

Distribution Agreement

In presenting this thesis or dissertation as a partial fulfillment of the requirements for an advanced degree from Emory University, I hereby grant to Emory University and its agents the non-exclusive license to archive, make accessible, and display my thesis or dissertation in whole or in part in all forms of media, now or hereafter known, including display on the world wide web. I understand that I may select some access restrictions as part of the online submission of this thesis or dissertation. I retain all ownership rights to the copyright of the thesis or dissertation. I also retain the right to use in future works (such as articles or books) all or part of this thesis or dissertation.

Signature:

Xavier A. Higgins

Date

Statistical Methods for Estimating and Analyzing Brain Functional Connectivity

By

Ixavier A. Higgins

Doctor of Philosophy

Biostatistics

Ying Guo, Ph.D.
Advisor

Suprateek Kundu, Ph.D.
Advisor

Ki Sueng Choi, Ph.D.
Committee Member

Helen S. Mayberg, M.D.
Committee Member

Lance Waller, Ph.D.
Committee Member

Accepted:

Lisa A. Tedesco, Ph.D.
Dean of the James T. Laney School of Graduate Studies

Date

Statistical Methods for Estimating and Analyzing Brain Functional Connectivity

By

Ixavier A. Higgins

M.S., Emory University, 2017

M.A., Wake Forest University, 2013

B.S., Wake Forest University, 2011

Advisor: Ying Guo, Ph.D. and Suprateek Kundu, Ph.D.

An abstract of

A dissertation submitted to the Faculty of the

James T. Laney School of Graduate Studies of Emory University

in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

in Biostatistics

2019

Abstract

Statistical Methods for Estimating and Analyzing Brain Functional Connectivity

By

Ixavier A. Higgins

In recent years, network science has become an increasingly popular approach for investigating interdependencies between spatially distinct brain regions. The network based paradigm is particularly useful in neuroimaging applications because of the ease of representing complex spatiotemporal relationships via a finite set of regions. Analyses of functional and structural architecture have elucidated the mechanisms by which neurological disorders distort local and global functional organization, inhibiting normal brain processes.

In chapter two, we propose a Differential Degree Test (DDT) which detects regions of interest (ROIs) incident to a statistically significant number of edges that are differentially weighted across healthy and depressed populations. We achieve this by generating null networks in which edge weights match distributional properties of edges in the observed difference network. Extensive numerical studies demonstrate superior performance relative to popular network comparison methods. We apply the method to major depressive disorder patients and age-matched healthy controls. Our method selects ROIs commonly implicated in studies of depression.

In chapter three, we propose a structurally informed Gaussian Graphical Model (siGGM) that incorporate structural connectivity into the estimation of functional connections between all region pairs. Although the exact relationship between brain function and structure is not completely known, anatomical wiring certainly constrains cortical activity. Our multimodal approach requires rs-fMRI and diffusion tensor imaging (DTI) which maps the orientation of all white matter fiber tracks in the brain. Our efficient optimization algorithm admits a MAP solution of subject-specific functional brain networks. Numerical studies and an application to sixty-nine individuals in the Philadelphia Neurodevelopment Cohort demonstrate our method's superior performance to state of the art competitors.

In chapter four, we investigate rapidly changing functional connectivity. Recent work suggests that the brain utilizes a finite set of connectivity states that are common across all health conditions. We propose a semi-parametric dictionary learning method to simultaneously estimate the shared set of brain networks as well as classify individuals into disease groups based upon usage of the basis set. We assess the method's performance on simulated data and detect biologically meaningful brain networks in a study of posttraumatic stress disorder.

Statistical Methods for Estimating and Analyzing Brain Functional Connectivity

By

Ixavier A. Higgins

M.S., Emory University, 2017

M.A., Wake Forest University, 2013

B.S., Wake Forest University, 2011

Advisor: Ying Guo, Ph.D. and Suprateek Kundu, Ph.D.

A dissertation submitted to the Faculty of the
James T. Laney School of Graduate Studies of Emory University
in partial fulfillment of the requirements for the degree of
Doctor of Philosophy
in Biostatistics

2019

Acknowledgement

First and foremost, I would like to thank my Lord and Savior Jesus Christ. It is truly through Him that all things are possible. I am extremely grateful for the many individuals He has placed in my life who have been essential to my success.

