-FC correlation, and statistically significant improvements can be achieved using super-resolved dMRI data.

To investigate why trilinear interpolation resulted in a lower AC-FC correlation compared to the original data, we calculated the number of tracts reconstructed with each method. The local intra-parcel connections were excluded since they have no effect on AC-FC correlation. Fig. 2.3 shows the number of inter-parcel tracts averaged across the group along with the corresponding standard deviations. As observed from this figure, performing tractography on volumes upsampled with trilinear interpolation resulted in a lower number of tracts compared to the original volumes, even though the same number of seed points were used to initiate tracking for all of the methods we compared. We speculate that the reason of this phenomenon is the additional partial volume effects introduced by the blurring of the data during trilinear interpolation, which hamper the tractography quality. Spline interpolation, however, is known to cause less blurring compared to trilinear interpolation, and our results suggest that upsampling dMRI data using spline interpolation can be beneficial for tractography. The overall trend of inter-parcel tract counts closely resembles to that of the AC-FC correlation, with our proposed method outperforming all other methods tested. This shows that dictionary based super-resolution is a viable post-processing solution for dMRI that can help in mapping the white matter brain architecture more accurately.

2.1.5 Discussion

Low spatial resolution is a known limitation of dMRI, which often hinders the performance of tractography significantly. We proposed the use of a

2.1. Super-resolution for dMRI

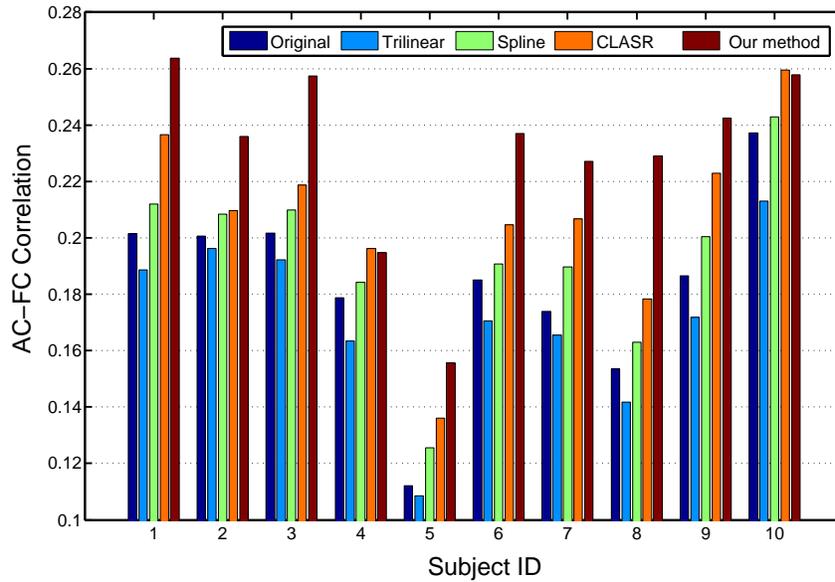


Figure 2.2: AC-FC correlation for 10 subjects with AC estimated from the data at its original resolution (2 mm isotropic), and high-resolution data (1 mm isotropic) obtained using trilinear interpolation, spline interpolation, CLASR and the proposed method. Our method outperforms all other methods tested for eight of the subjects, and performs comparable to CLASR for two subjects (subjects 4 and 10).

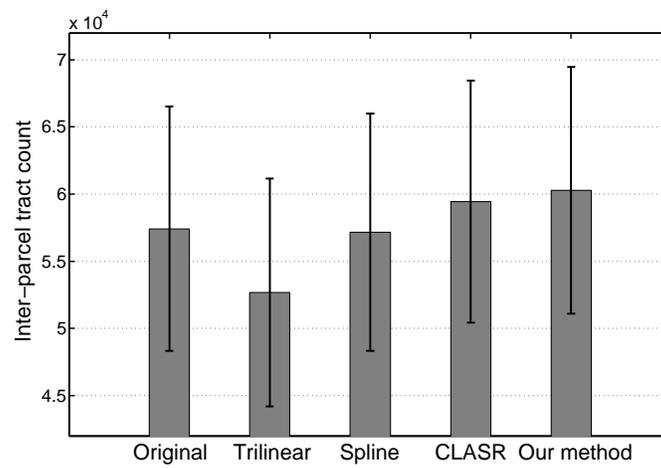


Figure 2.3: Number of inter-parcel tracts reconstructed from the data at its original resolution (2 mm isotropic), and high-resolution data (1 mm isotropic) obtained using trilinear interpolation, spline interpolation, CLASR and the proposed method. Intra-parcel tracts are not included here since they do not contribute to AC-FC correlation. We emphasize that tractography is initiated with the same number of seeds (750,000) for each method.

simple yet effective super-resolution technique in dMRI to capture a more accurate portrayal of the white matter architecture. The advantage of this method is three-fold. First, the super-resolved DW images can be used with any diffusion model as permitted by the number of gradient directions in the original dataset. Second, this method does not rely on repeated acquisitions from the same subject, allowing it to be used with legacy data and under various clinical acquisition schemes. Third, this method may still be readily applied when the imaging protocol involves multiple acquisitions, as an additional step after reconstructing a single image from multiple low resolution acquisitions [102, 105, 106]. Quantitatively, we demonstrated that AC-FC consistency can be markedly increased with the use of our approach in estimating AC. We also qualitatively illustrated that the gain in spatial resolution remarkably improves the fiber tracts and tract-density maps generated. Taken collectively, our results suggest that dictionary based super-resolution holds great promise in enhancing the spatial resolution in dMRI, without requiring additional scans or any modifications of the acquisition protocol.

2.2 Estimating Anatomical Connectivity

The success of techniques that integrate dMRI and fMRI for multimodal connectivity estimation highly depends on the consistency between the AC and FC estimates. To this end, we analyze the impact of AC estimation strategies on the consistency between AC and FC estimates. Specifically, we compare conventional streamline tractography on conventional diffusion tensors and ODFs. Our rationale is that ODFs can model multiple fiber directions, thus may lead to increased AC-FC consistency by more accurately depicting the anatomical wiring structure of the brain. Further, we compare streamline tractography against global tractography [42]. Global tractography accounts for spatial interactions among fiber orientation estimates by reconstructing all fibers simultaneously and finding the configuration that best explains the data [42]. Global tractography thus alleviates error accumulation along reconstructed tracts, which is a major drawback of the streamline approach. We hence hypothesize that fiber tracts extracted with global tractography would better resemble the macroscale neural circuitry,

Section 2.2 is adapted from: B. Yoldemir, B. Ng, R. Abugharbieh. Effects of tractography approach on consistency between anatomical and functional connectivity estimates. In Proceedings of the IEEE International Symposium on Biomedical Imaging (ISBI'14), pp. 250-253, Beijing, China, April 2014.

and thus provide higher consistency between AC and FC estimates. We use the most common methods for estimating FC and quantifying AC, namely Pearson’s correlation between fMRI time courses and fiber count, respectively. We focus our attention on AC estimation, i.e. diffusion modelling and tractography.

2.2.1 Modelling of Local Diffusion Properties

When local diffusion is modelled using diffusion tensors, only the dominant diffusion direction can be recovered at each voxel. In contrast, ODFs are more sensitive in capturing intricate intravoxel fiber geometries. ODFs are typically estimated from Q-ball imaging reconstruction of high angular resolution diffusion imaging (HARDI) data acquired on one or multiple spherical shells in q-space [15]. We opt to use CSA-ODF [41] since this reconstruction method intrinsically provides sharp ODFs, which enables multiple intravoxel fiber orientations to be more easily resolved. Moreover, this method takes advantage of the additional information provided by multi-shell acquisition, which has been empirically shown to better resolve fiber crossings compared to single q-shell models [41]. Using CSA-ODF thus better exploits the multiple q-shells available in the data used in our experiments. Further details on CSA-ODF estimation can be found in [41].

2.2.2 Tractography

Streamline Tractography

Conventional streamline tractography is a deterministic approach, in which a single fiber trajectory is generated for each seed point [36]. Seed points are defined on a dense grid at the voxel or subvoxel level. Fiber tracts are then reconstructed by bidirectionally traversing the direction of maximum diffusion, i.e. the largest ODF peak or the principal diffusion direction in the case of diffusion tensors. In this work, we use Euler’s method for numerically performing tractography. Starting from a seed point \mathbf{p}_0 , fiber tracts are reconstructed using

$$\mathbf{p}_{n+1} = \mathbf{p}_n + \mathbf{v}(\mathbf{p}_n)\Delta s, \tag{2.5}$$

where \mathbf{p}_n is the position of the propagator at step n , $\mathbf{v}(\mathbf{p}_n)$ is the propagation direction at \mathbf{p}_n , and Δs is the fixed step size [122]. To generate smooth tracts, we compute $\mathbf{v}(\mathbf{p}_n)$ using trilinear interpolation which combines the information of voxels within the 8-neighbourhood of \mathbf{p}_n [122].

Global Tractography

In contrast to the streamline approach, global tractography reconstructs all fiber tracts simultaneously, which enables fiber trajectories to be jointly considered in determining the most possible fiber configuration. In this approach, short fiber segments are bridged together to form the set of fiber tracts that best explains the measured dMRI data. This globally optimal tract set is found by minimizing an energy function consisting of an internal and an external energy term: $E(\mathbf{m}) = E_{int}(\mathbf{m}) + E_{ext}(\mathbf{m}, \mathbf{D})$. Here, \mathbf{m} is the set of all short fiber segments and \mathbf{D} is the observed dMRI signal. E_{int} encourages the short fiber segments having consistent orientations to form chains and E_{ext} forces the set of fiber tracts to be in agreement with \mathbf{D} . We note that random changes are introduced into the reconstructed fiber set to avoid getting trapped in local minima. These random changes are accepted or rejected through a Metropolis-Hastings sampler [42]. This global approach is particularly beneficial in resolving cases where local diffusion characteristics are ambiguous due to crossing fibers. Streamline tractography tends to terminate tracking at such locations, whereas global tractography exploits the geometry of surrounding fibers by considering the whole-brain fiber configuration in aggregate to compensate for the lack of reliable local information.

2.2.3 Experiments

Materials

Minimally preprocessed RS fMRI and dMRI data from 40 unrelated healthy subjects of the Human Connectome Project (HCP) dataset [123] were used.² The dataset comprised four RS fMRI scans of 15 minutes per subject, with a TR of 0.72 s and a voxel size of 2 mm (isotropic). The dMRI data had a voxel size of 1.25 mm (isotropic), 3 shells (b=1000, 2000 and 3000 s/mm²) and 288 gradient directions. Further details on acquisition can be found in [123].

In addition to the minimal preprocessing that was already performed on the HCP RS fMRI data which included gradient distortion correction, motion correction, spatial normalization and intensity normalization [124], we further regressed out motion artifacts, white matter and cerebrospinal fluid confounds, and principal components of high variance voxels found using CompCor [125]. Band-pass filtering (0.01 Hz - 0.1 Hz) was subsequently ap-

²This dataset is available online at: <http://www.humanconnectome.org/data>

plied. To define ROIs, we functionally divided the brain into $d=200$ parcels by temporally concatenating RS fMRI voxel time series across scans of all subjects and applying Ward clustering [119]. This value of d provides higher functional homogeneity than typical anatomical atlases, while maintaining reasonable ROIs to time samples ratio for reliable connectivity estimation. The reason why we used a higher number of ROIs than that in Section 2.1.4 is that HCP dataset has higher resolution, hence facilitates a more fine-grained parcellation. ROI time series were then generated by averaging voxel time series within each ROI.

The HCP dMRI data, which have been corrected for EPI distortion, eddy current, gradient nonlinearity and motion artifacts [124], were downsampled to 2 mm isotropic resolution to reduce the computational cost. Streamline (on both diffusion tensors and ODFs) and global tractography on ODFs were carried out using Dipy [122] and MITK [126], respectively. For all methods, approximately 65,000 fibers per subject were reconstructed. The functionally derived group parcellation map was warped onto the dMRI space of each subject to facilitate the computation of fiber count.

Results

We first present a qualitative comparison between streamline and global tractography for a representative subject. Fig. 2.4 shows the tracts running through an ROI within the corpus callosum as estimated using streamline tractography on diffusion tensors and ODFs, and global tractography on ODFs (in yellow, green and red, respectively). We chose this ROI since most tractography algorithms have difficulties tracking the callosal projections [42]. As apparent in Fig. 2.4, streamline tractography only captured the dominant U-shaped callosal radiation. This result is expected for the case of diffusion tensors since the principal diffusion directions of the tensors are ambiguous at fiber crossings. However, the same pattern was observed even with diffusion modelled using CSA-ODF, which is able to delineate multiple intravoxel diffusion directions. This shows that streamline tractography is unable to make full use of the information in the ODFs. In contrast, global tractography was able to better exploit the information in CSA-ODF and tracked the lateral transcallosal fibers.

Quantitatively, the AC-FC correlation, computed by finding the Pearson's correlation between intra-subject AC and FC estimates and averaging these values across the group, was 0.2234 ± 0.018 with streamline tractography on ODFs and 0.1643 ± 0.012 with streamline tractography on diffusion tensors. The difference was found to be statistically significant at $p < 0.01$

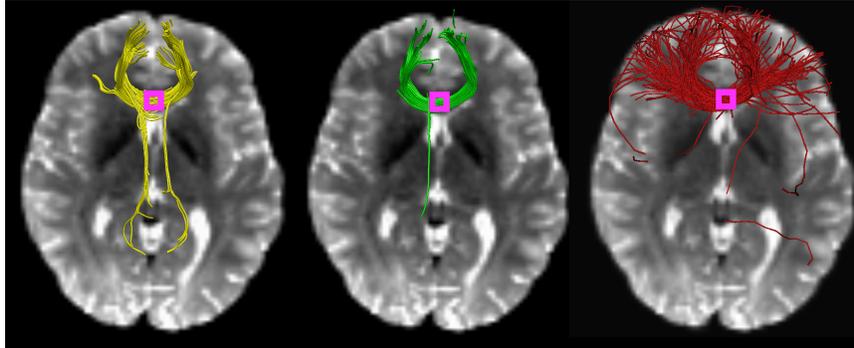


Figure 2.4: Fiber tracts reconstructed for a representative subject using streamline tractography on diffusion tensors (left), streamline tractography on ODFs (middle), and global tractography on ODFs (right) with an ROI placed in the corpus callosum (shown with a square). While streamline tractography results were limited to the callosal radiation, global tractography also tracked bilateral projections.

based on the Wilcoxon signed rank test, showing that using ODFs can notably increase AC-FC consistency over diffusion tensors. With global tractography on ODFs, AC-FC consistency further increased to 0.2619 ± 0.018 . The difference between global and streamline tractography on ODFs was statistically significant at $p < 0.01$. These findings confirm our hypothesis that global tractography can better estimate the macroscale neural circuitry, leading to higher agreement between estimates of AC and FC, though the gain compared to using enhanced diffusion modelling, i.e. ODFs over tensors, was smaller. We speculate that uncertainty in fiber endpoint locations is the limiting factor. Global tractography is driven by the measurements in areas with anisotropic diffusion, whereas diffusion around white-grey matter boundaries is near-isotropic. Thus, reconstructed fibers might not terminate in the correct location. We hence argue that modelling fiber endpoint uncertainty is an under-investigated problem that has high potential in improving AC estimation.

2.2.4 Discussion

Streamline tractography, which remains to be the most commonly used tractography approach, is known for its inability to resolve fiber crossings and premature termination of the tracking process. On high quality multimodal data, we demonstrated that this can lead to underestimation of the corre-

2.3. Quantifying Anatomical Connectivity

lation between estimates of AC and FC. We illustrated that this problem is more pronounced when streamline tractography is performed on conventional diffusion tensors as opposed to ODFs. We also showed that AC-FC consistency can be improved by reconstructing the fiber tracts using global tractography, which provides higher stability against noise and imaging artifacts by jointly incorporating the geometry of all fiber tracts.

Even though we showed that the consistency between AC and FC measures can be improved by the choice of AC estimation method, higher correlation values have been previously reported in the literature [50, 63, 64]. There are four possible reasons of this discrepancy. Firstly, we report our values over all brain region pairs, instead of excluding absent anatomical connections as was the case in [63] or dividing the connection structure into subsets as was done in [64]. Since anatomically unconnected brain region pairs can still show high FC [97], analyzing only anatomically connected brain regions implies discarding part of AC-FC inconsistency, which can potentially inflate the AC-FC correlation values. Secondly, we do not apply any post-processing steps on the AC estimates and directly use fiber counts to assess the consistency between AC and FC measures. Fiber counts were resampled into a Gaussian distribution in [63], which increased the average AC-FC correlation from 0.156 to 0.36. Even though this strategy increased the consistency empirically, such resampling has no theoretical or physiological grounding. Thirdly, we do not apply distance correction to weigh down the effect of longer anatomical connections in generating AC estimates as was done in [50]. Such distance correction would lead to a positive bias in AC-FC correlation since proximal brain region pairs tend to have higher FC compared to distal region pairs on average [97]. Finally, we do not apply global RS fMRI signal regression unlike [50, 63], since global signal regression artificially introduces negative RS correlations [127]. We observed that the AC-FC correlation increased by more than 0.05 for both streamline and global tractography when the global signal was removed. However, whether this increase has a neural basis remains to be a point of debate.

2.3 Quantifying Anatomical Connectivity

Section 2.3 is adapted from: M. Bajammal, B. Yoldemir, R. Abugharbieh. Comparison of structural connectivity metrics for multimodal brain image analysis. In Proceedings of the IEEE International Symposium on Biomedical Imaging (ISBI'15), pp. 934-937, Brooklyn, NY, USA, April 2015.

Multimodal analyses typically require the quantification of connectivity, and inaccurate connectivity metrics may lead to flawed inferences. In contrast to quantification of FC, no well-analyzed and accepted metric exists for quantifying AC. To this end, we present a comparison of the commonly used AC metrics assessed from the perspective of the largely accepted inherent relationship between brain anatomy and function.

As described in Section 1.4.1, the most commonly used AC metrics include (i) fiber count, (ii) fiber count normalized by the total volume of the region pairs they connect, (iii) total length of reconstructed fiber tracts, and (iv) average FA along fiber tracts connecting regions. Among these, arguably the most popular metric has been fiber count to date. We reiterate that all of these AC metrics are confounded by several factors such as the number of seeds used for tracking the fibers and the structural features of the pathway (e.g. curvature, length and width). Nonetheless, these metrics are increasingly used in the literature driven by a practical need for quantifying the degree of connection between brain regions. We thus perform a quantitative analysis to determine which AC metric has the highest potential of being of practical use in multimodal brain image analysis efforts, with the understanding that these AC metrics do not truly convey the “strength” of the anatomical connections.

2.3.1 Anatomical Connectivity Metrics

Let $r_{i,j}^k$ be the k^{th} reconstructed fiber between a pair of anatomically connected regions P_i and P_j . We consider four widely used anatomical connectivity measures in this thesis: fiber count ($f_{i,j}$), fiber count normalized by the total volume of the connected regions ($N_{i,j}$), total length of fibers connecting region pairs ($L_{i,j}$), and average FA along the fibers. For each subject, we compute the Pearson’s correlation between the FC and AC estimates to quantify the AC-FC relationship for each AC metric. More formally, the metrics can be expressed as

$$N_{i,j} = \frac{f_{i,j}}{V(P_i) + V(P_j)} \quad (2.6)$$

$$L_{i,j} = \sum_k \ell(r_{i,j}^k), \quad (2.7)$$

where $V(\cdot)$ is the volume of the corresponding region, and $\ell(r_{i,j}^k)$ is the length of $r_{i,j}^k$.

Prior to the computation of AC metrics, we reconstruct the fibers via global tractography on CSA-ODF using MITK [126]. Global tractography

was chosen over the more common streamline tractography since we have shown in Section 2.2 that it facilitates higher AC-FC consistency.

2.3.2 Experiments

Materials

dMRI, RS fMRI and task fMRI data from HCP dataset were used in the experiments [123]. Details and preprocessing steps of dMRI and RS fMRI data and the generation of the brain parcellation map are presented in Section 2.2.3. Each subject had 7 task fMRI datasets (two 2-5 minute scans each for working memory, gambling, motor, language, social cognition, relational processing, and emotional processing) obtained using the same acquisition parameters as the RS fMRI data. Preprocessing of the task fMRI data closely followed that of RS fMRI, except a high-pass filter at 1/128 Hz was used instead of a bandpass filter.

Results

The results of the AC-FC correlation for RS and task fMRI are shown in Figs. 2.5 and 2.6. As observed from these figures, average FA has lower correlation with FC compared to the rest of the examined AC metrics (fiber count, volume-normalized fiber count, and total fiber length). This is true for both RS and task fMRI. We speculate that the observed low average FA correlation can be attributed to the large number of factors affecting local diffusion anisotropy [53]. These results also show that volume-normalized fiber count has the highest correlation with FC compared to the rest of the examined AC metrics for both RS and task fMRI. The pairwise differences between AC-FC correlation assessed using normalized fiber count and other AC metrics were found to be statistically significant at $p < 0.001$ based on the Wilcoxon signed rank test. Our results thus imply that the compensation for the differences in number of fibers due to the variable size of brain regions yields better depiction of anatomical networks.

We also note that the relatively consistent AC-FC correlation levels across a variety of tasks and RS data support the notion that AC forms the backbone of the brain connectivity around which functional reorganization occurs to respond to different tasks. A diverse repertoire of functional brain connectivity patterns can thus arise constrained by the same anatomical substrate.

AC-FC correlation is not the only factor to consider while choosing an AC metric. It is known that proximal brain region pairs tend to have higher

2.3. Quantifying Anatomical Connectivity

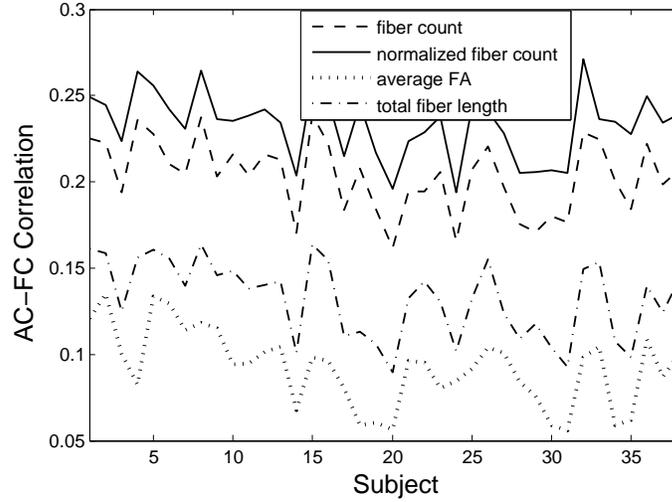


Figure 2.5: Correlation between AC and FC using four commonly used AC metrics. FC was estimated from RS fMRI data, thus depicts intrinsic connections.

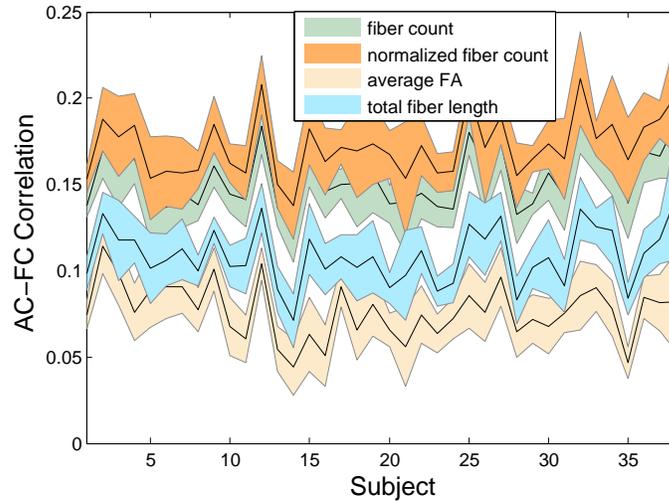


Figure 2.6: Correlation between AC and FC using four commonly used AC metrics. FC was estimated from task fMRI data acquired during 7 different tasks. The narrow shaded bands in represent the standard deviation of AC-FC correlation across different tasks.

2.3. Quantifying Anatomical Connectivity

Table 2.1: Comparison of fiber length bias for AC metrics.

AC metric	Fiber length bias
Fiber count	0.14
Normalized fiber count	0.17
Average FA	0.15
Total fiber length	0.10

FC compared to distal region pairs on average [97]. An AC metric considering mostly shorter fibers and ignoring the long fibers might thus lead to an inflated AC-FC consistency, even though there is no reason to believe long fibers should be weighed less than short fibers when quantifying AC. Indeed, AC metrics have an inherent bias in favor of shorter fibers since shorter fibers are easier to be reconstructed. This in turn induces a dependency of the AC-FC relationship on distance. To quantify the extent of this phenomenon, we computed the difference between group-averaged AC-FC correlation values of the longest and shortest 20% of all fibers, which we refer to as *fiber length bias*. As shown in Table 2.1, total fiber length metric has the smallest fiber length bias. This finding confirms that total fiber length indeed partially corrects for length bias fiber count metrics suffer from. However, this metric results in a considerably lower AC-FC correlation than fiber count and normalized fiber count. We thus used fiber count to quantify AC in this thesis, which provides a lower fiber length bias than normalized fiber count, and a higher AC-FC correlation than total fiber length.

2.3.3 Discussion

We presented a comparison of four AC metrics computed from tractography results with respect to their relationship to FC. Among the metrics considered, we showed that volume-normalized fiber count has the highest correlation with FC for both RS and task data. On the other hand, our results showed that average FA has the lowest correlation with FC. We speculate that the reason of this low correlation is the non-specificity of FA, with several inadvertent factors (such as axonal density, axonal ordering, and amount of myelination) modulating it along with the reliability of local diffusion information [53]. In addition, we also demonstrated that total fiber length metric reduces the fiber length bias associated with shorter fibers. To strike a balance between high AC-FC consistency and low fiber length bias, we chose to employ fiber count to quantify AC in this thesis.

2.4 Summary

In this chapter, we investigated several steps leading to quantitative AC estimates, particularly from the perspective of their effect on the consistency between AC and FC. We showed that part of the inconsistency between AC and FC estimates can be attributed to the analysis methods used to obtain measures of AC. In particular, we showed that increasing the spatial resolution of DW images using a dictionary based super-resolution approach improves the reconstructed fiber tracts and tract-density maps qualitatively, and leads to improved AC-FC consistency. Further, we showed that modelling local diffusion using ODFs over diffusion tensors, and reconstructing the fibers using global tractography over streamline tractography help better resolve fiber crossings around corpus callosum, and provide higher AC-FC consistency over the whole brain. Finally, we showed that quantifying AC using fiber count or volume-normalized fiber count results in higher AC-FC consistency compared to average FA and total fiber length. Our experiments revealed that fiber count has less bias in favor of shorter fibers than volume-normalized fiber count, and thus fiber count is the metric used to quantify AC in this thesis.

Chapter 3

Multimodal Brain Activation Detection

To map brain areas to function, the standard analysis approach models fMRI observations as a combination of expected temporal responses using GLM [26]. However, the strong noise in fMRI data arising from confounds, such as scanner drifts, motion artifacts, and physiological effects, greatly hampers reliable detection of brain activation. Exploring other sources of extra information to regularize activation detection has thus proven to be beneficial. To model the integrative property of the brain, which is known to facilitate brain function [128], the use of local neighbourhood information has been proposed to regularize activation detection [69, 72]. Although such methods help suppress false spatially-isolated activations by encouraging spatial continuity, they completely ignore long-range functional interactions. The incorporation of functional connectivity information into task activation detection has also been put forth [74], but the approach taken estimates both activation effects and functional connectivity from the same dataset, hence the information gain might be limited. Other works investigated the use of RS functional connectivity information to inform task activation detection [79] as motivated by the similarity of RS subnetworks and those engaged during task [78]. Adding to the arsenal of prior-informed brain activation detection methods, we propose two novel multimodal activation detection approaches in this chapter.

Firstly, we explore the feasibility of deconfounding the effects of RS activity from task fMRI data. Recent studies suggest that part of the noise in task fMRI data actually pertains to ongoing RS brain activity. Due to the sporadic nature of RS temporal dynamics, pre-specifying temporal regressors to reduce the confounding effects of RS activity on task activation detection is far from trivial. We propose a novel approach that exploits the intrinsic task-rest relationships in brain activity for addressing this challenging problem. With an approximate task activation pattern serving as a seed, we first infer areas in the brain that are intrinsically connected to this seed using RS fMRI data. We then apply PCA to extract the RS component within the

task fMRI time courses of the identified intrinsically-connected brain areas. Using the learned RS modulations as confound regressors, we re-estimate the task activation pattern, and repeat this process until convergence. On real data, we show that removal of the estimated RS modulations from task fMRI data significantly improves activation detection (Section 3.1).

Secondly, we propose incorporating anatomical connectivity into brain activation detection as motivated by how the functional integration of distinct brain areas is facilitated via neural fiber pathways. We formulate activation detection as a probabilistic graph-based segmentation problem with fiber networks estimated from dMRI data serving as a prior. Our approach is reinforced with a data-driven scheme for refining the connectivity prior to reflect the fact that not all fibers are necessarily deployed during a given cognitive task as well as to account for false fiber tracts arising from limitations of dMRI tractography. Validating on real clinical data collected from schizophrenia patients and matched healthy controls, we show that incorporating anatomical connectivity significantly increases sensitivity in detecting task activation in controls compared to existing univariate techniques. Further, we illustrate how our model enables the detection of significant group activation differences between controls and patients that are missed with standard methods (Section 3.2).

3.1 Modelling Intrinsic Functional Connectivity

The frequency range at which RS activity resides is typically found to be between 0.01 to 0.1 Hz [76], which overlaps with the stimulus frequencies employed in most task fMRI studies. Thus, standard high pass filtering e.g. at 1/128 Hz, which is the default cutoff frequency in the SPM software, for removing temporal drifts in task fMRI data, would not account for ongoing RS modulations. Also, unlike task-evoked responses, which are time-locked to stimulus, the time at which RS activity peaks and troughs is difficult to predict. Hence, pre-specifying temporal regressors to model RS activity is non-trivial.

Albeit its seemingly sporadic temporal dynamics, RS activity is not random [76]. Rather, strong synchrony in RS modulations between specific brain areas has been observed in numerous studies [76]. In fact, many of

Section 3.1 is adapted from: B. Yoldemir, B. Ng, R. Abugharbieh. Deconfounding the effects of resting state activity on task activation detection in fMRI. In Proceedings of the MICCAI Workshop on Multimodal Brain Image Analysis, LNCS, vol. 7509, pp. 51-60, Nice, France, October 2012.

the detected RS subnetworks exhibit high resemblance to subnetworks seen in task experiments [78]. Further supporting this finding is a recent work [79] that demonstrated enhanced sensitivity in task activation detection by incorporating an RS connectivity prior. In addition, studies that jointly examined RS fMRI and dMRI data indicate an anatomical basis for RS activity [97, 98]. Thus, the spatial patterns of RS subnetworks would presumably be constrained by the underlying fiber pathways [97, 98]. Taken together, these findings suggest that there is spatial structure in RS activity and that the spatial structure of ongoing RS activity during task would likely remain similar to that during rest [75].

In this section, we propose a novel approach for RS activity removal in a general setting. The key challenge to this problem is that RS activity is internally-driven by the brain, as opposed to being evoked by external stimulus with known timing. It is thus not obvious how the temporal dynamics of RS modulations that occurred during task performance can be determined a priori. Representative time courses reflective of ongoing RS activity must hence be extracted from the task fMRI data itself. Since the brain comprises multiple subnetworks [78], the RS modulations superimposed on the fMRI responses of the task-activated brain areas would be specific to the RS subnetwork in which these brain areas belong. Extracting RS activity from task fMRI data would thus require knowing the parts of the brain that are activated and their intrinsically-connected areas, which introduces a circular problem. To deal with this issue, we employ an iterative strategy in which we first apply seed-based analysis [22] with an approximate task activation pattern being the seed to infer the intrinsically-connected brain areas from RS fMRI data. Assuming the spatial structure of RS subnetworks is sustained during task performance [75], we extract RS modulations from the task fMRI time courses of the identified brain areas and re-estimate the task activation pattern with the learned RS activity as confound regressors. The overall method is summarized in Fig. 3.1.

3.1.1 Group Activation Detection

Our approach begins with the estimation of an approximate task activation pattern (i.e. without accounting for ongoing RS modulations). A standard GLM is first applied to compute the intra-subject activation effects [26]:

$$\mathbf{Y}^i = \mathbf{X}^i \boldsymbol{\beta}^i + \mathbf{E}^i, \quad (3.1)$$

where \mathbf{Y}^i is a $t \times d$ matrix containing the task fMRI time courses of d brain areas of subject i , $\mathbf{X}^i = [\mathbf{X}_{task} | \mathbf{X}_{confounds}^i]$ is a $t \times p$ matrix with \mathbf{X}_{task}

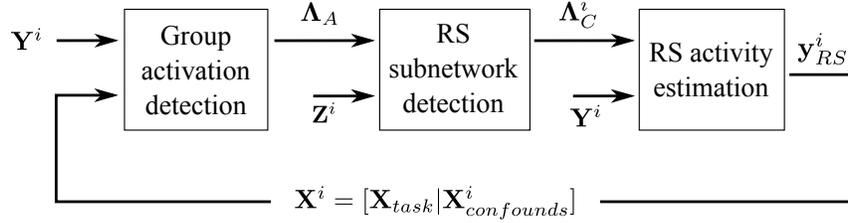


Figure 3.1: Graphical depiction of proposed RS activity removal approach. $\mathbf{X}^i = [\mathbf{X}_{task} | \mathbf{X}_{confounds}^i]$ is a regressor matrix, where \mathbf{X}_{task} corresponds to task regressors and $\mathbf{X}_{confounds}^i$ corresponds to confound regressors specific to subject i . \mathbf{Y}^i are the task fMRI time courses of subject i . Λ_A is the set of activated brain areas common across a group of subjects. \mathbf{Z}^i are the RS fMRI time courses of subject i . Λ_C^i is the set of brain areas estimated to be intrinsically connected to Λ_A for subject i . \mathbf{y}_{RS}^i is the estimated RS activity time course of subject i , which is entered into \mathbf{X}^i as a confound regressor for re-estimating the group activation pattern Λ_A . The three steps: group activation detection, RS subnetwork detection, and RS activity estimation, are repeated until Λ_A stabilizes.

corresponding to task regressors and $\mathbf{X}_{confounds}^i$ corresponding to confound regressors specific to subject i , β^i is a $p \times d$ activation effect matrix to be estimated, and \mathbf{E}^i is a $t \times d$ residual matrix. Due to the strong noise in task fMRI data, activation patterns estimated at the intra-subject level might be inaccurate [129]. Therefore, we opt to combine information across subjects in generating a group activation map, which is then used as a seed for identifying intrinsically-connected brain areas. To infer group activation, we apply a max-t permutation test [67] on β^i of all subjects, which implicitly accounts for multiple comparisons and provides strong control over false detections. Group activation is declared at a p-value threshold of 0.05. We denote the set of detected brain areas as Λ_A .

3.1.2 RS Subnetwork Detection

With the detected group activation pattern taken as a seed, our goal is to identify brain areas that belong to the same RS subnetwork as the seed, so that we can extract RS modulations specific to this RS subnetwork. To proceed, we first average the RS fMRI time courses within the detected activated brain areas in generating a seed time course for each subject.

3.1. Modelling Intrinsic Functional Connectivity

We then apply the standard seed-based analysis [22] to find brain areas intrinsically-connected to this seed for each subject i :

$$\mathbf{Z}_{\sim S}^i = \mathbf{z}_s^i \mathbf{w}^i + \mathbf{\Omega}^i, \quad (3.2)$$

where $\mathbf{Z}_{\sim S}^i$ is an $n \times (d - |\mathbf{\Lambda}_A|)$ matrix containing the RS fMRI time courses of all brain areas except those in $\mathbf{\Lambda}_A$, $|\mathbf{\Lambda}_A|$ is the number of brain areas in $\mathbf{\Lambda}_A$, \mathbf{z}_s^i is an $n \times 1$ vector containing the seed time course, \mathbf{w}^i is a $1 \times (d - |\mathbf{\Lambda}_A|)$ vector with each element reflecting the correlation between the seed and each brain area not in $\mathbf{\Lambda}_A$, and $\mathbf{\Omega}^i$ is an $n \times (d - |\mathbf{\Lambda}_A|)$ residual matrix. Statistical significance of each element of \mathbf{w}^i is declared at a p-value threshold of 0.05 with false discovery rate (FDR) correction [130] to account for multiple comparisons. FDR correction is used instead of max-t permutation test due to correlations between brain volumes at adjacent time points, which violates the independent sample assumption in max-t permutation test [67]. A max-t permutation test could be applied to identify brain areas that are significantly correlated with the seed at the group level. However, compared to task-based experiments, RS experiments are less prone to motion artifacts, which constitute a major part of fMRI noise. Reliable RS subnetworks could thus potentially be extracted at the intra-subject level. We hence opt to perform intra-subject seed-based analysis to retain subject-specific information. We denote the set of brain areas significantly correlated with the seed as $\mathbf{\Lambda}_C^i$.