I would like to thank my dissertation advisors, Drs. Ying Guo and Suprateek Kundu, who have pushed me to achieve things I didn't think possible. I've learned tremendous amounts from you not only about neuroimaging research, but more importantly about my capacity to learn quickly, think critically, and write effectively. I consider this dissertation process a success because of the many skills I have learned under your guidance. I am also extremely lucky to have Drs. Lance Waller, Helen Mayberg and Ki Sueng Choi serve on my dissertation committee. This dissertation would be nothing without your amazing insights on the scientific questions I've addressed. Your thoughtful comments always helped me refocus on the broader goals of the developed methods when I became bogged down in the technical details.

I owe much of my success to my mentors, Dr. DuBois Bowman, Dr. Renee Moore, and Dr. Sean Simpson. Despite your busy schedules, you've always made time to talk with me about academic planning, professional development, and career prospects. No matter how minor the issue, you always made me feel that my concerns were valid and provided advice on how to proceed. My hope is to share with future generations of scholars all of the knowledge that you have poured into with me.

I am also grateful for the friends and family I have developed within the department and across the Emory community. First, I want to thank Mary, Bob, Melissa and Angela for all that they do for this department. Every success of the students and faculty is a credit to your professionalism and excellent knowledge of Emory regulations. There's no greater group at Emory and I am lucky to have had the chance

to work with you and get to know you over the last few years. My current officemates, Jeong Jang, Jay Soh, Behzad Kianian, and Gavin Tian, as well as past officemates Teng Fei, Dave Van Domelen and Shawnee Anderson, made the office a fun place to spend many hours each day. You guys have pushed me to work harder and kept me sane throughout the process. There are also many, many others that have been instrumental to my happiness and success. To my fellow friends in CBIS—in particular Josh, Raphiel, Yikai, Ben Wu, Xin, Jin, and Yingtian—thank you for your help over the years working on various projects. I learned a tremendous amount from each of you. I'd like to thank my dear friend, Dr. Caprichia Jeffers, who always looked out for me from my first day in the department. In watching you conquer your challenges, I found the strength to push through my obstacles. There are many, many others that have helped me push through the dissertation writing process, worked out with me in the gym, shared meals with me after long days in the office, and worked with me to plan social and academic activities. I am eternally grateful for each and every one of you.

I don't know where I would be today without "The Circle"— Terrol Graham, Geoffrey Curtis, Dr. Lavern Keitt, Dr. Arif Ibrahim, Damien Boston, Arturo Caesar, Chris Faulkner, Ajee Elliott, Rakim Austin, and Dr. Jason Phillippe. You guys have helped me work through countless personal and professional decisions while guaranteeing laughs day in and day out. You encouraged me to fight through the many unforeseen obstacles that arise during the course of a doctoral program. I'm truly grateful that we all chose Wake Forest in 2007 and made the effort to get to know each other. I'm even more proud that twelve years later, we all still talk to each other every day. Thank you for being who you are and I look forward to many more decades of friendship. I also owe many thanks to Toni Bolar, who has been a steady, reliable friend for as long as I can remember. And finally, I have to thank my best friend from childhood, James Rapley, and his wife Leah. You two have been

everything I needed at every stage of my life thus far. I've learned a lot from watching you two and look forward to our kids growing up together one day.

Finally, I want to dedicate this dissertation to my village, my praise team, my family. To my wife and best friend, Mikela Gordon-Higgins, I want to thank you for your unwavering support, kindness, understanding, and amazing jokes! I would have lost my mind 1000 times by now if it had not been for you. You kept me steady, focused, and in good spirits. I cannot say thank enough times to my mother-in-law and father-in-law, Debra and Michael Gordon. Your kind words, timely advice, and warm meals on the weekends were exactly what I needed. You opened your doors to me when Mikela and I were just friends that enjoyed each others' company. You are the reason I call Atlanta home and I'm so blessed that we are now family! Most importantly, this would not have been possible without my parents, Cheryl and Kenny Ware. You have made so many sacrifices for me, Buck and LB. You certainly don't get enough credit for what you have done for us and I know that I would not be here today if not for the hard decisions you made to give us the best future. This journey is a testament of your commitment to my success. You've never told me a dream was too big, even when I spent years planning a career as a NASA astronaut. You've been my biggest cheerleaders and my sounding board as I waded through various career and personal decisions. I couldn't have been blessed with better parents.