3.1.3 RS Activity Estimation and Removal

After finding the set of brain areas that is intrinsically-connected to the estimated task activation pattern for each subject i , the next step is to extract the RS components from the task fMRI time courses of these brain areas, which we denote as \mathbf{Y}_C^i . To target the specific frequency range at which RS activity resides, we first band-pass filter each column of \mathbf{Y}_C^i at cut-off frequencies of 0.01 and 0.1 Hz. Since the estimated task activation pattern is only an approximation without accounting for the confounding effects of RS activity, some intrinsically-connected brain areas might in fact be activated considering the resemblance between task and RS subnetworks [78]. Thus, \mathbf{Y}_C^i might contain task signals. To remove the task-related response in \mathbf{Y}_C^i , we apply PCA through eigen-decomposition to separate \mathbf{Y}_C^i into task and non-task components:

$$\mathbf{C}^i = \mathbf{U}^i \mathbf{D}^i \mathbf{U}^{iT}, \quad (3.3)$$

3.1. Modelling Intrinsic Functional Connectivity

where \mathbf{C}^i is the $d \times d$ covariance matrix of \mathbf{Y}_C^i , \mathbf{U}^i is a $d \times d$ matrix containing the eigenvectors of \mathbf{C}^i , and \mathbf{D}^i is a $d \times d$ matrix containing the eigenvalues of \mathbf{C}^i along the diagonal. The columns of \mathbf{U}^i are ordered such that the first column, \mathbf{U}_1^i , corresponds to the largest eigenvalue. To identify the task-related components, we compute the correlation between each column of \mathbf{U}^i and the task regressor. Statistical significance in correlation is declared at a p-value threshold of 0.05 with FDR correction. For the data examined in this work, the task regressor is found to be most significantly correlated with \mathbf{U}_1^i . This high correlation between \mathbf{U}_1^i and the task stimulus, as shown in Fig. 3.2, signifies a definite need for task response removal from \mathbf{Y}_C^i . We remove the task components by reconstructing \mathbf{Y}_C^i with the significantly correlated columns of \mathbf{U}^i discarded.

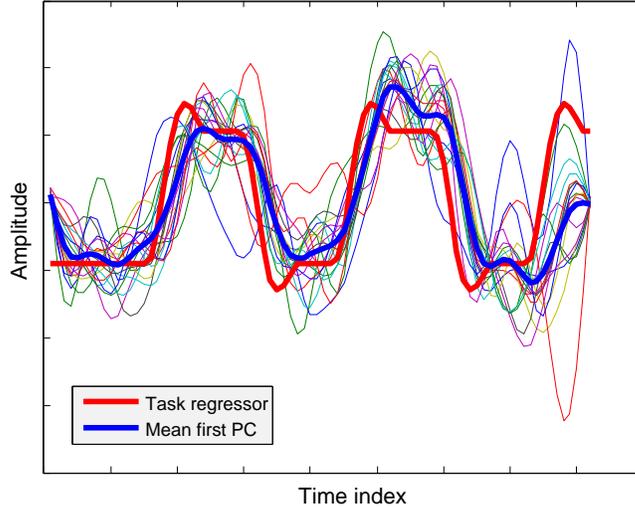


Figure 3.2: Dominant principal component (PC) extracted from time courses of intrinsically-connected brain areas. The thick red and blue lines correspond to task regressor and mean dominant PC across subjects. Each thin line corresponds to the dominant PC of a single subject.

Denoting the reconstructed \mathbf{Y}_C^i as $\mathbf{V}_{\sim task}^i$, we take the mean of $\mathbf{V}_{\sim task}^i$ over brain areas to generate a representative RS activity time course for each subject i , \mathbf{y}_{RS}^i , which we enter into (3.1) as a confound regressor to re-estimate the task activation pattern. This process is repeated until the group activation pattern $\mathbf{\Lambda}_A$ stabilizes.

3.1.4 Experiments

Materials

For testing our proposed RS activity removal approach, we used the publicly available Multiband Test-Retest Pilot Dataset, which was released as a part of the 1000 Functional Connectomes Project.³ Excluding subjects with missing brain volumes, the dataset comprises 19 subjects (14 men, 5 women, mean age 33.1 ± 13.2 years). Each subject performed a passive viewing task in which a checkerboard was displayed on a monitor for 20 s, with 20 s of rest interleaved between stimulus blocks. The total task duration was approximately 2.5 minutes. Task fMRI data were acquired with a TR of 1.4 s and a voxel size of 2 mm (isotropic). RS fMRI data of 5 minutes duration were also collected with a TR of 2.5 s and a voxel size of 3mm (isotropic). Preprocessing of the data followed the pipeline presented in Section 2.2.3.

Results

To validate our proposed approach, we compared applying the standard univariate analysis [26] with and without RS activity removal in detecting group activation. We denote these two cases as RSR and GLM, respectively. For increased group activation detection to be a legitimate validation criterion, strong control over false positive rate is critical. For this, we used the max-t permutation test [67] for both RSR and GLM, which implicitly accounts for multiple comparisons, provides strong control on false positive rate, and generates less conservative t-value thresholds than Bonferroni correction [67].

Fig. 3.3 shows the number of detected parcels for different p-value thresholds. For the same specificity, our approach provided higher detection sensitivity than GLM in general. To assess whether the increased detection was statistically significant, we employed a “parcel-label” permutation test. Specifically, for each permutation, we first randomly selected half of the parcels and exchanged the labels (i.e. active or non-active) assigned by RSR and GLM for each p-value threshold. We then computed the difference in the number of detected parcels with and without RS removal, which we denote as N_{diff} . This procedure was repeated 1000 times to generate a null distribution. The original N_{diff} was found to be greater than the 95th percentile of the null distribution for all corrected p-value thresholds within the typical range of [0.01, 0.05]. The detection improvement with

³This dataset is available online at: http://fcon_1000.projects.nitrc.org/indi/pro/eNKI_RS_TRT/FrontPage.html

RSR compared to GLM was thus statistically significant. We note that improvement in detection was observed even with just one iteration of RSR, and the detected activation pattern stabilized within two iterations, i.e. no more than a couple of parcels changing labels in subsequent iterations.

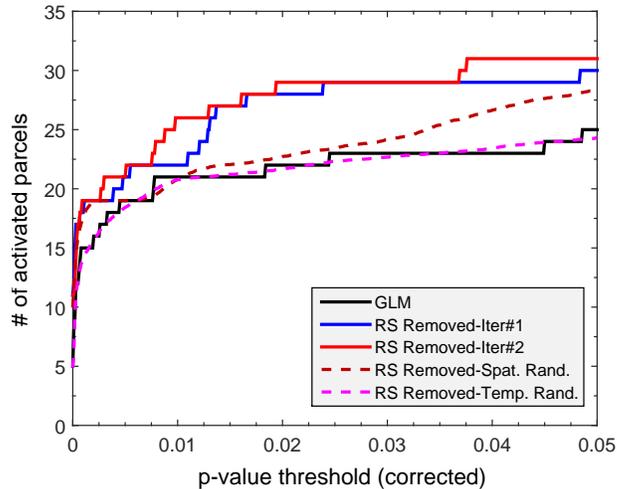


Figure 3.3: Comparison of activation detection results with and without the proposed RS activity removal method. Number of parcels detected with significant activation vs. p-value thresholds.

To show that there is specific temporal structure in the estimated RS activity time courses that gave rise to the observed detection improvement, we applied RSR on temporally-permuted RS activity time courses 100 times. The average number of detected parcels over the 100 permutations (RS Removed-Temp. Rand. in Fig. 3.3) was found to be similar to that of GLM. The difference in detection performance between RSR and temporally-permuted case was statistically significant based on the parcel-label permutation test with a threshold set at the 95th percentile of the null distribution. Our results thus indicate that there is certain temporal structure in the estimated RS activity time courses that is critical for successful RS activity removal.

Since the brain comprises multiple subnetworks [78], not all parcels would contain the same RS modulations as those superimposed on the underlying task activated brain areas. To illustrate this point, we applied RSR with RS modulations extracted from N_c randomly selected parcels (excluding parcels identified by RSR), where N_c is the number of intrinsically-connected

brain areas originally determined with RSR. The average number of detected parcels over 100 random subsets of parcels (RS Removed-Spat. Rand. in Fig. 3.3) was found to be similar to that of GLM for p-value thresholds between 0 and 0.02, and modestly better than GLM for p-value thresholds above 0.02. We suspect the increased detection arises from how some parcels might be intrinsically-connected to the task-evoked brain areas, but the estimated correlations were declared not significant due to noise. Such parcels would contain RS modulations common to the task activation pattern, hence the increased activation detection observed. Nevertheless, the increase was not statistically significant based on the parcel-label permutation test with a threshold set at the 95th percentile of the null distribution.

Qualitatively, RSR additionally detected brain areas adjacent to those found by GLM (Fig. 3.4). More bilateral activation was also found with RSR. The detected brain areas lie within the primary visual cortex and the extrastriate cortex, which are known to pertain to visual checkerboard stimulus [131], hence confirming our results.

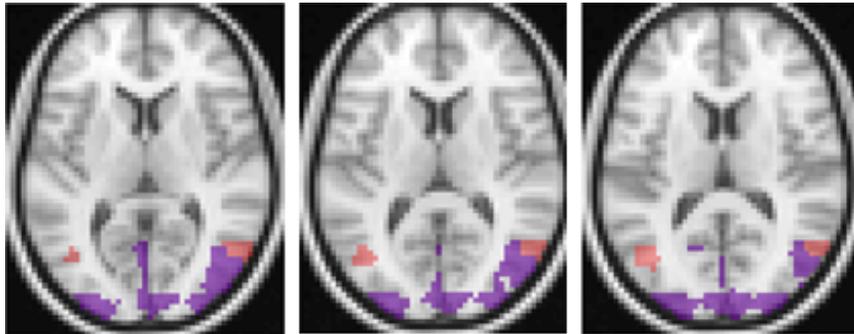


Figure 3.4: Detected activation patterns with and without the proposed RS activity removal method. Three axial slices shown. Parcels detected at p-value < 0.05 (corrected). Red = detected by RSR only. Purple = detected by both GLM and RSR.

3.1.5 Discussion

We proposed a novel approach for the estimation and removal of continual RS activity in task fMRI data. Exploiting how the spatial structure of RS subnetworks are constrained by the underlying fiber pathways hence would remain similar during task performance, our approach first extracts RS modulations from task fMRI time courses within brain areas that are

significantly correlated with an approximate task activation pattern at rest. The estimated RS modulations are then entered into a GLM as confound regressors to model the effects of RS activity. Applying our approach on real data resulted in statistically significant improvement in task activation detection. In agreement with the seminal work by Fox *et al.* [75], our results indicate that RS activity also contributes to the noise seen in task fMRI data, in contrast to the traditional belief that only scanner artifacts, head motions, and physiological confounds contribute to fMRI noise. It is thus important to model RS activity in task fMRI studies.

3.2 Modelling Anatomical Connectivity

In this section, we propose incorporating AC information into task activation detection. Given that fiber pathways serve as the physical substrate for functional interactions, we hypothesize that intrinsically connected brain areas would likely be in similar state, e.g. co-activate, during task [79]. Thus, informing activation detection with AC should presumably improve the detection sensitivity. For this, we employ the graph-theoretic random walker (RW) formulation [132], which easily permits such an integrated scheme for estimating activation probabilities. Posterior activation probabilities estimated by the RW formulation are guaranteed to be unique and globally-optimal [132], which makes RW an eminent choice. RW has been previously applied to task activation detection with functional connectivity taken as the prior [74]. Here, we investigate the implications of complementing task activation detection analysis with AC information. To infer group activation from posterior activation probabilities, we devise a permutation test with activation probabilities as attributes, which we empirically show to provide stronger control on false positive rate than simply comparing the posterior probability of being activated, not activated, and de-activated. After estimating a group activation pattern, we iteratively refine the anatomical connectivity prior by removing the links between non-active brain regions and re-estimate the posterior activation probabilities until the detected activation pattern stabilizes. An overview of our multimodal task activation detection approach is shown in Fig. 3.5.

On real data, we demonstrate that incorporating AC increases sensi-

Section 3.2 is adapted from: B. Yoldemir, B. Ng, T. Woodward, R. Abugharbieh. Fiber connectivity integrated brain activation detection. In Proceedings of the 23rd Biennial International Conference on Information Processing in Medical Imaging (IPMI'13), LNCS, vol. 7917, pp. 135-146, Asilomar, CA, June 2013.

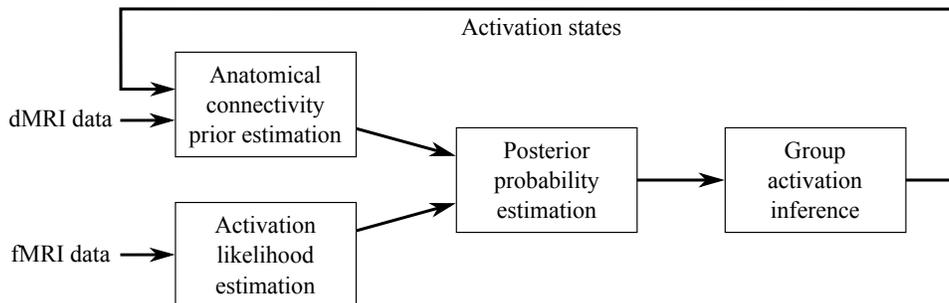


Figure 3.5: Overview of the proposed multimodal task activation detection approach. Anatomical connectivity prior and activation likelihood estimates, i.e. label priors, are integrated under the RW formulation to find the posterior activation probabilities. A permutation test is applied on these probabilities to infer group activation. The anatomical connectivity prior is then iteratively refined based on the activation states of the brain regions until convergence.

tivity in detecting group activation over using univariate techniques. We further show that our method is able to detect significant group activation differences between schizophrenia patients and healthy controls, which are missed by standard analysis approaches.

3.2.1 Random Walker for Activation Estimation

RW is a graph-based image segmentation approach, in which graph nodes (vertices) correspond to image voxels and graph edges connecting neighbouring nodes are assigned weights reflecting node similarity. In its original formulation [133], RW labels nodes based on the probability that a random walker starting from each node will first reach a pre-labeled seed given edge weights that bias the paths. This requires user-specified seeds and does not utilize local observations at each vertex. Thus, in the context of activation detection with brain regions being the graph nodes and their interactions modelled through graph edges, this would require user-defined seeds for every functionally-disparate region and only functional interactions would be considered without accounting for activation effects. We thus adopt the formulation in [132], which overcomes the need for user interaction and integrates activation effects into the formulation as label priors. This is equivalent to adding a floating vertex for each label and connecting these floating vertices to every vertex in the original graph with label priors being the

3.2. Modelling Anatomical Connectivity

weights of the added edges [132]. In this formulation, posterior probabilities are calculated by the minimization of the following energy functional:

$$E(\mathbf{x}^s) = \mathbf{x}^{sT} \mathbf{L} \mathbf{x}^s + \sum_{k=1, k \neq s}^K \mathbf{x}^{kT} \mathbf{\Lambda}^k \mathbf{x}^k + (\mathbf{x}^s - 1)^T \mathbf{\Lambda}^s (\mathbf{x}^s - 1), \quad (3.4)$$

where \mathbf{x}^s is an $M \times 1$ vector of unknown posterior probabilities of each ROI belonging to class s , \mathbf{L} is an $M \times M$ weighted Laplacian matrix, $\mathbf{\Lambda}^s$ is an $M \times M$ diagonal matrix having the prior probabilities of the ROIs belonging to class s on its diagonal, K is the number of class labels, and M is the number of brain regions. The first term in (3.4) is a spatial term for modelling the interactions between graph vertices as characterized by \mathbf{L} . The second term denotes the aspatial component for modelling the local observations at the vertices. In our context of activation detection, the spatial term models the anatomical connectivity information and the aspatial term models the activation effects. The method can thus be thought of as grouping brain regions into classes via random walk on an augmented graph, where edges in the original graph are weighted by AC information and edges leading to the floating nodes are weighted by activation effects. Assuming equal weighting between the spatial and aspatial energy terms in (3.4), it has been shown [132] that the posterior probabilities can be found by solving

$$\left(\mathbf{L} + \sum_{k=1}^K \mathbf{\Lambda}^k \right) \mathbf{x}^s = \boldsymbol{\lambda}^s, \quad (3.5)$$

where $\boldsymbol{\lambda}^s$ is an $M \times 1$ vector consisting of the diagonal elements of $\mathbf{\Lambda}^s$. Since \mathbf{L} is positive semi-definite and $\mathbf{\Lambda}^k$ is strictly positive definite, their summation would be diagonally dominant. Hence, matrix inversion is possible for solving (3.5). Following [134], we set $K = 3$ and define the class labels as deactive, nonactive, and active. For clarity, we explicitly denote the posterior probabilities for each class as \mathbf{p}_D , \mathbf{p}_N and \mathbf{p}_A , corresponding to deactive, nonactive and active classes, respectively.

3.2.2 Anatomical Connectivity Prior Estimation

Let \mathbf{D} be an $M \times M$ weighted adjacency matrix, where each element \mathbf{D}_{ij} is an estimate of the AC between brain regions i and j , set to fiber count following our discussion in Section 2.3. The corresponding $M \times M$ normalized

Laplacian matrix, \mathbf{L} , of \mathbf{D} is given by

$$\mathbf{L}_{ij} = \begin{cases} 1 & \text{if } i = j \text{ and } d_j \neq 0 \\ -\frac{1}{\sqrt{d_i d_j}} & \text{if } i \text{ and } j \text{ are adjacent ,} \\ 0 & \text{otherwise} \end{cases}, \quad (3.6)$$

where $d_i = \sum_j \mathbf{D}_{ij}$ is the degree of node i . A major difficulty with tractography is resolving fiber crossing regions where accuracy of most algorithms is seriously affected. In [64], it was proposed that multiplying \mathbf{D} by itself may help address this problem by generating multi-step fibers from parts of fibers that might be split at crossing regions:

$$\mathbf{D}^{MS} = \exp(\mathbf{D}) = \sum_{k=0}^{\infty} \frac{1}{k!} \mathbf{D}^k, \quad (3.7)$$

where $\exp(\cdot)$ denotes the matrix exponential. $\mathbf{D}^k = \mathbf{D} \cdot \mathbf{D} \cdot \mathbf{D} \cdot \dots$ and \mathbf{D}_{ij}^k is the number of paths of length k connecting regions i and j . \mathbf{D}^{MS} hence comprises all possible paths between each region pair, where indirect paths are more heavily penalized as these paths are potentially artifactual [64].