Contents

1	Introduction	1
1.1	Overview	2
1.1.1	Basics of functional magnetic resonance	5
1.1.1.1	Resting state fMRI	8
1.1.2	Basics of Diffusion Tensor Imaging	8
1.1.3	Brain structure and function	10
1.2	Brain Connectome	11
1.2.1	Static Functional connectivity Methods	14
1.2.2	Dynamic functional connectivity Methods	16
1.2.3	Challenges in functional connectivity	17
1.3	Motivating Data	18
1.3.1	Major Depressive Disorder dataset	19
1.3.2	Philadelphia Neurodevelopmental Cohort dataset	20
1.3.3	Grady Trauma Project dataset	22
1.4	Proposed Research	23
1.4.1	Topic 1: Difference degree test for comparing brain networks .	24
1.4.2	Topic 2: Anatomically informed estimation of brain functional networks	24
1.4.3	Topic 3: Semi-parametric Bayesian hierarchical dictionary learn- ing	25

2	Comparison of nodal differential degree centrality	26
2.1	Introduction	27
2.2	Method	31
2.2.1	Brain network construction	31
2.2.2	Differential Degree Test	33
2.2.2.1	Difference Network Construction	33
2.2.2.2	Deriving p-value from between-group tests	35
2.2.2.3	Null distribution generation	36
2.2.2.4	An adaptive threshold selection method	38
2.2.2.5	The DDT Test	40
2.3	Simulation	41
2.3.1	Results	44
2.4	Data Application	48
2.4.1	Data	49
2.4.2	Results	50
2.5	Discussion	55
3	Anatomically Informed Estimation of Functional Brain Networks	57
3.1	Introduction	58
3.2	Materials and methods	62
3.2.1	Gaussian graphical model for brain networks	62
3.2.2	Structurally informed Bayesian Gaussian graphical model	64
3.2.3	Model Estimation	68
3.3	Results and Discussion	70
3.3.1	Simulations	70
3.3.1.1	Simulation Setting	70
3.3.1.2	Results	74
3.3.2	PNC Data Application	79

3.3.2.1	Data preprocessing	80
3.3.2.2	Results	80
3.4	Conclusion	86
4	Semi-parametric Bayesian hierarchical dictionary learning	88
4.1	Introduction	89
4.2	Methodology	93
4.2.1	Dictionary learning	93
4.2.2	<u>B</u> ayesian <u>D</u> ictionary <u>L</u> earning (BayDiL)	94
4.2.3	Posterior Sampling	96
4.2.4	Subcase: Regional time series as Observed Signal	98
4.3	Results and Discussion	100
4.3.1	Simulations	100
4.3.1.1	Simulation Setting	100
4.3.1.2	Results	102
4.3.2	Posttraumatic Stress Disorder Data Application	106
4.3.2.1	Results	106
4.4	Conclusion	112
5	Summary and Future Directions	115
A	Appendix for Chapter 2	119
A.1	Proof for HQS procedure	119
A.2	Tables	122
B	Appendix for Chapter 3	123
B.1	Parameter Updates	123
B.2	Hyperparameter Choice and Initial Values	125
B.3	Measure for computing between module differences	126

B.4	Calculation of ICC	126
B.5	Tables	127
B.6	Network Metrics	128
C	Appendix for Chapter 4	130
C.1	Hyper-parameter selection and initial values	130
C.2	Figures	132
	Bibliography	135

List of Figures

1.1	The brain is composed of grey matter, white matter, and cerebrospinal fluid.	3
1.2	Communication between neurons in the brain.	11
1.3	Estimation of sliding windows correlations.	17
2.1	Comparison of eDDT, aDDT, t-test ($T_{(10\%)}$), binomial tests (Bin_F , Bin_B) on simulated data with three differentially connected nodes incident to 4, 7, and 11 DWEs	46
2.2	Performance of the methods as the network size increases with a fixed density	47
2.3	Performance of the methods in selecting DWEs	48
2.4	DWEs detected by eDDT under four model assumptions	52
2.5	Heat map of $X^2_{(g1,g2)}$ statistic	54
3.1	FC and SC association observed in the data	60
3.2	Distribution of ω_{jk} as a function of hyperparameters and SC	67
3.3	Graphical illustration of the siGGM	68
3.4	Comparison of siGGM, G-Wishart, and aGlasso for p=100 regions	76
3.5	Results for data simulated from non-Gaussian ICA model	77
3.6	siGGM computation time	77
3.7	Gender-stratified correlates of FC and SC	81