3.2.3 Activation Likelihood Estimation

To estimate the activation likelihoods, which are used as label priors $\boldsymbol{\lambda}^s$ in RW, we first apply the classical GLM [10] to compute the intra-subject activation statistics:

$$\begin{aligned} \mathbf{y}_j &= \mathbf{X}\boldsymbol{\beta}_j + \mathbf{e}_j \\ \mathbf{t}_j &= \hat{\boldsymbol{\beta}}_j / se(\hat{\boldsymbol{\beta}}_j), \end{aligned} \quad (3.8)$$

where \mathbf{y}_j is the $n \times 1$ time course of ROI j , \mathbf{X} is the $n \times r$ design matrix of expected responses, $\hat{\boldsymbol{\beta}}_j$ is an $r \times 1$ vector containing estimates of the activation effects, $\boldsymbol{\beta}_j$, \mathbf{e}_j is the $n \times 1$ residual assumed to be white Gaussian noise after preprocessing, $se(\hat{\boldsymbol{\beta}}_j)$ is the standard error of $\hat{\boldsymbol{\beta}}_j$, \mathbf{t}_j is the $r \times 1$ vector of sought activation statistics, n is the number of time points, and r is the number of experimental conditions. Columns of the design matrix \mathbf{X} are generated by convolving the canonical hemodynamic response function (HRF) with a boxcar time-locked to stimulus [10]. To compute the prior probabilities of ROIs belonging to each class, we fit a Gamma-Gaussian-Gamma (GGG) mixture model separately to the t-values of each condition, \mathbf{t}^c , i.e. c^{th} element of \mathbf{t}_j assembled across all ROIs [134]. The Gaussian

distribution models the nonactive state and the Gamma distributions model the deactive and active states:

$$\mathbf{t}_j^c \sim \pi_D \Gamma(k_D, \theta_D) + \pi_N \mathcal{N}(\mu, \sigma) + \pi_A \Gamma(k_A, \theta_A), \quad (3.9)$$

where μ is the mean and σ is the standard deviation of the Gaussian component, k_D and k_A are the shape parameters, and θ_D and θ_A are the scale parameters of the deactivation (D) and activation (A) components, respectively. We employ the expectation-maximization (EM) algorithm [135] to estimate the model parameters separately for each experimental condition c with the probabilities of \mathbf{t}^c given the parameter estimates used as the label priors λ^s . These label priors and the Laplacian matrix are combined through (3.5) to estimate the posterior activation probabilities.

3.2.4 Group Activation Inference

The high dimensionality of fMRI data elicits a high risk of false detection. To infer group activation from activation statistics, such as t-values, several methods that control for false positive rate have been proposed, e.g. Bonferroni correction, Gaussian random field theory, max-t permutation test [67]. For group activation inference from posterior activation probability maps, most studies directly threshold the posterior activation probabilities at $1/K$, where K is the number of classes, arguing that activation inference from posterior activation probabilities does not suffer from false positives [136]. However, as we will empirically demonstrate in the next subsection, directly thresholding the posterior probabilities is actually prone to false detection, necessitating a more rigorous activation inference method. To this end, we propose a permutation test for group inference from activation probabilities that controls for the false positive rate. Specifically, for each permutation, we first randomly select one third of the subjects and swap the posterior probabilities \mathbf{p}_A and \mathbf{p}_D of each selected subject. We note that this swap is done at the intra-subject level for all ROIs, hence the spatial pattern of the activation probabilities is preserved. Similarly, we swap \mathbf{p}_A and \mathbf{p}_N for another third of randomly selected subjects. We then compute the z-scores of \mathbf{p}_A for each permutation across all subjects. This procedure is repeated 10,000 times to generate the null distribution of activation probabilities of each ROI. We assign the p-value for the activation likelihood of each ROI as the number of times the z-scores of permuted \mathbf{p}_A are greater than the z-scores of the original \mathbf{p}_A values divided by the total number of trials (in this case, 10,000). Under the null hypothesis that there is no activation, the z-score of the original \mathbf{p}_A value of an ROI would lie around

the mean of the generated null distribution, resulting in a non-significant p-value around 0.5. Finally, FDR [137] is applied to these p-values to account for multiple comparisons. As shown in the next subsection on a synthetic case, this permutation test offers a much lower false positive rate compared to posterior probability thresholding while keeping the true positive rate at the same (or even at a higher) rate.

3.2.5 Re-estimation of Group Activation Using Partial Connectome

Since some of the estimated fiber tracts might be false due to tractography errors and not all fibers are necessarily employed during a given cognitive task, we propose restricting the anatomical connectivity prior to the subset of fibers that are likely to be active by iteratively refining the original connectivity prior, \mathbf{D}^{orig} , as

$$\mathbf{D}^{m+1} = \mathbf{D}^{task} + \frac{1}{\log(m+1)} \mathbf{D}^m \quad (3.10)$$

$$\mathbf{D}_{ij}^{task} = \begin{cases} \mathbf{D}_{ij}^{orig} & \text{if } \ell_i^m \vee \ell_j^m = 1 \\ 0 & \text{if } \ell_i^m = \ell_j^m = 0 \end{cases}, \quad (3.11)$$

where \mathbf{D}^m is the $M \times M$ partial adjacency matrix found in iteration m with \mathbf{D}^1 set to \mathbf{D}^{orig} . ℓ_i^m is the estimated activation state of ROI i during iteration m with a value of 1 denoting activated and 0 denoting not activated or deactivated. We note that updating \mathbf{D}^{orig} using the above scheme enables \mathbf{D}^{orig} to gradually evolve as opposed to having its information from previous iterations completely discarded. The overall process proceeds by estimating the group activation map with \mathbf{D}^1 as the prior. We then refine \mathbf{D}^1 using (3.10) and (3.11), and repeat the process until the group activation map stabilizes.

3.2.6 Experiments

Materials

After obtaining informed consent, fMRI data were collected from 13 healthy subjects (6 men, 7 women, mean age 27.46 ± 6.38 years) and 7 schizophrenia patients (5 men, 2 women, mean age 30.57 ± 10.08 years). Each subject was first presented with words in four different contexts: associating, hearing, solving and reading, during a non-scanned encoding session. The subjects

were then presented the same set of words in a subsequent recall run during which fMRI data were acquired and subjects were asked to indicate the context in which the presented words were previously encountered. Image acquisition was performed on a Philips Achieva 3.0 T MRI scanner using a T2*-weighted gradient-echo spin pulse sequence with a repetition time of 2000 ms, an echo time of 30 ms, a flip angle of 90° , a field of view of 240×240 mm, and an in-plane resolution of 80×80 pixels. Each volume comprised 36 axial slices of 3 mm thickness with a 1 mm gap. Each scan lasted for 920 s, which tallies to 460 fMRI volumes.

dMRI data were collected from the same subjects using a Philips Achieva 3.0 T MRI scanner with a TR of 7500 ms, a TE of 54 ms, an EPI factor of 59, an FOV of 224×224 mm and an in-plane resolution of 256×256 pixels. Fifteen diffusion weighted volumes were acquired at a b-value of 800 s/mm^2 in addition to a volume with no diffusion sensitization. Each volume consisted of 72 slices of 2 mm thickness with no gap. Acquisition time was 480 s. Preprocessing of the fMRI and dMRI data followed the pipeline presented in Section 2.2.3 and Section 2.3.2. We used MedINRIA [138] for fiber tractography. To facilitate the computation of fiber count, we warped our functionally derived group parcel map to the $b=0$ volume of each subject.

Results

We first present the synthetic test performed to assess our proposed permutation test as compared to the posterior probability thresholding approach commonly employed in the literature [136]. For evaluation on real data, we compare the sensitivity of our proposed approach in detecting group activation in controls against that of univariate techniques. We then contrast our method against classical schemes in detecting group activation differences between schizophrenia patients and controls.

We performed a synthetic test carefully designed from activation probabilities calculated from the real data of healthy controls. Specifically, after estimating the group activation map with our approach, we used the highest one third of \mathbf{p}_N among nonactive ROIs and the highest one third of \mathbf{p}_A among active ROIs to generate a pseudo ground truth of nonactivation and activation probability distributions at the intra-subject level. These pseudo ground truth distributions are assumed to be Gaussian with means and standard deviations set to that of the respective thirds of \mathbf{p}_N and \mathbf{p}_A . Out of a total of 100 synthetic datasets, each dataset comprised 13 subjects having 500 ROIs each, with 100 of them defined to be active. Random samples of \mathbf{p}_N and \mathbf{p}_A were drawn for each subject from the corresponding

3.2. Modelling Anatomical Connectivity

probability distributions. Deactivation probabilities were computed based on how posterior probabilities should sum to 1. Assessing group activation with the proposed permutation test resulted in a true positive rate (TP) of 0.819 ± 0.038 and a false positive rate (FP) of 0.008 ± 0.004 , whereas thresholding \mathbf{p}_A at $1/3$ gave a TP of 0.374 ± 0.049 and an FP of 0.097 ± 0.015 . One-sample t-tests among TPs and FPs of these two strategies declared the differences to be significant at $p < 10^{-6}$, demonstrating superior sensitivity and specificity.

We compare the sensitivity of our method against classical GLM and against inferring group activation from the activation likelihoods given by the GGG mixture model. We further compare the effect of using \mathbf{D} and \mathbf{D}^{MS} as the anatomical connectivity estimates. Fig. 3.6 shows the number of detected ROIs for different p-value thresholds. Our approach is denoted as RW for the first iteration where the full anatomical adjacency matrix is used, and RW_{it} for the following iterations with partial adjacency matrices. For the same specificity, our approach provided higher detection sensitivity than using the activation likelihoods given by the GGG mixture model, implicating the advantage of incorporating AC into activation detection. Iterating the procedure with refined AC estimates considerably improved detection. The procedure was iterated 100 times and for clarity, only the mean and standard deviations for the iterations with partial adjacency matrices were provided. It was observed that group activation map stabilized after 100 iterations, with only a couple of parcels changing labels in further iterations. This additional improvement provided by the refinement of the AC estimate suggests that only a subset of fibers is employed during a given task. Hence, isolating the utilized fibers can be beneficial. The improvement could also be partly due to removal of false fiber tracts arising from tractography errors. GLM (FDR corrected) provides more detection than our method at more liberal thresholds, but its performance varies considerably with p-value thresholds. In contrast, our method provides more consistent results across p-values. Comparing the results of using the original anatomical adjacency matrix versus the multi-step counterpart (denoted as RW-MS for the first, RW-MS_{it} for the following iterations) suggests that some of the artifactual connections generated by the multi-step approach likely do not pertain to task activation.

Qualitatively, the ROIs detected across all experimental conditions in healthy controls largely match regions known to be involved in context memory tasks [139], which further validates our method. As shown in Fig. 3.7, the areas detected across all conditions include superior frontal gyrus, middle frontal gyrus, inferior frontal gyrus, supramarginal gyrus, precentral gyrus,

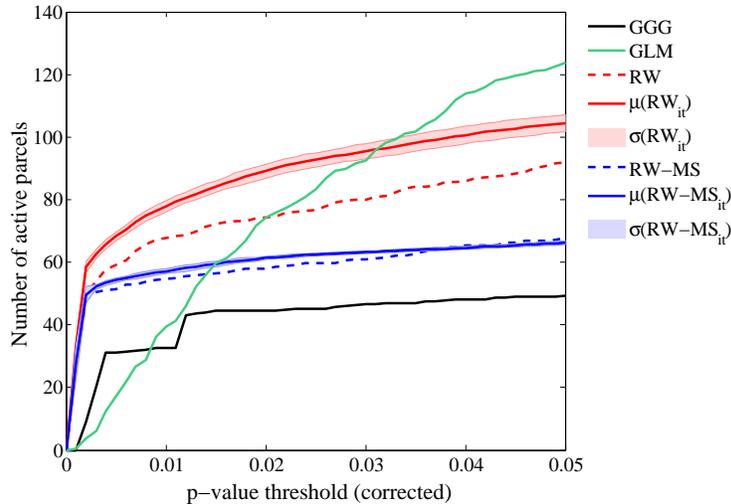


Figure 3.6: Quantitative results of the proposed multimodal activation detection technique.

postcentral gyrus, angular gyrus, lateral occipital cortex, occipital pole, cingulate gyrus, lingual gyrus, hippocampus and insula.

To assess the significant activation differences between schizophrenia patients and control subjects, we compared two different types of context memory recall: (1) self-other source information monitoring (did I say this word or did I hear it?) and (2) task information monitoring (did I produce a semantic associate of this word or did I read it?). Both types of context memory are thought to be impaired in schizophrenia. This comparison corresponds to the contrast between associating/reading and solving/hearing conditions. We employed a max-t permutation test [67] on the \mathbf{p}_A values of the two groups for this contrast and observed a significant difference ($p < 0.05$, corrected) in the left hippocampus. Activity in this region was higher in the source monitoring condition for schizophrenia patients relative to controls but lower in the task monitoring condition relative to controls. Applying the same test on the results of GLM or GGG mixture modelling approach failed to detect any group differences up to $p < 0.2$. The left hippocampus is involved in reactivation and association of stored semantic knowledge to consolidate new information into existing semantic frameworks [140–142]. The implication of the significant difference in the left hippocampus is that, aberrant activity in this region could lead to different manifestations of poor

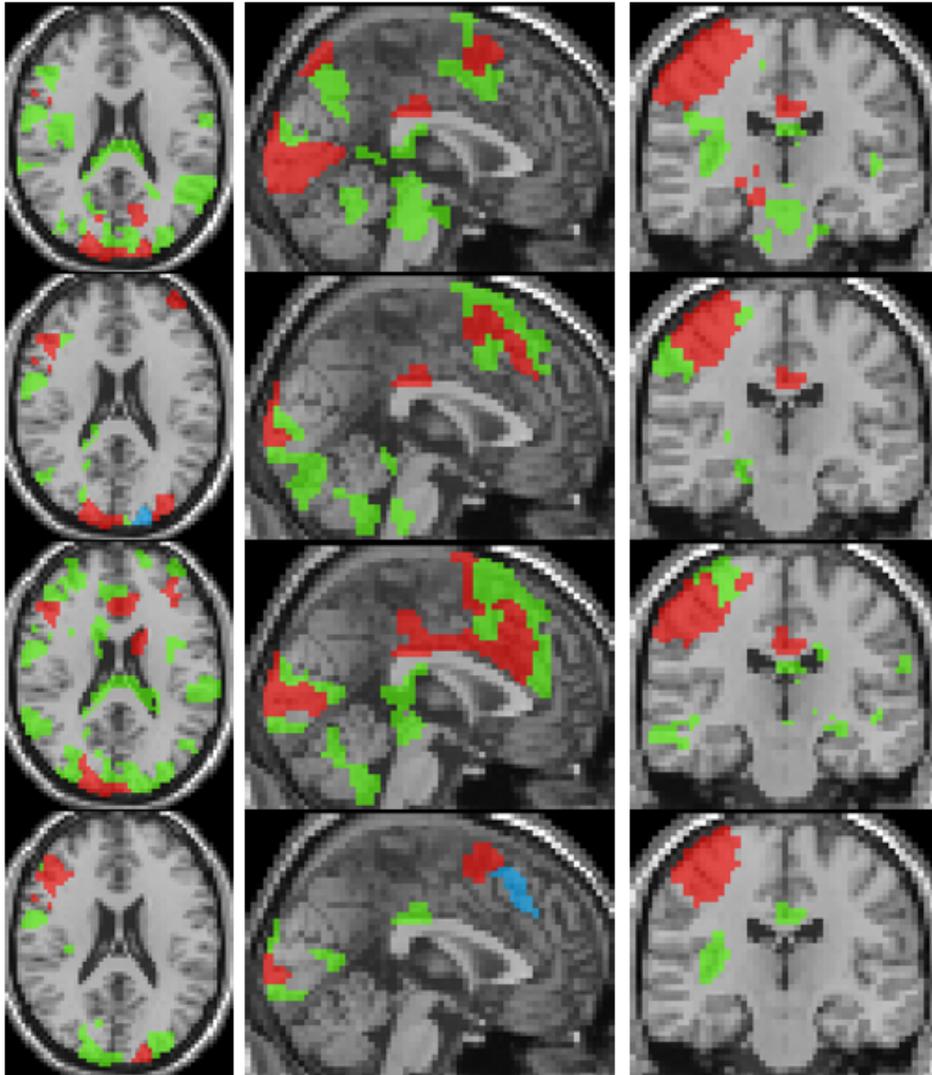


Figure 3.7: Qualitative results of the proposed multimodal activation detection method at p -value < 0.05 . Green=detected by RW only. Blue=detected by GGG only. Red=detected both by RW and GGG. Each row corresponds to an experimental condition, top-to-bottom: associating, hearing, solving, reading.

performance in context memory. Reduced activity during task monitoring could relate to the episodic memory impairments commonly observed in schizophrenia [143], which implies that underactivity in this region would lead to reduced reactivation of stored semantic knowledge for context memory. Increased activity during source monitoring could relate to the source memory impairments commonly observed in schizophrenia [144], such that overactivity in this region would lead to increased perceptual context for both self and other source information, leading to more difficulty in distinguishing between these two sources in context memory.

3.2.7 Discussion

In this section, we proposed a novel fiber connectivity integrated approach for group activation inference. On real data of healthy subjects, we demonstrated that integrating dMRI and fMRI significantly increases sensitivity in detecting group activation compared to analyzing fMRI data alone. We further showed that incorporating a refined connectome comprising anatomical connections linked only to the estimated active brain regions results in improved performance over using the full estimated connectome. Finally, we presented novel findings in activation differences between healthy controls and schizophrenia patients that were missed with standard methods. Our multimodal integration strategy thus holds great promise for brain activity analysis.

3.3 Summary

In this chapter, we proposed two multimodal integration methods to improve the sensitivity in activation detection from fMRI data. The first method involves the fusion of RS fMRI and task fMRI data, and aims at deconfounding the effects of RS modulations from task fMRI data. Our results imply that RS activity also contributes to the noise seen in task fMRI data, and that its effects can be diminished by considering RS fMRI and task fMRI jointly in a multimodal framework. Our second proposed method involves the fusion of dMRI and task fMRI data. We incorporated AC information into task activation detection motivated by the fact that fiber pathways serve as the physical substrate for functional interactions, and showed that this multimodal approach increases sensitivity in activation detection. Given the strong noise in fMRI data, we illustrated that multimodal methods provide promising solutions to regularize the challenging problem of brain activation detection.

Chapter 4

Multimodal Brain Subnetwork Identification

The brain is known to consist of spatially disparate regions that are functionally connected as subnetworks [145]. Unraveling this modular structure of the brain is a challenging endeavor. To identify subnetworks from fMRI data, seed-based correlation [22] and ICA [86] are primarily used. More recently, graph theoretical approaches [146] are gaining popularity as they facilitate compact brain representation and enable the reduction of complex brain interactions into intuitive summary network measures. Under this formalism, the brain is modelled as a graph, where brain regions and their connections correspond to nodes and edges, respectively.