3.8	Network estimates for females and males in the PNC data application	82
3.9	Comparison of network features across males and females	83
3.10	Reliability of network metrics estimated by siGGM, aGlasso, and glasso	86
4.1	Comparison of the the BayDiL and DL methods in recovering the basis signals when the signals share 1, 2 and 3 active regions in common.	104
4.2	BayDiL performance with misspecified number of atoms	105
4.3	Covariance matrices of the five basis components estimated by the BayDiL method	108
4.4	Covariance matrices of the five basis components estimated by the DL method	108
4.5	Maps of co-activation in each atom for the BayDiL	109
4.6	Posterior distributions of the length scale parameters for the BayDiL estimates	111
C.2.1	Comparison of network estimates for Square Exponential Kernel sim- ulations	132
C.2.2	Comparison of network estimates for Sinusoidal simulations	133
C.2.3	Twenty seven brain regions organized by functional module	134

List of Tables

2.1	False positive and true positive rates for the Random, Small World, and Hybrid network structures considered.	45
2.2	Top 20 differentially connected nodes in the MDD data application .	51
2.3	Consistency of identified DWEs	53
3.1	Performance of SC-informed methods on simulated data	78
3.2	Within- and between- module differences in FC between males and females	85
4.1	Dictionary atom recovery success rate in simulations	103
4.2	Average classification accuracy as measured by the adjusted rand index for the BayDiL method.	104
4.3	Average score in the intrusion, avoidance/numbing, and hyperarousal domains. Reported values are the mean and standard deviation.	111
A.1	Within and between functional module DWE in the major depressive disorder study. Bolded values indicate statistically significant number of DWE between the respective functional modules	122
B.1	Performance of siGGM and SC naive approaches on simulated network data with $p=100$ and 200 nodes. Eglob is the bias in global efficiency.	127

Chapter 1

Introduction

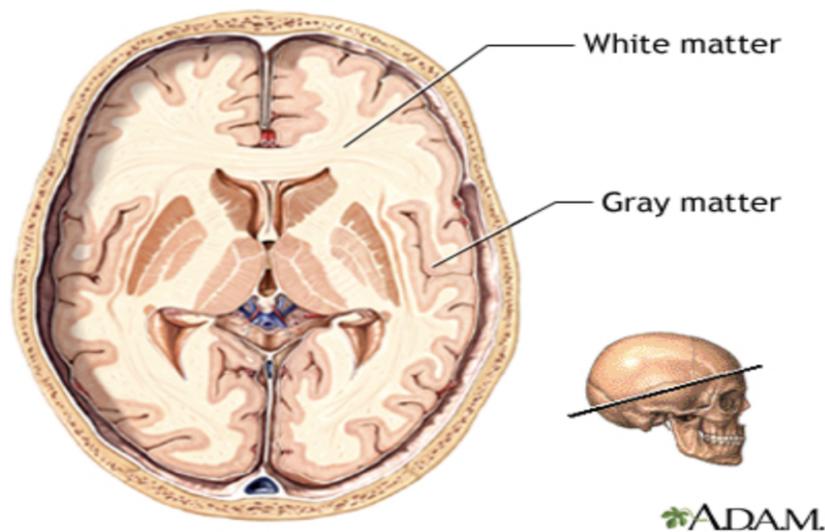
1.1 Overview

Mental health disorders is an oft used umbrella term encompassing a heterogeneous pool of illnesses impacting an individual's mood, behavior, and ability to think. Recent estimates indicate that these disorders affect as many as one in five adults and children in a given year, while one in two will experience a mental disorder at least once in life. The disorders can diminish one's quality of life, increasing the risk of other health problems, poverty, social problems, and financial strain. Sadly, those that suffer from depression are at an increased risk of suicide, the second leading cause of death for individuals ages 10 to 34. The side effects of mental disorders are not the sole responsibility of affected individuals. In 2010, mental health disorders produced a \$200 billion burden on the US economy (Roehrig, 2016). Clearly, mental health disorders are a public health crisis requiring comprehensive investigations into the biological mechanisms of diseases and the identification of biomarkers that can guide treatment.