As described in Section 1.6.2, AC and FC estimates suffer from opposing artifacts. Namely, FC estimates tend to contain many false positive connections due to several confounding factors involved, whereas AC estimates often suffer from false negatives due to premature termination of reconstructed tracts at fiber crossings. We thus hypothesize that multimodal integration for subnetwork identification should prove beneficial since incorporating AC would reduce the effects of noise-induced FC and vice versa.

In this chapter, we propose a technique based on replicator dynamics (RD) that facilitates integration of fMRI and dMRI data in identifying brain subnetworks. We start by reviewing the properties of RD [147], which only

Chapter 4 is adapted from the following papers:

B. Yoldemir, B. Ng, R. Abugharbieh. Overlapping replicator dynamics for functional subnetwork identification. In Proceedings of the International Conference on Medical Image Computing and Computer Assisted Intervention (MICCAI'13), LNCS, vol. 8150, pp. 682-689, Nagoya, Japan, September 2013.

B. Yoldemir, B. Ng, R. Abugharbieh. Coupled stable overlapping replicator dynamics for multimodal brain subnetwork identification. In Proceedings of the 24th Biennial International Conference on Information Processing in Medical Imaging (IPMI'15), LNCS, vol. 9123, pp. 770-781, Isle of Skye, Scotland, June 2015.

B. Yoldemir, B. Ng, R. Abugharbieh. Stable overlapping replicator dynamics for brain community detection. IEEE Transactions on Medical Imaging, vol. 35, no. 2, pp. 529-538, February 2016.

permits unimodal subnetwork identification (Section 4.1). We then describe a sex-differentiated formulation of RD [148] (Section 4.2) that permits FC and AC to be jointly modelled. In particular, this formulation extracts coupled subnetwork pairs based on fitness measures of both females and males, which in the present context, correspond to FC and AC estimates. We refer to this formulation as coupled RD (CRD). Importantly, we prove that for dense graphs, CRD always selects the same set of nodes for each pair of identified subnetworks (Section 4.2). Hence, the effective output of CRD is actually a single set of subnetworks with FC and AC integrated. CRD alone is prone to noise and tends to falsely divide the subnetworks. We thus adopt a graph incrementation scheme to facilitate merging of subnetwork components that might be falsely split and present a procedure for incorporating stability selection [149] to statistically control for inclusion of false nodes arising from noise and overmerging (Section 4.3). We refer to this extension of CRD as coupled stable RD (CSRD). Further, we describe how a graph augmentation strategy can be integrated into CSRD to enable overlapping subnetwork identification (Section 4.4). In aggregate, the resulting technique: (i) facilitates multimodal integration, (ii) can operate on weighted graphs, (iii) allows for subnetwork overlaps, (iv) has an intrinsic criterion for determining the number of subnetworks, (v) does not force all nodes to be part of a subnetwork, and (vi) statistically controls for false node inclusion. We refer to this technique as coupled stable overlapping RD (CSORD).

4.1 Replicator Dynamics

RD is a concept that originated from theoretical biology for modelling the evolution of interacting and self-replicating entities [147]. Let $\mathbf{w}(k) = (\mathbf{w}_1(k), \dots, \mathbf{w}_d(k))^T$ be a $d \times 1$ vector, where $\mathbf{w}_j(k)$ is the proportion of allele that are of type j in the gene pool during generation k , and d is the number of alleles. Further, let \mathbf{C} be a $d \times d$ matrix with each element, \mathbf{C}_{ij} , reflecting the fitness of a genotype, i.e. a pair of alleles i and j . Under the assumption of natural selection, $\mathbf{w}(k+1)$ is given by the replicator equation [147]:

$$\mathbf{w}(k+1) = \frac{\mathbf{w}(k) \circ \mathbf{C}\mathbf{w}(k)}{\mathbf{w}(k)^T \mathbf{C}\mathbf{w}(k)}, \quad (4.1)$$

where \circ denotes element-wise multiplication. Eq. (4.1) is shown to solve the following non-convex optimization problem [129]:

$$\max_{\mathbf{w}} \mathbf{w}^T \mathbf{C}\mathbf{w} \text{ s.t. } \|\mathbf{w}\|_1 = 1, \mathbf{w} \geq 0. \quad (4.2)$$

The fundamental theorem of natural selection states that if \mathbf{C} is real-valued, symmetric, and non-negative, $\mathbf{w}(k)^T \mathbf{C} \mathbf{w}(k)$ strictly increases with k along any nonstationary trajectory $\mathbf{w}(k)$ until it converges to a strict local maximum as (4.1) is iterated [150]. In the more general case of any real-valued \mathbf{C} , local maxima of $\mathbf{w}(k)^T \mathbf{C} \mathbf{w}(k)$ are guaranteed to attract nearby trajectories $\mathbf{w}(k)$ asymptotically as (4.1) is iterated [151]. $\mathbf{w}_j(k+1)$ will equal $\mathbf{w}_j(k)$ upon convergence thus $(\mathbf{C} \mathbf{w}(k))_j = \mathbf{w}(k)^T \mathbf{C} \mathbf{w}(k)$ for all j , where $(\cdot)_j$ denotes the j^{th} element of the corresponding vector. In graph-theoretic terms with \mathbf{C} being a similarity matrix, this implies that all selected nodes (i.e. nodes with non-zero $\mathbf{w}(k)$) will have the same weighted average correlation with one another. Hence, the grouping of the nodes is based on mutual similarity among the nodes. With \mathbf{w} constrained to lie on the standard simplex, i.e. $\|\mathbf{w}\|_1 = 1$, $\mathbf{w} \geq 0$, sparse \mathbf{w} is encouraged since restricting the l_1 norm induces sparsity [112]. $\mathbf{w}_j(k) > 0$ hence indicates that allele j persisted in generation k . The local maximum to which RD converges highly depends on $\mathbf{w}(0)$. Since the vicinity at which the local maxima reside is unknown a priori, an unbiased way of setting $\mathbf{w}(0)$ is to assign all of its elements to $1/d$, which tends to find the community with the highest mutual similarity among nodes [129]. Other subnetworks can be found by reapplying (4.1) after removing the nodes in the identified subnetworks [152].

4.2 Coupled Replicator Dynamics

Separating \mathbf{C} into $d \times d$ genotype fitness matrices \mathbf{F} for females and \mathbf{M} for males to model sex differences, the survival probabilities of the alleles in the female gene pool, $\mathbf{p}(k)$, and in male gene pool, $\mathbf{q}(k)$, based on natural selection can be estimated using the following multiplicative updates [148]:

$$\mathbf{p}(k+1) = 0.5 \frac{\mathbf{p}(k) \circ \mathbf{F} \mathbf{q}(k) + \mathbf{q}(k) \circ \mathbf{F} \mathbf{p}(k)}{\mathbf{p}(k)^T \mathbf{F} \mathbf{q}(k)} \quad (4.3)$$

$$\mathbf{q}(k+1) = 0.5 \frac{\mathbf{q}(k) \circ \mathbf{M} \mathbf{p}(k) + \mathbf{p}(k) \circ \mathbf{M} \mathbf{q}(k)}{\mathbf{p}(k)^T \mathbf{M} \mathbf{q}(k)}. \quad (4.4)$$

Analogous to RD, \mathbf{F} and \mathbf{M} are assumed to be real, symmetric, and non-negative. The non-negativity constraint enforces $\mathbf{p}(k)$ and $\mathbf{q}(k)$ to be

4.2. Coupled Replicator Dynamics

greater than or equal to 0. Also, $\sum_j \mathbf{p}_j(k) = 1$ since:

$$\begin{aligned}
 & 0.5 \cdot \sum_j (\mathbf{p}_j(k)(\mathbf{F}\mathbf{q}(k))_j + \mathbf{q}_j(k)(\mathbf{F}\mathbf{p}(k))_j) = \\
 & 0.5 \cdot (\mathbf{p}(k)^T \mathbf{F}\mathbf{q}(k) + \mathbf{q}(k)^T \mathbf{F}\mathbf{p}(k)) = \\
 & 0.5 \cdot (\mathbf{p}(k)^T \mathbf{F}\mathbf{q}(k) + \mathbf{p}(k)^T \mathbf{F}\mathbf{q}(k)) = \\
 & \mathbf{p}(k)^T \mathbf{F}\mathbf{q}(k),
 \end{aligned} \tag{4.5}$$

where we have exerted the property that a scalar and its transpose are equal, i.e. $\mathbf{q}(k)^T \mathbf{F}\mathbf{p}(k) = \mathbf{p}(k)^T \mathbf{F}\mathbf{q}(k)$. Similarly, $\sum_j \mathbf{q}_j(k) = 1$. Thus, by matching the denominator of (4.3) and (4.4) with that of (4.1), one can see that (4.3) and (4.4) together form a local maximizer of the following optimization problem:

$$\max_{\mathbf{p}, \mathbf{q}} \mathbf{p}^T \mathbf{F}\mathbf{q} + \mathbf{p}^T \mathbf{M}\mathbf{q} \text{ s.t. } \|\mathbf{p}\|_1 = 1, \mathbf{p} \geq 0, \|\mathbf{q}\|_1 = 1, \mathbf{q} \geq 0, \tag{4.6}$$

where the l_1 constraint on \mathbf{p} and \mathbf{q} enforces sparsity. Since \mathbf{p} is multiplied to both \mathbf{F} and \mathbf{M} in (4.6), the optimal \mathbf{p} would jointly account for the genotype fitness of both sexes, and the same goes for \mathbf{q} . This coupling property of CRD can be more clearly seen by examining (4.3) and (4.4). Specifically, information in \mathbf{M} is propagated to \mathbf{p} through \mathbf{q} , and vice versa. Further, we highlight that the sparsity patterns of \mathbf{p} and \mathbf{q} tend to perfectly match upon convergence for dense graphs. Specifically, assume $\mathbf{p}_i(k) = 0$ upon convergence. Based on (4.3), this implies $\mathbf{p}_i(k) \circ \mathbf{F}\mathbf{q}_i(k) + \mathbf{q}_i(k) \circ \mathbf{F}\mathbf{p}_i(k) = 0$. $\mathbf{p}_i(k) \circ \mathbf{F}\mathbf{q}_i(k)$ must equal 0 since $\mathbf{p}_i(k) = 0$, thus $\mathbf{q}_i(k) \circ \mathbf{F}\mathbf{p}_i(k)$ must also equal 0. This condition is satisfied only if $\mathbf{q}_i(k) = 0$ or $\mathbf{F}\mathbf{p}_i(k) = 0$. $\mathbf{F}\mathbf{p}_i(k) = 0$ requires \mathbf{F}_{ij} to be exactly zero for each surviving allele j . In both our synthetic and real data, \mathbf{F} is never sparse enough to warrant this condition. Hence, $\mathbf{q}_i(k) \circ \mathbf{F}\mathbf{p}_i(k) = 0$ can only arise by having $\mathbf{q}_i(k) = 0$, which enforces \mathbf{p} and \mathbf{q} to have the same sparsity patterns. The nodes in the identified subnetwork are thus given by the nonzero elements of either \mathbf{p} or \mathbf{q} . Again, other subnetworks can be found by reapplying (4.3) and (4.4) after removing the nodes of the identified subnetworks from the graphs. However, this strategy modifies the original graph structure, thereby affecting the remaining local maxima. We describe a more principled way of identifying the other subnetworks that permits subnetwork overlaps and preserves the location of the remaining local maxima in Section 4.4.

4.3 Coupled Stable Replicator Dynamics

4.3.1 Graph Incrementation

Sparsity of the subnetworks identified by CRD is intrinsically determined based on the mutual similarity among nodes. This property is desirable under the ideal case of noiseless graphs. In practice, such a strict criterion for grouping nodes could easily result in a subnetwork being falsely split into components [153], since even small perturbations to \mathbf{F} and \mathbf{M} would render certain nodes non-mutually connected. To reduce this effect, one strategy is to add a constant, η , to the off-diagonal elements of \mathbf{F} and \mathbf{M} . Using an extreme example for intuition, if we add 10,000 to \mathbf{F} and \mathbf{M} with values ranging from 0 to 1, the original differences between the elements of \mathbf{F} (and \mathbf{M}) would be negligible. Applying CRD to such matrices would result in all nodes being assigned to a single subnetwork, i.e. all elements of \mathbf{p} and \mathbf{q} would attain nonzero values. If we instead use a η of magnitude similar to the values in \mathbf{F} and \mathbf{M} , only the smaller differences in \mathbf{F} and \mathbf{M} would be negligible. This strategy thus artificially increases the mutual similarity among all nodes and allows adjusting the sparsity level of the identified subnetworks by changing η . Choosing the optimal η is non-trivial. Also, noise in \mathbf{F} and \mathbf{M} could result in false nodes being included in the subnetworks. To jointly deal with these problems, we present next a procedure for integrating stability selection into CRD.

4.3.2 Stability Selection

The idea behind stability selection is that if we bootstrap the data many times and perform subnetwork identification on each bootstrap sample, nodes that truly belong to the same subnetwork are likely to be jointly selected over a large fraction of bootstrap samples, whereas false nodes are unlikely to be persistently selected [149]. To incorporate stability selection in refining each subnetwork identified by CRD, we first generate a set of matrices $\{\mathbf{F}_\eta\}$ by adding η to the off-diagonal elements of \mathbf{F} . A range of η from 0 to $d \cdot \max_{i,j} |\mathbf{F}_{ij}|$ at a step size of κ is used, where κ is chosen small enough to ensure that no more than one node is added to a subnetwork at each increment of η . The same procedure is applied to \mathbf{M} to generate $\{\mathbf{M}_\eta\}$. We then apply CRD to each $(\mathbf{F}_\eta, \mathbf{M}_\eta)$ pair with $\mathbf{p}(0)$ and $\mathbf{q}(0)$ set to the CRD solution of the previous η increment (with $1/d$ added to all elements and renormalized to sum to 1), which enforces CRD to converge to a local maximum in the vicinity of the previous η increment for retaining subnetwork correspondence. We note that adding $1/d$ is critical since zeros in $\mathbf{p}(k)$

and $\mathbf{q}(k)$ remain as zeros as (4.3) and (4.4) are iterated, which prohibits new nodes from being added. We discard subnetworks that contain more than 10% of the nodes assuming a subnetwork would not span more than 10% of the brain. Next, we generate 100 bootstrap samples of \mathbf{F} and \mathbf{M} (as explained in Section 4.5), and compute $\{\mathbf{F}_\eta^b\}$ and $\{\mathbf{M}_\eta^b\}$ from each bootstrap sample \mathbf{F}^b and \mathbf{M}^b . We then apply CRD on each $(\mathbf{F}_\eta^b, \mathbf{M}_\eta^b)$ pair. For each η , we estimate the selection probability of each node as the proportion of bootstrap samples in which the given node attains a nonzero weight [149]. This estimation can be done based on either \mathbf{p} or \mathbf{q} since their sparsity patterns perfectly match. Finally, we threshold the selection probabilities to identify nodes that belong to the same subnetwork. A threshold τ that bounds the expected number of false nodes in a subnetwork, E , is given by [149]

$$\tau \leq \left(\frac{q^2}{Ed} + 1 \right) / 2, \quad (4.7)$$

where q is the average number of nodes per subnetwork, which can be estimated as the sum of selection probabilities of all nodes averaged over η . We set E to 1 so that the average number of false nodes included in a subnetwork is statistically controlled to be less than or equal to 1. We declare the set of nodes with selection probabilities higher than the resulting τ for any η to be a subnetwork [149]. We note that each subnetwork identified by CRD is refined independently using the above procedure.

4.4 Coupled Overlapping Stable Replicator Dynamics

Since CRD is deterministic, reapplying it with the same initialization will converge to the same local maximum. Here, we adopt a graph augmentation strategy that enables destabilization of previously-found local maxima without altering the location of the others. Reapplying CRD on this augmented graph would thus converge to one of the remaining local maxima.

4.5. Parameter Selection and Implementation

The required graph augmentation is as follows [151]:

$$\mathbf{F}_{ij}^{aug} = \begin{cases} \mathbf{F}_{ij} & \text{if } i, j \leq d \\ \alpha & \text{if } j > d \text{ and } i \notin S_{j-d} \\ \beta & \text{if } i, j > d \text{ and } i = j \\ \gamma_{ij} & \text{if } i > d \text{ and } j \in S_{i-d}, \\ 0 & \text{otherwise} \end{cases}, \quad (4.8)$$

$$\gamma_{ij} = \sigma_{ij} \frac{1}{|S_{i-d}|} \sum_{m \in S_{i-d}} \mathbf{F}_{mj}, \quad \sigma_{ij} > 1$$

where $\alpha > \beta$, $\beta = \max_{i \neq j} \mathbf{F}_{ij}$, S_l is the set of nodes of the l^{th} identified sub-network, and $|\cdot|$ denotes cardinality. \mathbf{M} is similarly augmented to generate \mathbf{M}^{aug} . In effect, this is equivalent to adding artificial nodes to the graphs by appending \mathbf{F} and \mathbf{M} with new rows and new columns, and extending weighted edges to the original nodes such that the previously-found local maxima provably become unstable. Given S_1 found using CRD and refined by stability selection, finding S_2 proceeds by adding a new row and a new column to \mathbf{F} and \mathbf{M} based on (4.8), applying CRD on the new $(\mathbf{F}^{aug}, \mathbf{M}^{aug})$ pair, and refining the resulting subnetwork using stability selection. Remaining S_l can be similarly extracted. We highlight that an intrinsic criterion for terminating further subnetwork extraction is to stop if after convergence, $\mathbf{p}(k)^T \mathbf{F} \mathbf{q}(k) + \mathbf{p}(k)^T \mathbf{M} \mathbf{q}(k) \leq \mathbf{p}(0)^T \mathbf{F} \mathbf{q}(0) + \mathbf{p}(0)^T \mathbf{M} \mathbf{q}(0)$, since this suggests that no further solutions of (4.6) are present. In practice, we found on synthetic data that using $a \cdot (\mathbf{p}(0)^T \mathbf{F} \mathbf{q}(0) + \mathbf{p}(0)^T \mathbf{M} \mathbf{q}(0))$, $a > 1$, as the stopping criterion is more robust to noisy \mathbf{F} and \mathbf{M} . A rigorous way of choosing a is discussed in Section 4.5.

4.5 Parameter Selection and Implementation

In the context of brain subnetwork identification, \mathbf{F} and \mathbf{M} correspond to estimates of FC and AC, respectively, and alleles correspond to brain ROIs. Our goal is to find subnetworks of ROIs that are highly inter-connected both functionally and anatomically. For \mathbf{F} , we use the conventional Pearson's correlation between average ROI time series. To ensure \mathbf{F} is non-negative as required for the properties of CRD to hold, we set negative \mathbf{F}_{ij} to zero given the currently unclear interpretation of negative FC. For \mathbf{M} , we use the fiber count between ROIs following our discussion in Section 2.3.2. The main diagonals of \mathbf{F} and \mathbf{M} are set to zero to avoid self-connections [129]. Given

the substantially larger magnitudes in \mathbf{M} compared to \mathbf{F} , the subnetwork detection process would be dominated by \mathbf{M} . To mitigate this effect, we linearly rescale \mathbf{M} such that the values in \mathbf{F} and \mathbf{M} are in the same range.

An open problem in subnetwork identification is the choice on the number of subnetworks. Although $\mathbf{p}(k)^T \mathbf{F} \mathbf{q}(k) + \mathbf{p}(k)^T \mathbf{M} \mathbf{q}(k) \leq \mathbf{p}(0)^T \mathbf{F} \mathbf{q}(0) + \mathbf{p}(0)^T \mathbf{M} \mathbf{q}(0)$ is theoretically justified, noisy \mathbf{F} and \mathbf{M} could lead to spurious subnetworks with $\mathbf{p}(k)^T \mathbf{F} \mathbf{q}(k) + \mathbf{p}(k)^T \mathbf{M} \mathbf{q}(k)$ being slightly above the threshold. Hence, there is a need for a stricter threshold, $a \cdot (\mathbf{p}(0)^T \mathbf{F} \mathbf{q}(0) + \mathbf{p}(0)^T \mathbf{M} \mathbf{q}(0))$, $a \geq 1$. To set a , we generate 500 synthetic datasets and extract subnetworks over a range of a from 1 to 10. The value of a that minimizes the difference between the number of estimated subnetworks and the ground truth is declared as the optimal a . We note that the synthetic datasets used for selecting a are separate from those used for evaluating the performance of CSORD.

To generate bootstrap samples of \mathbf{F} for applying stability selection, we adopt a parametric bootstrap approach [154] since the spatiotemporal correlation structure of fMRI data would not be retained with regular bootstrap, where time points are independently sampled with replacement. Let \mathbf{X} be a $t \times d$ fMRI time series matrix, where t is the number of time samples. We randomly draw 100 \mathbf{X}^b of size $t \times d$ from a multivariate normal distribution with zero mean and covariance, \mathbf{F}_{pd} , where \mathbf{F}_{pd} is a positive definite approximation of \mathbf{F} , computed by setting non-positive eigenvalues of \mathbf{F} to 10^{-16} . \mathbf{F}^b is then computed as the Pearson’s correlation of \mathbf{X}^b . As for \mathbf{M} , one approach for generating bootstrap samples is to randomly select diffusion gradient directions with replacement to generate bootstrapped dMRI volumes and reapply tractography [155], but this is very computationally intensive. Instead, we use the same parametric bootstrap procedure to generate bootstrap samples of \mathbf{M} .

4.6 Materials

4.6.1 Synthetic Data

We generated 500 synthetic datasets that cover a wide variety of network configurations. Each dataset comprised $d = 200$ regions and 4 scans of 1200 time points as in the real data. We set the number of subnetworks, N , in each dataset to a random value between 10 and 20. The number of ROIs in each subnetwork was set to $\lceil d/N \rceil + c$ with c being a random number between 1 and 5, which ensures that the total number of ROIs across subnetworks is greater than d , hence guaranteeing the presence of subnetwork

overlaps. ROIs were randomly assigned to subnetworks, and the same ROI was allowed to repeat across subnetworks. We opted to use random network configurations since the ground truth number of subnetworks, number of nodes within each subnetwork, and amount of subnetwork overlap are unknown. Our strategy is thus to try to exhaust the possibilities. For each network configuration, we generated a $d \times d$ adjacency matrix, Σ , which was taken as the ground truth. Next, we built a noise-free AC matrix, Σ_a , by randomly setting $p_1\%$ of the values in Σ to 0 to model how AC estimates are prone to false negatives, and built a noise-free FC matrix, Σ_f , by randomly setting $p_2\%$ of the values in Σ to 1 to model how FC estimates are prone to false positives. p_1 and p_2 were randomly chosen from $[0, 20]$. Two sets of time series were then generated by drawing random samples from $\mathcal{N}(0, \Sigma_a)$ and $\mathcal{N}(0, \Sigma_f)$. Finally, AC and FC matrices were simulated by adding Gaussian noise to the time series with SNR randomly set between -6 and -3 dB, and computing the Pearson’s correlation of these noisy time series. We chose the typical SNR levels seen in task fMRI data [23, 24] since SNR of RS fMRI data is hard to determine.

4.6.2 Real Data

dmMRI, RS fMRI and task fMRI data from the HCP dataset were used in the experiments [123]. Details and preprocessing of the data and the generation of the brain parcellation map are presented in Sections 2.2.3 and 2.3.2.

4.7 Results

4.7.1 Synthetic Data Results

We compared CSORD against CSRD, CRD, stable RD (SRD), RD, NC [156], and multi-view spectral clustering (MVSC) [52]. SRD is analogous to CSRD, except (4.1) is used instead of (4.3) and (4.4) in finding subnetworks. Among the contrasted techniques, CSORD, CSRD, CRD and MVSC permit multimodal subnetwork identification, whereas SRD, RD, and NC are limited to a single modality. For these unimodal techniques, we focused on their performance in extracting subnetworks from FC estimates. On 500 synthetic datasets, we assessed each contrasted technique by computing the Dice coefficient (DC) between the estimated and ground truth subnetworks as

$$\text{DC} = 2 \frac{|S_{est} \cap S_{gnd}|}{|S_{est}| + |S_{gnd}|}, \quad (4.9)$$

where S_{est} is the set of ROIs in an estimated subnetwork, and S_{gnd} is the set of ROIs in the corresponding ground truth subnetwork matched to S_{est} using Hungarian clustering [157]. DC of unmatched subnetworks was set to zero. CSORD achieved significantly higher DC than each contrasted technique at $p < 10^{-40}$ based on Wilcoxon signed rank test (Fig. 4.1).

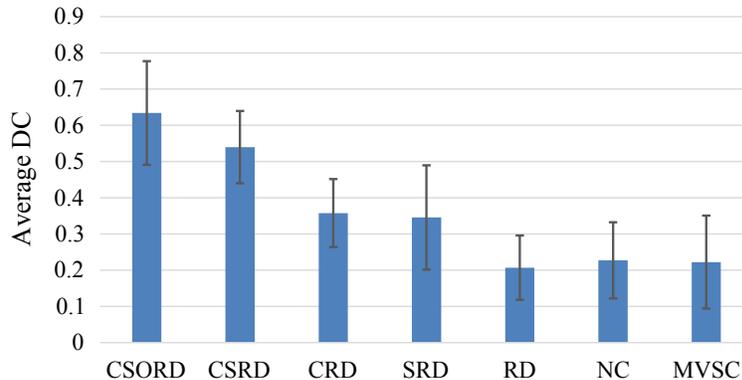


Figure 4.1: Subnetwork identification accuracy on synthetic data.

4.7.2 Real Data Results

Since ground truth subnetworks are unknown for real data, we based our validation on test-retest reliability and task classification accuracy. We estimated test-retest reliability as the intra-class correlation (ICC) over four graph metrics. Specifically, we first computed the clustering coefficient, transitivity, global efficiency, and local efficiency [158] for each subnetwork extracted from RS fMRI data (and dMRI data for multimodal methods) of each scan session of a given subject. We then averaged the values of a given metric across each subject’s subnetworks for each of the two scan sessions. Collecting these average values of each metric across subjects, we computed the ICC between the two scan sessions. We finally estimated test-retest reliability as the average of resulting four ICC values.

High test-retest reliability is a desired property, but could arise from common noise between scan sessions. We thus additionally assessed the task classification accuracy of the contrasted techniques as follows. We first extracted group subnetworks from RS fMRI and dMRI data. We then computed \mathbf{F} of each task from each subject’s task fMRI dataset, and scaled

its elements as

$$\mathbf{F}_{ij}^s = \mathbf{F}_{ij} \cdot \frac{|m_i \cap m_j|}{|m_i \cup m_j|}, \quad (4.10)$$

where m_i and m_j are the group subnetwork memberships of nodes i and j based on RS fMRI and dMRI data. This scaling weighs down \mathbf{F}_{ij} if nodes i and j were assigned to different group subnetworks, which in effect, propagates the group subnetwork structure onto \mathbf{F} . The weighted degree of each brain region computed from \mathbf{F}^s was taken as a feature, resulting in a $d \times 1$ feature vector for each task. A support vector machine was used for classification. For estimating classification accuracy, we subsampled the data over 100 random splits with 35 subjects used for training and 5 for testing.

We adopted the above scaling scheme of \mathbf{F} , which was originally used for extending the definition of modularity to overlapping subnetworks [159], since the notion of within and between subnetwork connections become fuzzy when subnetworks are allowed to overlap. Also, this evaluation scheme ensures the data used for subnetwork identification (RS fMRI and dMRI data) are independent from the data used for classification performance assessment (task fMRI data). The assumption is that tasks induce only small functional changes to intrinsic brain connectivity. This hypothesis is supported by the strong resemblance between RS and task subnetworks [78], as well as how energy consumption in the brain is mainly for supporting ongoing activity with task-evoked response constituting less than 5% of the total [76].

Results on test-retest reliability and task classification accuracy are shown in Fig. 4.2. Comparing RD’s and CRD’s performance suggests that simply incorporating AC without considering the limitations of RD produces only marginal benefits. In fact, a much larger gain was obtained by controlling for false node inclusions with SRD. Nevertheless, by jointly incorporating AC and controlling for false node inclusion as well as enabling subnetwork overlaps, CSORD achieved the highest task classification accuracy and is on par with the best contrasted techniques in terms of test-retest reliability. We note that average FC has been argued to contain no dynamical information. It is true that average connectivity matrix does not capture dynamical connectivity, but the overlapping structure of subnetworks is well reflected by the relative magnitude of its elements, and can be readily extracted as shown by our results.

Qualitatively, CSORD identified all commonly found subnetworks in the literature [160]. The extracted visuospatial subnetwork and right executive control subnetwork are shown in Figs. 4.3 and 4.4 as exemplar re-

4.7. Results

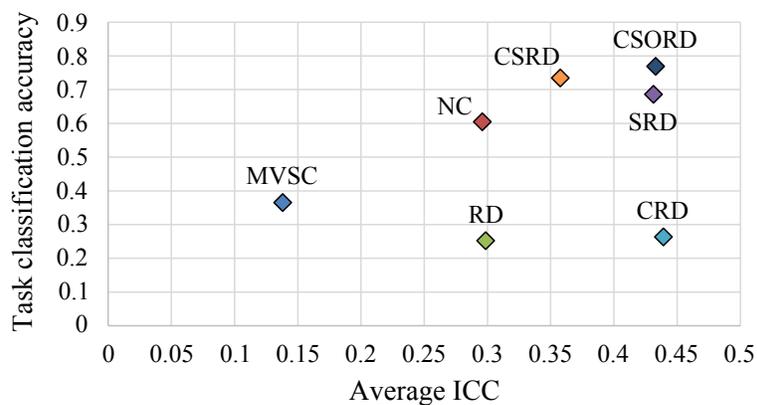


Figure 4.2: Subnetwork identification results on real data. CSORD achieved the best classification accuracy and high test-retest reliability.

sults. CSORD also identified a subnetwork (Fig. 4.5) that comprises the basal ganglia, thalamus, higher visual cortex, motor areas, and part of the cerebellum. These regions constitute the visual corticostriatal loop [161], striato-thalamo-cortical (STC) loop, and cerebello-thalamo-cortical (CTC) loop [162], which are rarely identified as a single subnetwork. Instead, state-of-the-art techniques typically declare the basal ganglia and cerebellum as separate subnetworks [78, 163], failing to capture cortical projections from subcortical regions.

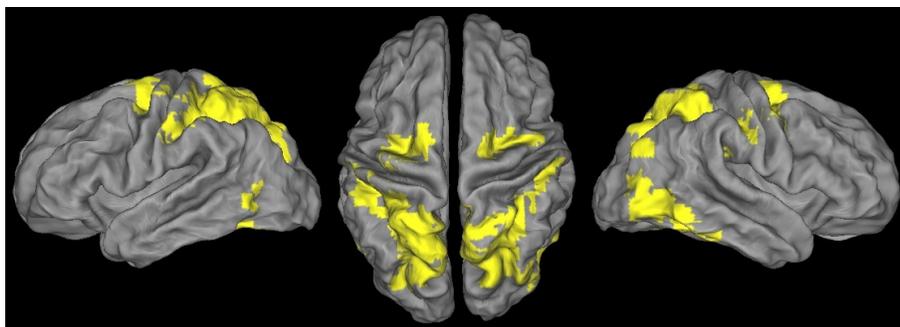


Figure 4.3: Visuospatial subnetwork identified by CSORD at the group level.

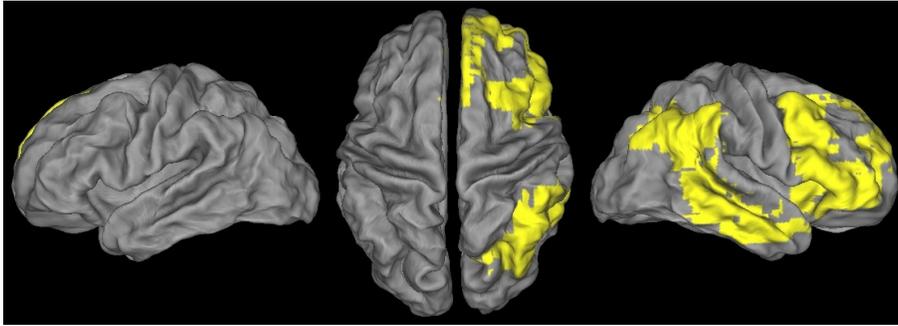


Figure 4.4: Right executive control subnetwork identified by CSORD at the group level.

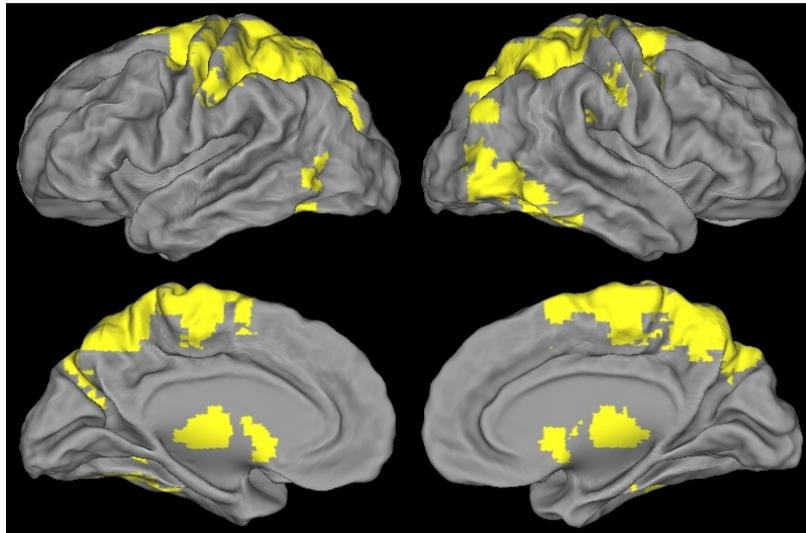


Figure 4.5: A group subnetwork identified by CSORD that comprises the visual corticostriatal, CTC, and STC loops. Note that the cerebellum is outside the presented views.

4.8 Discussion

We proposed CSORD to combine information from fMRI and dMRI data in identifying overlapping brain subnetworks. On synthetic data covering a diverse set of network configurations, we showed that CSORD provides significantly higher subnetwork identification accuracy than a number of state-of-the-art techniques. On real data, we illustrated that CSORD is not only more robust to inter-session variability than most of the contrasted techniques, but it also achieved the highest classification accuracy. Importantly, standard techniques typically declare the basal ganglia, cerebellum, and motor cortex as separate subnetworks, despite their known joint involvements in various motor loops. In contrast, CSORD was able to extract a single subnetwork comprising these well-known interacting motor regions. Collectively, our results show that multimodal integration is a promising approach for brain subnetwork identification.

Chapter 5

Conclusions

We presented in this thesis our technical contributions for the assessment of AC, and joint analysis of AC and FC for investigating functional segregation and integration in the human brain. Our techniques were based on multimodal integration of RS fMRI, task fMRI and dMRI data. We conclude in this section with a discussion of the strengths and weaknesses of our methods and some promising directions to explore.

In Chapter 2, we proposed the use of a super-resolution technique in dMRI to ameliorate the adverse effects of low spatial resolution on AC estimation. We showed that our approach qualitatively improves the fiber tracts and tract-density maps generated and significantly improves the correlation between AC and FC estimates on real data. We further showed that the traditional way of estimating AC from dMRI data, i.e. streamline tractography on diffusion tensors, leads to understating the AC-FC consistency, and using ODFs over diffusion tensors and global tractography over streamline tractography help increase the consistency between AC and FC estimates. Finally, we illustrated the effect of the choice of AC metric on the AC-FC consistency and showed that fiber count provides a good compromise between high AC-FC consistency and low fiber length bias.

In Chapter 3, we addressed the challenging problem of brain activation detection and proposed two multimodal strategies to increase sensitivity in detecting task activation. Our first method involved the fusion of RS fMRI and task fMRI data with the aim of deconfounding the effects of RS activity from task fMRI data. We showed that the removal of estimated RS modulations from task fMRI data significantly improves activation detection. We used a linear regression model for the removal of spontaneous fluctuations motivated by the work of Fox *et al.* [75], where the results imply that task-evoked signal and RS activity are roughly linearly superimposed in the brain. Our second proposed method involved the fusion of dMRI and task fMRI data. We showed that our fiber connectivity integrated activation detection approach increases sensitivity in detecting activation in controls, and also enables the detection of significant group differences between controls and schizophrenia patients. To the best of our knowledge, this is the first work

in the literature to regularize activation detection using AC information.

In Chapter 4, we proposed CSORD to identify groups of brain regions that exhibit high inter-connectivity both functionally and anatomically. Our proposed method draws on advances from theoretical biology [147], evolutionary game theory [151] and statistics [149]. There are several advantages of the proposed method: (i) it facilitates multimodal integration, (ii) it can operate on weighted graphs, (iii) it allows for subnetwork overlaps, (iv) it has an intrinsic criterion for determining the number of subnetworks, (v) it does not force all nodes to be part of a subnetwork, and (vi) it statistically controls for false node inclusion. We demonstrated how this method enables more accurate detection of subnetworks on synthetic data compared to a number of state-of-the-art approaches. Also, we showed that CSORD identifies more reproducible subnetworks and facilitates successful classification of the tasks subjects are involved in by analyzing the subnetworks recruited.