Mental health disorders are thought to be caused by a nebulous combination of inherited genetic abnormalities, environmental exposures, and chemical imbalances in the brain. Since the disorders are generally characterized by poor mood and decreased cognitive abilities, many efforts have investigated the structure and function of the human brain in health and disease. The brain is a small organ weighing approximately three pounds and is composed of 75% water. As the primary operating system in the body, the brain is responsible for all bodily functions, interpretation of the external world, response to external stimuli, emotion regulation, and cognitive functioning. Many of these duties are the responsibility of the cerebrum, which is subdivided into four lobes— frontal, temporal, parietal, occipital— in the left and right hemispheres. The outermost layer of the cerebrum, called the cortex, is composed of approximately 16 billion neurons densely packed into the gray matter brain tissue (<https://mayfieldclinic.com/pe-anatbrain.htm>). Neurons are specialized nerve

cells that are responsible for communicating information via chemical and electrical signals. All brain function results from coordinated communication between these small cells. Myelinated shafts attached to each neuron transmit the cell body's action potential to other local or long range neurons in the brain. These bundles of axons compose the white matter tissue.

Figure 1.1: The brain is composed of grey matter, white matter, and cerebrospinal fluid.



(<https://medlineplus.gov/ency/imagepages/18117.htm>)

Our understanding of the brain's anatomy and function has rapidly expanded over the last century due to the rapid development of advanced imaging technologies. Consequentially, we have invaluable insights into the complexity of human cognition. Prior to the advent of powerful brain imaging machines, analysis of the human brain was only a possibility postmortem, when the brain could be safely removed and dissected. While this produced crucial insights into the brain's structural architecture, little was known about brain function in health and in disease.

Within the last fifty years, non-invasive imaging technologies have provided well

defined, in vivo images of the brain. In the 1970s, computerized axial tomography (CAT) was introduced to define anatomic structures based on differential absorption of x-rays into the three brain tissues (grey matter, white matter, and cerebrospinal fluid (CSF)) . Similarly, T1-weighted images derived from fMRI provide high resolution visuals of the tissues in the brain and diffusion weighted imaging permits estimation of the direction of white matter fiber bundles connecting cortical brain regions.

During the 1970s, the discovery of radioligands led to the development of positron emission tomography (PET). PET scans measure the emitted signal from radioactively labelled chemicals injected into the patient's blood stream. Emissions within brain tissue indicate functional activity within the region. During this same era, magnetic resonance imaging (MRI) was developed to measure functional activity in the brain. MRI gained significant traction in the 1990s and has become the dominant imaging modality due to its low invasiveness, limited radiation exposure, and ease of acquisition across a wide spectrum of disorders.

Although PET and fMRI leverage blood flow as a proxy for functional activity, other technologies exist to measure brain function. Magnetoencephalography (MEG) quantifies brain activity via changes in the magnetic field due to altered electrical currents in the brain. Electroencephalography (EEG) detects electrical currents in neurons, which are considered the basic functional processing unit in the brain. While MEG and EEG do not expose patients to radiation, they lack spatial specificity and signal detection can be hindered by the skull. Each of these modalities provide different insights into brain function.

With these tools, researchers have begun to enumerate the processes involved in simple motor control such as finger tapping to more complex tasks such as encoding memories. In healthy populations, spatially remote regions efficiently communicate in order to accomplish the various tasks. However, in many disorders, the brain does

not relay signals efficiently which can severely impact cognitive ability. Imaging technologies permit localization of irregular functional activity and damage to structural pathways that define various mental health disorders and diseases.

1.1.1 Basics of functional magnetic resonance

Functional magnetic resonance imaging (fMRI) is a noninvasive method for measuring neuronal activity in the brain. Due to the ease of data acquisition, it has become one of the most popular imaging techniques. Functional magnetic resonance imaging is a technique for measuring the quantity of oxygenated blood present throughout the brain. Biologically, neurons require oxygen to perform any action. Thus, demands for oxygen rich blood, which fMRI is designed to detect, serves as a proxy for functional activity.

Measuring changes in the blood oxygen level dependent (BOLD) signal is an extension of magnetic resonance imaging (MRI). Hydrogen atoms, which are abundant in water that is widespread throughout the brain, are key to the success of MRI. In MRI, a strong magnetic field (B_0) causes hydrogen atoms to orient parallel to the direction of field and imposes a constant phase rate on all atoms. The strength of the magnetic field is typically 1.5T or 3T, although higher Tesla fields have been safely applied to human subjects to acquire images with finer spatial and temporal resolution. The alignment of the atoms creates a net magnetization in the direction of (B_0). Subsequently, a radio frequency (RF) pulse is repeatedly applied to the field to “excite” the atoms. In this excited state, atoms are oriented perpendicular to (B_0) and exhibit different phases rate. Between RF pulses, atoms attempt to reorient to the low energy state parallel to (B_0), which produces signals that are detected by the scanner coils. These signals are referred to as the T1 relaxation time, T2 and T2* relaxation measures. T1 relaxation defines the rate at which the atoms become reoriented to (B_0) following a RF pulse and T2/T2* measure the rate of signal decay

transverse to the constant magnetic field. The three brain tissues exhibit different T1 and T2/T2* relaxation times, producing images that clearly delineate their locations throughout the brain. T1-weighted images have high spatial resolution but take significantly more time to acquire than T2-weighted images. Conversely, T2-weighted images can be acquired quickly at regular intervals, making them highly useful in studies of the brain over time, i.e. a scanning session.

As previously mentioned, fMRI works by measuring fluctuations in the BOLD signal. The consumption of glucose by active neurons requires an influx of oxygenated blood, which creates an overabundance of oxygenated hemoglobin relative to deoxygenated hemoglobin. Hemoglobin has differential responses to magnetic fields that is directly related to whether it is bound to oxygen. Deoxygenated hemoglobin is attracted to magnetic fields, leading to distortions in the magnetic field, (B_0). However, oxygenated hemoglobin loses magnetization slower via the T2* decay, leading to stronger signals in areas where the oxygen is in abundance. Contrasts in this BOLD signal over time are key to the success of fMRI in measuring neuronal activity.

The change in the fMRI BOLD signal is assumed to be a smooth process, termed the hemodynamic response function (HRF) in the imaging community. Interestingly, the HRF reaches its peak value approximately five seconds after stimulus onset. The delayed response is a result of the vascular system requiring time to respond to the demand for resources. Ideally, the HRF is perfectly observed for active neurons in the brain. Unfortunately, the true signal is mixed with various sources of noise. First, the hardware introduces noise, such as scanner drift, heat up, and inhomogeneities in the magnetic field, that can produce BOLD fluctuations unrelated to neuronal activity. A second source of noise severely impacting BOLD fMRI is that of physiological noise. Patient movement, cardiovascular and respiratory processes lead to incorrect voxel-to-neuron mapping during the scanning session. These noise artifacts can completely obscure the signals of interest. In a 1.5T scanner, neuronal induced fluctuations in the

BOLD signal rarely exceed 3%. Thus, it is of extreme importance to appropriately clean the data to reduce the influence of noise prior to analysis.

Traditional preprocessing pipelines include many steps to increase the signal-to-noise ratio. Although there are variations in which steps are performed, most preprocessing pipelines include the following steps: (1) extraction of the brain matter from the skull; (2) slice timing correction; (3) motion correction to align all images to a reference image; (4) co-registration of the T2*-weighted image to the T1 anatomical image; (5) spatial registration to map the brain volumes to a standard space; (6) spatial smoothing to blur residual differences and boost signal-to-noise ratio; and (7) temporal smoothing to remove system and physiological noise artifacts. Software, such as FSL and AFNI, provide a suite of functions for performing these steps. Additionally, researchers often standardize the preprocessed BOLD signals to reflect a % change in the signal.

Ideally, the processed data reflects the desired signal in each voxel sampled at a regular interval across the duration of the scan. At each acquisition point (i.e. brain volume), there are signals for hundreds of thousands of voxels, each of which is a volumetric unit containing thousands of spatially proximal neurons densely packed in the cortex. A typical fMRI dataset consists of 100-200 brain volumes for each subject, where each volume contains signals for hundreds of thousands of voxels. Clearly, fMRI data is high dimensional and requires significant storage capacity for even moderately-size sample populations. These factors must be accounted for when developing statistical methods. The network-based methods presented in this dissertation do not sacrifice data quality in order to reduce the computational burdens of working with these large datasets. To accomplish this, we harness network theory (see section 1.2 for an overview).