5.1 Future Work

Throughout Chapter 2, we used Pearson’s correlation to assess the consistency between AC and FC. This linear model was used instead of higher-order nonlinear models in estimating AC-FC consistency for easier interpretability, and based on the empirical finding that AC and FC in the human brain are roughly linearly related [64, 97] when AC and FC are quantified using the metrics we have adopted in this thesis. Nonetheless, understanding the exact nature of this relationship warrants further research. It should be emphasized that there is currently no well-accepted definition of “anatomical connectivity strength”. All AC metrics we have considered in this thesis constitute indirect measures based on tractography that capture different aspects of the connectivity pattern. Morphological measures such as cortical thickness computed from structural MRI data have also been employed for this purpose [164], however these are also proxies for AC strength. Also, cortical thickness cannot be computed for subcortical areas unlike tractography-based metrics. Devising more direct AC metrics is currently an open research problem beyond the scope of this thesis.

In Chapter 3, we used a linear model to deconfound RS modulations from task fMRI data. Our work was motivated by prior work which implies a largely linear superposition of task-evoked signal and RS modulations in the brain [75]. Our results in Section 3.1.4 support this linearity claim. However, we believe that understanding whether there are additional nonlinear interactions between the task-evoked signal and RS activity warrants further

5.1. *Future Work*

research. Sensitivity in brain activation detection can then presumably be further increased by modelling such additional interactions.

A limitation of CSORD presented in Chapter 4 is its high computational complexity. The method involves running an iterative scheme for each bootstrap sample, thus may not easily extend to voxel-level analysis. Nevertheless, the refinement of each subnetwork using stability selection is independently performed, hence the method may benefit from parallel programming to speed up the computation. Identifying subject-specific subnetworks using CSORD and examining whether the topology of these subnetworks can be linked to e.g. disease severity or genetic factors is a promising direction to explore in the future.

Bibliography

- [1] J. Hooper and D. Teresi. *The three-pound universe*. Macmillan, 1986.
- [2] Mapping connections: An understanding of neurological conditions in Canada. Public Health Agency of Canada, 2014.
- [3] The global burden of disease: Generating evidence, guiding policy. Institute for Health Metrics and Evaluation, University of Washington, 2012.
- [4] V. Hesselmann, B. Sorger, R. Girnus, K. Lasek, M. Maarouf, C. Wedekind, J. Bunke, O. Schulte, B. Krug, K. Lackner, and V. Sturm. Intraoperative functional MRI as a new approach to monitor deep brain stimulation in Parkinson’s disease. *European Radiology*, 14(4):686–690, 2004.
- [5] G.D. Honey, E. Pomarol-Clotet, P.R. Corlett, R.A.E. Honey, P.J. Mckenna, E.T. Bullmore, and P.C. Fletcher. Functional dysconnectivity in schizophrenia associated with attentional modulation of motor function. *Brain*, 128(11):2597–2611, 2005.
- [6] T.C. Baghai, H.-J. Moller, and R. Rupprecht. Recent progress in pharmacological and non-pharmacological treatment options of major depression. *Current Pharmaceutical Design*, 12(4):503–515, 2006.
- [7] S.Y. Bookheimer, M.H. Strojwas, M.S. Cohen, A.M. Saunders, M.A. Pericak-Vance, J.C. Mazziotta, and G.W. Small. Patterns of brain activation in people at risk for Alzheimer’s disease. *New England Journal of Medicine*, 343(7):450–456, 2000.
- [8] M.D. Greicius, K. Supekar, V. Menon, and R.F. Dougherty. Resting-state functional connectivity reflects structural connectivity in the default mode network. *Cerebral Cortex*, 19(1):72–78, 2009.
- [9] S. Jbabdi and H. Johansen-Berg. Tractography: Where do we go from here? *Brain Connectivity*, 1(3):169–183, 2011.

- [10] K.J. Friston. Models of brain function in neuroimaging. *Annual Review of Psychology*, 56(1):57–87, 2005.
- [11] G. de Marco. Effective connectivity and brain modeling by fMRI. *Advanced Studies in Biology*, 1(3):139–144, 2009.
- [12] O. Sporns, G. Tononi, and R. Kötter. The human connectome: a structural description of the human brain. *PLoS Computational Biology*, 1(4):e42, 2005.
- [13] R.C. Craddock, S. Jbabdi, C.-G. Yan, J.T. Vogelstein, F.X. Castellanos, A. Di Martino, C. Kelly, K. Heberlein, S. Colcombe, and M.P. Milham. Imaging human connectomes at the macroscale. *Nature Methods*, 10(6):524–539, 2013.
- [14] P. J. Basser, J. Mattiello, and D. Le Bihan. Estimation of the effective self-diffusion tensor from the NMR spin echo. *Journal of Magnetic Resonance, Series B*, 103(3):247–254.
- [15] H.-E. Assemlal, D. Tschumperlé, L. Brun, and K. Siddiqi. Recent advances in diffusion MRI modeling: Angular and radial reconstruction. *Medical Image Analysis*, 15(4):369–396, 2011.
- [16] P. Hagmann, L. Jonasson, P. Maeder, J.-P. Thiran, V.J. Wedeen, and R. Meuli. Understanding diffusion MR imaging techniques: From scalar diffusion-weighted imaging to diffusion tensor imaging and beyond. *Radiographics*, 26:S205–S223, 2006.
- [17] D. Le Bihan, C. Poupon, A. Amadon, and F. Lethimonnier. Artifacts and pitfalls in diffusion MRI. *Journal of Magnetic Resonance Imaging*, 24(3):478–488, 2006.
- [18] J. Soares, P. Marques, V. Alves, and N. Sousa. A hitchhiker’s guide to diffusion tensor imaging. *Frontiers in Neuroscience*, 7(31), 2013.
- [19] S. Ogawa, T. M. Lee, A. R. Kay, and D. W. Tank. Brain magnetic resonance imaging with contrast dependent on blood oxygenation. *Proceedings of the National Academy of Sciences*, 87(24):9868–9872, 1990.
- [20] E. Amaro and G.J. Barker. Study design in fMRI: Basic principles. *Brain and Cognition*, 60(3):220–232, 2006.
- [21] D. Le Bihan, P. Jezzard, J. Haxby, N. Sadato, L. Rueckert, and V. Mattay. Functional magnetic resonance imaging of the brain. *Annals of Internal Medicine*, 122(4):296–303, 1995.

- [22] B. Biswal, F. Zerrin Yetkin, V.M. Haughton, and J.S. Hyde. Functional connectivity in the motor cortex of resting human brain using echo-planar MRI. *Magnetic Resonance in Medicine*, 34(4):537–541, 1995.
- [23] G.E. Sarty. *Computing brain activity maps from fMRI time-series images*. Cambridge University Press, 2007.
- [24] V. Calhoun. Preprocessing iii-iv: Anatomic transformation to standard brain spaces. http://ece.unm.edu/~vcalhoun/courses/fMRI_Spring10/. Accessed: 2015-08-18.
- [25] R.N.A. Henson, C. Buechel, O. Josephs, and K.J. Friston. The slice-timing problem in event-related fMRI. *NeuroImage*, 9:125, 1999.
- [26] K.J. Friston, J. Ashburner, C.D. Frith, J.-B. Poline, J.D. Heather, and R.S.J. Frackowiak. Spatial registration and normalization of images. *Human Brain Mapping*, 3(3):165–189, 1995.
- [27] A.P. Holmes, O. Josephs, C. Buchel, and K.J. Friston. Statistical modelling of low-frequency confounds in fMRI. *NeuroImage*, 5:S480, 1997.
- [28] K.J. Worsley, C.H. Liao, J. Aston, V. Petre, G.H. Duncan, F. Morales, and A.C. Evans. A general statistical analysis for fMRI data. *NeuroImage*, 15(1):1–15, 2002.
- [29] J.C. Mazziotta, A.W. Toga, A. Evans, P. Fox, and J. Lancaster. A probabilistic atlas of the human brain: Theory and rationale for its development: The international consortium for brain mapping (ICBM). *NeuroImage*, 2(2):89–101, 1995.
- [30] W.D. Penny, K.J. Friston, J.T. Ashburner, S.J. Kiebel, and T.E. Nichols. *Statistical parametric mapping: the analysis of functional brain images*. Academic Press, 2011.
- [31] S.M. Smith, M. Jenkinson, M.W. Woolrich, C.F. Beckmann, T.E.J. Behrens, H. Johansen-Berg, P.R. Bannister, M. De Luca, I. Drobnjak, and D.E. Flitney. Advances in functional and structural MR image analysis and implementation as FSL. *NeuroImage*, 23:S208–S219, 2004.

- [32] Y. Chao-Gan and Z. Yu-Feng. DPARSF: a MATLAB toolbox for pipeline data analysis of resting-state fMRI. *Frontiers in Systems Neuroscience*, 4, 2010.
- [33] P.J. Basser and C. Pierpaoli. Microstructural and physiological features of tissues elucidated by quantitative-diffusion-tensor MRI. *Journal of Magnetic Resonance*, 111(3):209–219, 1996.
- [34] P.J. Basser, J. Mattiello, and D. LeBihan. MR diffusion tensor spectroscopy and imaging. *Biophysical Journal*, 66(1):259, 1994.
- [35] P.J. Basser, S. Pajevic, C. Pierpaoli, J. Duda, and A. Aldroubi. In vivo fiber tractography using DT-MRI data. *Magnetic Resonance in Medicine*, 44(4):625–632, 2000.
- [36] T.E. Conturo, N.F. Lori, T.S. Cull, E. Akbudak, A.Z. Snyder, J.S. Shimony, R.C. McKinstry, H. Burton, and M.E. Raichle. Tracking neuronal fiber pathways in the living human brain. *Proceedings of the National Academy of Sciences*, 96(18):10422–10427, 1999.
- [37] D.K. Jones, A. Simmons, S.C.R. Williams, and M.A. Horsfield. Non-invasive assessment of axonal fiber connectivity in the human brain via diffusion tensor MRI. *Magnetic Resonance in Medicine*, 42(1):37–41, 1999.
- [38] S. Mori and J.-D. Tournier. *Introduction to diffusion tensor imaging and higher order models*. Academic Press, 2013.
- [39] P.T. Callaghan. NMR imaging, NMR diffraction and applications of pulsed gradient spin echoes in porous media. *Magnetic Resonance Imaging*, 14(78):701–709, 1996.
- [40] D.S. Tuch. Q-ball imaging. *Magnetic Resonance in Medicine*, 52(6):1358–1372, 2004.
- [41] I. Aganj, C. Lenglet, G. Sapiro, E. Yacoub, K. Ugurbil, and N. Harel. 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, 2010.
- [42] M. Reisert, I. Mader, C. Anastasopoulos, M. Weigel, S. Schnell, and V. Kiselev. Global fiber reconstruction becomes practical. *NeuroImage*, 54(2):955–962, 2011.

- [43] T.E.J. Behrens, H.J. Berg, S. Jbabdi, M.F.S. Rushworth, and M.W. Woolrich. Probabilistic diffusion tractography with multiple fibre orientations: What can we gain? *NeuroImage*, 34(1):144–155, 2007.
- [44] G.J.M. Parker, H.A. Haroon, and C.A.M. Wheeler-Kingshott. A framework for a streamline-based probabilistic index of connectivity (PICO) using a structural interpretation of MRI diffusion measurements. *Journal of Magnetic Resonance Imaging*, 18(2):242–254, 2003.
- [45] D.K. Jones and C. Pierpaoli. Confidence mapping in diffusion tensor magnetic resonance imaging tractography using a bootstrap approach. *Magnetic Resonance in Medicine*, 53(5):1143–1149, 2005.
- [46] M. Xu, X. Zhang, Y. Wang, L. Ren, Z. Wen, Y. Xu, G. Gong, N. Xu, and H. Yang. Probabilistic brain fiber tractography on GPUs. In *IEEE 26th International Parallel and Distributed Processing Symposium Workshops*, pages 742–751, 2012.
- [47] B. Jeurissen, A. Leemans, J.-D. Tournier, K. Jones, and J. Sijbers. Investigating the prevalence of complex fiber configurations in white matter tissue with diffusion magnetic resonance imaging. *Human Brain Mapping*, 34(11):2747–2766, 2013.
- [48] Marco Catani. From hodology to function. *Brain*, 130(3):602–605, 2007.
- [49] Q. Cao, N. Shu, L. An, P. Wang, L. Sun, M.-R. Xia, J.-H. Wang, G.-L. Gong, Y.-F. Zang, and Y.-F. Wang. Probabilistic diffusion tractography and graph theory analysis reveal abnormal white matter structural connectivity networks in drug-naive boys with attention deficit/hyperactivity disorder. *The Journal of Neuroscience*, 33(26):10676–10687, 2013.
- [50] P. Hagmann, L. Cammoun, X. Gigandet, R. Meuli, C.J. Honey, V.J. Wedeen, and O. Sporns. Mapping the structural core of human cerebral cortex. *PLoS Biology*, 6(7):e159, 07 2008.
- [51] V. Tomassini, S. Jbabdi, J.C. Klein, T.E.J. Behrens, C. Pozzilli, P.M. Matthews, M.F.S. Rushworth, and H. Johansen-Berg. Diffusion-weighted imaging tractography-based parcellation of the human lateral premotor cortex identifies dorsal and ventral subregions with anatomical and functional specializations. *The Journal of Neuroscience*, 27(38):10259–10269, 2007.

- [52] H. Chen, K. Li, D. Zhu, X. Jiang, Y. Yuan, P. Lv, T. Zhang, L. Guo, D. Shen, and T. Liu. Inferring group-wise consistent multimodal brain networks via multi-view spectral clustering. *IEEE Transactions on Medical Imaging*, 32(9):1576–1586, 2013.
- [53] D.K. Jones, T.R. Knösche, and R. Turner. White matter integrity, fiber count, and other fallacies: the do’s and don’ts of diffusion MRI. *NeuroImage*, 73:239–254, 2013.
- [54] K.J. Friston, C.D. Frith, P.F. Liddle, and R.S.J. Frackowiak. Functional connectivity: the principal-component analysis of large PET data sets. *Journal of Cerebral Blood Flow and Metabolism*, 13:5–14, 1993.
- [55] M.D. Fox and M. Greicius. Clinical applications of resting state functional connectivity. *Frontiers in Systems Neuroscience*, 4, 2010.
- [56] G. Marrelec, A. Krainik, H. Duffau, M. Plgrini-Issac, S. Lehicry, J. Doyon, and H. Benali. Partial correlation for functional brain interactivity investigation in functional MRI. *NeuroImage*, 32(1):228–237, 2006.
- [57] F.T. Sun, L.M. Miller, and M. D’Esposito. Measuring interregional functional connectivity using coherence and partial coherence analyses of fMRI data. *NeuroImage*, 21(2):647–658, 2004.
- [58] S.M. Smith, K.L. Miller, G. Salimi-Khorshidi, M. Webster, C.F. Beckmann, T.E. Nichols, J.D. Ramsey, and M.W. Woolrich. Network modelling methods for fMRI. *NeuroImage*, 54(2):875–891, 2011.
- [59] M. Fiecas, H. Ombao, D. van Lunen, R. Baumgartner, A. Coimbra, and D. Feng. Quantifying temporal correlations: A test-retest evaluation of functional connectivity in resting-state fMRI. *NeuroImage*, 65:231–241, 2013.
- [60] BrainVoyager. Functional analysis: Preparation - preprocessing. <http://support.brainvoyager.com/functional-analysis-preparation.html>. Accessed: 2015-09-28.
- [61] M.A. Koch, D.G. Norris, and M. Hund-Georgiadis. An investigation of functional and anatomical connectivity using magnetic resonance imaging. *NeuroImage*, 16(1):241–250, 2002.

- [62] M. van den Heuvel, R. Mandl, J. Luigjes, and H. Hulshoff Pol. Microstructural organization of the cingulum tract and the level of default mode functional connectivity. *The Journal of Neuroscience*, 28(43):10844–10851, 2008.
- [63] C.J. Honey, O. Sporns, L. Cammoun, X. Gigandet, J.P. Thiran, R. Meuli, and P. Hagmann. Predicting human resting-state functional connectivity from structural connectivity. *Proceedings of the National Academy of Sciences*, 106(6):2035–2040, 2009.
- [64] P. Skudlarski, K. Jagannathan, V.D. Calhoun, M. Hampson, B.A. Skudlarska, and G. Pearlson. Measuring brain connectivity: Diffusion tensor imaging validates resting state temporal correlations. *NeuroImage*, 43(3):554–561, 2008.
- [65] B. Ng, G. Varoquaux, J.B. Poline, and B. Thirion. Implications of inconsistencies between fMRI and dMRI on multimodal connectivity estimation. In *Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, pages 652–659. 2013.
- [66] K.J. Friston, A.P. Holmes, K.J. Worsley, J.-P. Poline, C.D. Frith, and R.S.J. Frackowiak. Statistical parametric maps in functional imaging: A general linear approach. *Human Brain Mapping*, 2(4):189–210, 1994.
- [67] T. Nichols and S. Hayasaka. Controlling the familywise error rate in functional neuroimaging: a comparative review. *Statistical Methods in Medical Research*, 12(5):419–446, 2003.
- [68] G.K. Aguirre, E. Zarahn, and M. D’Esposito. Empirical analyses of BOLD fMRI statistics. *NeuroImage*, 5(3):199–212, 1997.
- [69] W.D. Penny, N.J. Trujillo-Barreto, and K.J. Friston. Bayesian fMRI time series analysis with spatial priors. *NeuroImage*, 24(2):350–362, 2005.
- [70] L.M. Harrison, W. Penny, J. Ashburner, N. Trujillo-Barreto, and K.J. Friston. Diffusion-based spatial priors for imaging. *NeuroImage*, 38(4):677–695, 2007.
- [71] C. Gossel, D.P. Auer, and L. Fahrmeir. Bayesian spatio-temporal inference in functional magnetic resonance imaging. *Biometrics*, 57:554–562, 2001.