1.1.1.1 Resting state fMRI

Thus far, we refer to neuronal activity in response to stimuli. While task based study designs elicit response from targeted brain regions, BOLD signal fluctuations persist at rest (i.e. task-free). This paradigm, referred to as resting-state fMRI (rs-fMRI), reflects spontaneous neuronal activity. Despite the absence of a regulated external stimulus, various studies have shown that the brain organizes into well defined subcomponents at rest. Smith et al. (2009a) show that the brain at rest organizes into well defined subcomponents that are consistent with modules observed in subjects engaged in cognitive tasks. More specifically, the authors present convincing evidence of the following functional modules: medial, visual, and occipital visual network; default mode network; sensorimotor network; cerebellum; auditory network; executive control network; frontoparietal task control network. In recent years, the default mode network, which is active in the resting state, has been implicated in several disorders such as depression, schizophrenia, and ADHD. Resting state fMRI is also very useful because it is amenable to a variety of research hypotheses because of the task-free study design. As such, rs-fMRI has become a highly utilized technique for studying irregular brain function of many mental disorders and diseases. In this dissertation, we exclusively focus on rs-fMRI.

1.1.2 Basics of Diffusion Tensor Imaging

Diffusion tensor imaging (DTI) is a MRI technique for measuring the orientation of white matter fiber tracts in the brain. Although T1-weighted images are excellent at defining the white matter, it does not provide information on the orientation of the myelinated axons that transmit signals between neurons. Similarly to fMRI, the success of DTI is due to the abundance of water molecules in the brain. In DTI, magnetic gradients are applied to the brain in different directions to induce movement in the water molecules which diffuse randomly in unrestricted environments.

The three brain tissues constrain water diffusion in distinctly different manners and it is these differences that make DTI successful. The high density of gray matter prohibits diffusion whereas water generally moves freely in all directions in the CSF. Importantly, water diffusivity in the white matter tissue is constrained by the axons. Molecules diffuse very slowly across a bundle of axons but quickly parallel to a fiber.

As with fMRI, raw DTI data contains substantial sources of noise created by the subject and the shifting magnetization inside the scanner. Soares et al. (2013) state that the two main artifacts severely impacting tensor estimation are eddy currents and head motion, which can both be corrected by registering all acquired volumes to the diffusion free brain volume. Although there are many variations in DTI preprocessing pipelines, most include the following steps to remove nuisance signals in the data: (1) skull stripping; (2) eddy current correction; (3) motion correction; (4) registration to a standard brain space (i.e. MNI, Talairach); (5) fit diffusion tensors at each voxel in the brain; (6) build up a distribution of tensors at each voxel via Markov Chain Monte Carlo sampling; (7) perform tractography to estimate the fiber bundles connecting all user-defined seed regions (i.e. voxel, a collection of voxels).

We take a closer look at the tensor estimation step, that is required to transform signal changes along each gradient to information on the orientation of fiber bundles. At each voxel, a tensor, which is a three dimensional ellipsoid, is estimated. Mathematically, the tensor is defined by three eigenvectors (e^1, e^2, e^3) and associated eigenvalues ($\lambda_1, \lambda_2, \lambda_3$) (Le Bihan et al., 2001). The first eigenvector is the principle direction of diffusivity and the associated eigenvalue is the rate of diffusivity along that direction. The tensor is depicted as a three-dimensional ellipsoid, where a spherical volume (isotropic) indicates equal diffusivity in all directions and an elongated, slender volume (anisotropic) reflects directed diffusivity. In a simple environment, this information should be sufficient for capturing the dominant direction of diffusion at each voxel. However, voxels often contain crossing fibers. If the fibers are perpen-

dicular, it is possible to estimate an isotropic tensor which erroneously indicates no dominant direction of local diffusion. This local inaccuracy can potentially lead to erroneous global structural connectivity estimates.

In recent years, advanced modeling techniques have been developed to mitigate such issues, such as multi tensor models, High Angular Resolution Diffusion Imaging (HARDI), Q-Ball imaging, and Spherical Deconvolution (Soares et al., 2013). From the tensors estimates, tractography can be performed to estimate the number of bundles connecting all pairs of voxels or brain regions. In deterministic tractography, tensor orientations are fixed and fully determine the direction in which streamlines can pass through spatially proximal voxels. In probabilistic tractography, tensors at each voxel are sampled from a voxel-specific distribution. The variability in the tensor orientations allow the streamlines to evolve along different paths over many samples. Hence, probabilistic tractography permits quantification of the probability of structural connectivity between any two brain regions.