- [72] X. Descombes, F. Kruggel, and D.Y. von Cramon. Spatio-temporal fMRI analysis using Markov random fields. *IEEE Transactions on Medical Imaging*, 17(6):1028–1039, 1998.
- [73] W. Ou and P. Golland. From spatial regularization to anatomical priors in fMRI analysis. In *Information Processing in Medical Imaging (IPMI)*, pages 88–100. 2005.
- [74] B. Ng, R. Abugharbieh, G. Hamarneh, and M. J. McKeown. Random walker based estimation and spatial analysis of probabilistic fMRI activation maps. In *MICCAI fMRI Data Analysis Workshop*, pages 37–44. 2009.
- [75] M.D. Fox, A.Z. Snyder, J.L. Vincent, and M.E. Raichle. Intrinsic fluctuations within cortical systems account for intertrial variability in human behavior. *Neuron*, 56(1):171–184, 2007.
- [76] M.D. Fox and M.E. Raichle. Spontaneous fluctuations in brain activity observed with functional magnetic resonance imaging. *Nature Reviews Neuroscience*, 8:700–711, 2007.
- [77] G. Deco, V.K. Jirsa, and A.R. McIntosh. Emerging concepts for the dynamical organization of resting-state activity in the brain. *Nature Reviews Neuroscience*, 12:43–56, 2011.
- [78] S.M. Smith, P.T. Fox, K.L. Miller, D.C. Glahn, P.M. Fox, C.E. Mackay, N. Filippini, K.E. Watkins, R. Toro, A.R. Laird, and C.F. Beckmann. Correspondence of the brain’s functional architecture during activation and rest. *Proceedings of the National Academy of Sciences*, 106(31):13040–13045, 2009.
- [79] B. Ng, R. Abugharbieh, G. Varoquaux, J.-B. Poline, and B. Thirion. Connectivity-informed fMRI activation detection. In *Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, pages 285–292. 2011.
- [80] K.J. Friston. Functional and effective connectivity: A review. *Brain Connectivity*, 1(1):13–36, 2011.
- [81] K. Li, L. Guo, J. Nie, G. Li, and T. Liu. Review of methods for functional brain connectivity detection using fMRI. *Computerized Medical Imaging and Graphics*, 33(2):131–139, 2009.

- [82] C. Goutte, P. Toft, E. Rostrup, F.A. Nielsen, and L.K. Hansen. On clustering fMRI time series. *NeuroImage*, 9(3):298–310, 1999.
- [83] D. Cordes, V. Haughton, J.D. Carew, K. Arfanakis, and K. Maravilla. Hierarchical clustering to measure connectivity in fMRI resting-state data. *Magnetic Resonance Imaging*, 20(4):305–317, 2002.
- [84] P. Golland, Y. Golland, and R. Malach. Detection of spatial activation patterns as unsupervised segmentation of fMRI data. In *Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, pages 110–118, 2007.
- [85] L. Ma, B. Wang, X. Chen, and J. Xiong. Detecting functional connectivity in the resting brain: a comparison between ICA and CCA. *Magnetic Resonance Imaging*, 25(1):47–56, 2007.
- [86] M.J. McKeown, S. Makeig, G.G. Brown, T.P. Jung, S.S. Kindermann, A.J. Bell, and T.J. Sejnowski. Analysis of fMRI data by blind separation into independent spatial components. *Human Brain Mapping*, 6:160–188, 1998.
- [87] V. D. Calhoun, T. Adali, G. D. Pearlson, and J. J. Pekar. A method for making group inferences from functional MRI data using independent component analysis. *Human Brain Mapping*, 14(3):140–151, 2001.
- [88] S.M. Smith, K.L. Miller, S. Moeller, J. Xu, E.J. Auerbach, M.W. Woolrich, C.F. Beckmann, M. Jenkinson, J. Andersson, M.F. Glasser, D.C. Van Essen, D.A. Feinberg, E.S. Yacoub, and K. Ugurbil. Temporally-independent functional modes of spontaneous brain activity. *Proceedings of the National Academy of Sciences*, 109(8):3131–3136, 2012.
- [89] K. Wu, Y. Taki, K. Sato, Y. Sassa, K. Inoue, R. Goto, K. Okada, R. Kawashima, Y. He, A.C. Evans, and H. Fukuda. The overlapping community structure of structural brain network in young healthy individuals. *PLoS ONE*, 6(5):e19608, 2011.
- [90] V.D. Blondel, J.-L. Guillaume, R. Lambiotte, and E. Lefebvre. Fast unfolding of communities in large networks. *Journal of Statistical Mechanics: Theory and Experiment*, 2008(10):P10008, 2008.
- [91] X. Yan, S. Kelley, M. Goldberg, and B.B. Biswal. Detecting overlapped functional clusters in resting state fMRI with connected iterative scan:

- A graph theory based clustering algorithm. *Journal of Neuroscience Methods*, 199(1):108–118, 2011.
- [92] T. Madhyastha, Y. Cao, S. Sujitnapisatham, G. Presnyakov, W.A. Chaovaitwongse, and T. Grabowski. Link clustering to explore brain dynamics using resting state functional MRI. *Journal of Radiology and Radiation Therapy*, 1(2):1012, 2013.
- [93] Y.-Y. Ahn, J.P. Bagrow, and S. Lehmann. Link communities reveal multiscale complexity in networks. *Nature*, 466(7307):761–764, 2010.
- [94] J. Xie, S. Kelley, and B.K. Szymanski. Overlapping community detection in networks: The state-of-the-art and comparative study. *ACM Computing Surveys*, 45(4):43, 2013.
- [95] R. Guimera, M. Sales-Pardo, and L.A.N. Amaral. Modularity from fluctuations in random graphs and complex networks. *Physical Review E*, 70(2):025101, 2004.
- [96] A. Lancichinetti, F. Radicchi, J.J. Ramasco, and S. Fortunato. Finding statistically significant communities in networks. *PLoS ONE*, 6(4):e18961, 2011.
- [97] C.J. Honey, J.-P. Thivierge, and O. Sporns. Can structure predict function in the human brain? *NeuroImage*, 52(3):766–776, 2010.
- [98] J.S. Damoiseaux and M.D. Greicius. Greater than the sum of its parts: a review of studies combining structural connectivity and resting-state functional connectivity. *Brain Structure and Function*, 213:525–533, 2009.
- [99] B. Ng, G. Varoquaux, J.-B. Poline, and B. Thirion. A novel sparse graphical approach for multimodal brain connectivity inference. In *Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, pages 707–714. 2012.
- [100] A. Venkataraman, Y. Rathi, M. Kubicki, C. Westin, and P. Golland. Joint modeling of anatomical and functional connectivity for population studies. *IEEE Transactions on Medical Imaging*, 31(2):164–182, 2012.
- [101] L. Doderer, A. Gozzi, A. Liska, V. Murino, and D. Sona. Group-wise functional community detection through joint Laplacian diagonaliza-

- tion. In *Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, pages 708–715. 2014.
- [102] B. Scherrer, A. Gholipour, and S.K. Warfield. Super-resolution in diffusion-weighted imaging. In *Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, pages 124–132. 2011.
- [103] A.L. Alexander, K.M. Hasan, M. Lazar, J.S. Tsuruda, and D.L. Parker. Analysis of partial volume effects in diffusion-tensor MRI. *Magnetic Resonance in Medicine*, 45(5):770–780, 2001.
- [104] S. Mori and J. Zhang. Principles of diffusion tensor imaging and its applications to basic neuroscience research. *Neuron*, 51(5):527–539, 2006.
- [105] S. Peled and Y. Yeshurun. Superresolution in MRI: application to human white matter fiber tract visualization by diffusion tensor imaging. *Magnetic Resonance in Medicine*, 45(1):29–35, 2001.
- [106] D.H.J. Poot, B. Jeurissen, Y. Bastiaensen, J. Veraart, W. Van Hecke, P.M. Parizel, and J. Sijbers. Super-resolution for multislice diffusion tensor imaging. *Magnetic Resonance in Medicine*, 69(1):103–113, 2013.
- [107] S.N. Sotiropoulos, S. Jbabdi, J.L. Andersson, M.W. Woolrich, K. Ugurbil, and T.E.J. Behrens. RubiX: combining spatial resolutions for Bayesian inference of crossing fibers in diffusion MRI. *IEEE Transactions on Medical Imaging*, 32(6):969–982, 2013.
- [108] V. Gupta, N. Ayache, and X. Pennec. Improving DTI resolution from a single clinical acquisition: a statistical approach using spatial prior. In *Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, pages 477–484. 2013.
- [109] P. Coupé, J.V. Manjón, M. Chamberland, M. Descoteaux, and B. Hiba. Collaborative patch-based super-resolution for diffusion-weighted images. *NeuroImage*, 83:245–261, 2013.
- [110] F. Calamante, J.-D. Tournier, G.D. Jackson, and A. Connelly. Track-density imaging (TDI): super-resolution white matter imaging using whole-brain track-density mapping. *NeuroImage*, 53(4):1233–1243, 2010.

- [111] R. Zeyde, M. Elad, and M. Protter. On single image scale-up using sparse representations. In *Curves and Surfaces*, pages 711–730. 2012.
- [112] R. Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 267–288, 1996.
- [113] M. Aharon, M. Elad, and A. Bruckstein. K-SVD: An algorithm for designing overcomplete dictionaries for sparse representation. *IEEE Transactions on Signal Processing*, 54(11):4311–4322, 2006.
- [114] Y.C. Pati, R. Rezaifar, and P. S. Krishnaprasad. Orthogonal matching pursuit: recursive function approximation with applications to wavelet decomposition. In *Asilomar Conference on Signals, Systems and Computers*, pages 40–44, 1993.
- [115] B.A. Landman, A.J. Huang, A. Gifford, D.S. Vikram, I.A.L. Lim, J.A.D. Farrell, J.A. Bogovic, J. Hua, M. Chen, S. Jarso, S.A. Smith, S. Joel, S. Mori, J.J. Pekar, P.B. Barker, J.L. Prince, and P.C.M. van Zijl. Multi-parametric neuroimaging reproducibility: a 3-T resource study. *NeuroImage*, 54(4):2854 – 2866, 2011.
- [116] F. Deleus and M.M. Van Hulle. A connectivity-based method for defining regions-of-interest in fMRI data. *IEEE Transactions on Image Processing*, 18(8):1760–1771, 2009.
- [117] R.A. Poldrack. Region of interest analysis for fMRI. *Social Cognitive and Affective Neuroscience*, 2(1):67–70, 2007.
- [118] B. Ng, G. Hamarneh, and R. Abugharbieh. Detecting brain activation in fMRI using group random walker. In *Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, pages 331–338. 2010.
- [119] V. Michel, A. Gramfort, G. Varoquaux, E. Eger, C. Keribin, and B. Thirion. A supervised clustering approach for fMRI-based inference of brain states. *Pattern Recognition*, 45(6):2041–2049, 2012.
- [120] P.F. Neher, B. Stieltjes, I. Wolf, H.P. Meinzer, and K.H. Maier-Hein. Analysis of tractography biases introduced by anisotropic voxels. In *Annual Meeting of the International Society for Magnetic Resonance in Medicine (ISMRM)*. 2013.

- [121] J.V. Manjón, P. Coupé, A. Buades, D.L. Collins, and M. Robles. New methods for MRI denoising based on sparseness and self-similarity. *Medical Image Analysis*, 16(1):18–27, 2012.
- [122] E. Garyfallidis, M. Brett, B. Amirbekian, A. Rokem, S. Van Der Walt, M. Descoteaux, and I. Nimmo-Smith. Dipy, a library for the analysis of diffusion MRI data. *Frontiers in Neuroinformatics*, 8(8), 2014.
- [123] D.C. Van Essen, S.M. Smith, D.M. Barch, T.E.J. Behrens, E. Yacoub, and K. Ugurbil. The WU-Minn Human Connectome Project: an overview. *NeuroImage*, 80:62–79, 2013.
- [124] M.F. Glasser, S.N. Sotiropoulos, J.A. Wilson, T.S. Coalson, B. Fischl, J.L. Andersson, J. Xu, S. Jbabdi, M. Webster, and J.R. Polimeni. The minimal preprocessing pipelines for the Human Connectome Project. *NeuroImage*, 80:105–124, 2013.
- [125] Y. Behzadi, K. Restom, J. Liau, and T.T. Liu. A component based noise correction method (CompCor) for BOLD and perfusion based fMRI. *NeuroImage*, 37(1):90–101, 2007.
- [126] K.H. Fritzsche, P.F. Neher, I. Reicht, T. van Bruggen, C. Goch, M. Reisert, M. Nolden, S. Zelzer, H.-P. Meinzer, and B. Stieltjes. MITK diffusion imaging. *Methods of Information in Medicine*, 51(5):441, 2012.
- [127] K. Murphy, R.M. Birn, D.A. Handwerker, T.B. Jones, and P.A. Bandettini. The impact of global signal regression on resting state correlations: are anti-correlated networks introduced? *NeuroImage*, 44(3):893–905, 2009.
- [128] B.P. Rogers, V.L. Morgan, A.T. Newton, and J.C. Gore. Assessing functional connectivity in the human brain by fMRI. *Magnetic Resonance Imaging*, 25(10):1347–1357, 2007.
- [129] B. Ng, G. Hamarneh, and R. Abugharbieh. Modeling brain activation in fMRI using group MRF. *IEEE Transactions on Medical Imaging*, 31(5):1113–1123, 2012.
- [130] Y. Benjamini and Y. Hochberg. Controlling the false discovery rate: a practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 289–300, 1995.

- [131] Z. Liu, N. Zhang, W. Chen, and B. He. Mapping the bilateral visual integration by EEG and fMRI. *NeuroImage*, 46(4):989–997, 2009.
- [132] L. Grady. Multilabel random walker image segmentation using prior models. In *IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, volume 1, pages 763–770, 2005.
- [133] L. Grady. Random walks for image segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 28(11):1768–1783, 2006.
- [134] N.V. Hartvig and J.L. Jensen. Spatial mixture modeling of fMRI data. *Human Brain Mapping*, 11(4):233–248, 2000.
- [135] A.P. Dempster, N.M. Laird, and D.B. Rubin. Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society, Series B*, 39(1):1–38, 1977.
- [136] K.J. Friston and W. Penny. Posterior probability maps and SPMS. *NeuroImage*, 19:12401249, 2003.
- [137] C.R. Genovese, N.A. Lazar, and T. Nichols. Thresholding of statistical maps in functional neuroimaging using the false discovery rate. *NeuroImage*, 15(4):870–878, 2002.
- [138] N. Toussaint, J.-C. Souplet, and P. Fillard. MedINRIA: medical image navigation and research tool by INRIA. In *Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, volume 7, page 280, 2007.
- [139] L. Wang, P.D. Metzak, W.G. Honer, and T.S. Woodward. Impaired efficiency of functional networks underlying episodic memory-for-context in schizophrenia. *The Journal of Neuroscience*, 30(39):13171–13179, 2010.
- [140] H. Eichenbaum. Conscious awareness, memory and the hippocampus. *Nature Neuroscience*, 2:775–776, 1999.
- [141] K. Henke, B. Weber, S. Kneifel, H.G. Wieser, and A. Buck. Human hippocampus associates information in memory. *Proceedings of the National Academy of Sciences*, 96(10):5884–5889, 1999.
- [142] K. Sullivan Giovanello, D.M. Schnyer, and M. Verfaellie. A critical role for the anterior hippocampus in relational memory: evidence from an

- fMRI study comparing associative and item recognition. *Hippocampus*, 14(1):5–8, 2004.
- [143] R.W. Heinrichs and K.K. Zakzanis. Neurocognitive deficit in schizophrenia: a quantitative review of the evidence. *Neuropsychology*, 12(3):426, 1998.
- [144] F. Waters, T. Woodward, P. Allen, A. Aleman, and I. Sommer. Self-recognition deficits in schizophrenia patients with auditory hallucinations: a meta-analysis of the literature. *Schizophrenia Bulletin*, 38(4):741–750, 2012.
- [145] O. Sporns. *Networks of the brain*. MIT Press, 2011.
- [146] S. Fortunato. Community detection in graphs. *Physics Reports*, 486(3):75–174, 2010.
- [147] P. Schuster and K. Sigmund. Replicator dynamics. *Journal of Theoretical Biology*, 100:533–538, 1983.
- [148] S. Karlin. Mathematical models, problems, and controversies of evolutionary theory. *Bulletin of the American Mathematical Society*, 10(2):221–273, 1984.
- [149] N. Meinshausen and P. Bühlmann. Stability selection. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 72(4):417–473, 2010.
- [150] K. Sigmund and J. Hofbauer. *The theory of evolution and dynamical systems*. Cambridge University Press, 1988.
- [151] A. Torsello, S. Rota Bulò, and M. Pelillo. Beyond partitions: Allowing overlapping groups in pairwise clustering. In *IEEE 19th International Conference on Pattern Recognition*, pages 1–4, 2008.
- [152] G. Lohmann and S. Bohn. Using replicator dynamics for analyzing fMRI data of the human brain. *IEEE Transactions on Medical Imaging*, 21(5):485–492, 2002.
- [153] B. Thirion, P. Pinel, A. Tucholka, A. Roche, P. Ciuciu, J.-F. Mangin, and J.-B. Poline. Structural analysis of fMRI data revisited: Improving the sensitivity and reliability of fMRI group studies. *IEEE Transactions on Medical Imaging*, 26(9):1256–1269, 2007.

- [154] B. Efron. *The jackknife, the bootstrap and other resampling plans*, volume 38. SIAM, 1982.
- [155] M. Lazar and A.L. Alexander. Bootstrap white matter tractography (BOOT-TRAC). *NeuroImage*, 24(2):524–532, 2005.
- [156] M. van den Heuvel, R. Mandl, and H. Hulshoff Pol. Normalized cut group clustering of resting-state fMRI data. *PLoS ONE*, 3(4):e2001, 2008.
- [157] H.W. Kuhn. The Hungarian method for the assignment problem. *Naval Research Logistics Quarterly*, 2(1-2):83–97, 1955.
- [158] M. Rubinov and O. Sporns. Complex network measures of brain connectivity: uses and interpretations. *NeuroImage*, 52(3):1059–1069, 2010.
- [159] Q. Wang and E. Fleury. Overlapping community structure and modular overlaps in complex networks. In *Mining Social Networks and Security Informatics*, pages 15–40. 2013.
- [160] M.P. Van Den Heuvel and H.E.H. Pol. Exploring the brain network: a review on resting-state fMRI functional connectivity. *European Neuropsychopharmacology*, 20(8):519–534, 2010.
- [161] C.A. Seger. The visual corticostriatal loop through the tail of the caudate: circuitry and function. *Frontiers in Systems Neuroscience*, 7, 2013.
- [162] S.J. Palmer, J. Li, Z.J. Wang, and M.J. McKeown. Joint amplitude and connectivity compensatory mechanisms in Parkinson’s disease. *Neuroscience*, 166(4):1110–1118, 2010.
- [163] X. Wen, L. Yao, T. Fan, X. Wu, and J. Liu. The spatial pattern of basal ganglia network: A resting state fMRI study. In *IEEE International Conference on Complex Medical Engineering*, pages 43–46, 2012.
- [164] Y. He, Z.J. Chen, and A.C. Evans. Small-world anatomical networks in the human brain revealed by cortical thickness from MRI. *Cerebral Cortex*, 17(10):2407–2419, 2007.