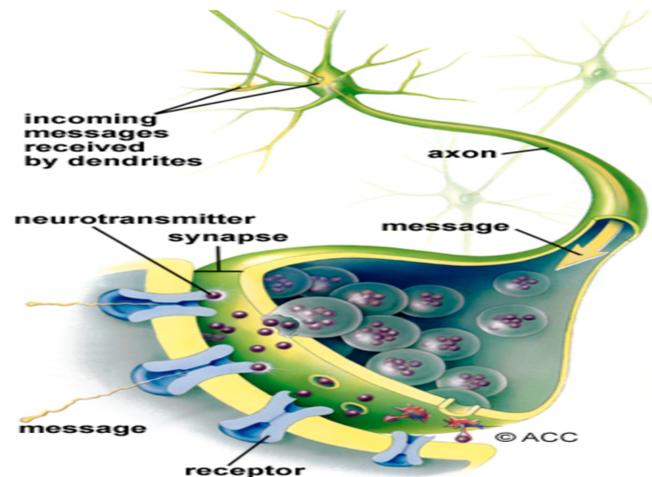
1.1.3 Brain structure and function

The previous sections discuss the acquisition and preprocessing of functional and structural MRI data. In recent years, the neuroimaging community has begun to investigate the relationship between brain function and structure Honey et al. (2009, 2010); Skudlarski et al. (2008). Figure 1.2 shows how neurons, the fundamental functional processing units in the brain, relay messages along myelinated axons. This arrangement permits communication between spatially remote cortical areas. Many studies agree that brain function is constrained by structure. Kemmer et al. (2018) show that coherent functional activity in patients at rest is strongest among brain areas with strong structural connections. However, strong functional connections have also been observed between structurally disconnected brain areas and weak functional connectivity between strongly connected regions (Greicius et al., 2009;

Van Den Heuvel and Pol, 2010).

Although the relationship is apparent, quantifying brain function co-activity as a deterministic function of structural connectivity has proven difficult. Correct specification should not only lead to more robust estimates of functional connectivity, but also will usher new insights into the physiology of disease in the brain. We investigate the structure-function relationship in chapter three.

Figure 1.2: Communication between neurons in the brain.



(<https://mayfieldclinic.com/pe-anatbrain.htm>)

1.2 Brain Connectome

The high definition images acquired from modern MRI scanners offer important insights into cortical activity. Shortly after the invention of fMRI, the first connectivity studies examined the level of co-activity in the time courses of brain regions (Van Den Heuvel and Pol, 2010). Biswal et al. (1995, 1997) were the first to report a high degree of correlation between regions in the motor network at rest, suggesting that the brain is functionally organized even at rest. Subsequent investigations found similar patterns in the visual and auditory networks as well as higher order cognitive networks (Greicius et al., 2003; Damoiseaux et al., 2006; Cordes et al., 2000). Recent

efforts seek to link cognition with functional connectivity patterns. However, this is extremely difficult due to the complexity of interactions in the brain responsible for human actions and behaviors. Over the last twenty years, researchers have investigated the brain as an integrative network of functionally coupled brain regions (Van Den Heuvel and Pol, 2010). In this simplistic representation of functional co-activation, spatio-temporal relationships between brain regions can be accounted for by a network structure, $\mathcal{G} = \langle V, E \rangle$. The vertex set, V , enumerates the brain regions of interest (ROI) and should be defined based upon the research hypothesis. Many parcellation schemes exist for defining V , including the Automated Anatomical Labelling Atlas, which consists of 90 contiguous cortical regions, as well as clustering based approaches to group voxels into homogenous regions based on functional similarity (Power et al., 2011a; Craddock et al., 2012; Yeo et al., 2011; Shen et al., 2013). In rare instances, each voxel could serve as a ROI, although this would lead to computational complications in downstream analyses as well as spurious interregional associations that are byproducts of in spatially smoothing raw MRI data.

The edge set, E , reflects all pairwise associations between nodes in V . In resting-state analysis, we acquire the BOLD signals in every brain region at T discrete, regularly sampled timepoints throughout the scan acquisition. Without loss of generality, assume that $v_1, v_2 \in V$ where for each region we have $v_1(1), \dots, v_1(T)$ and $v_2(1), \dots, v_2(T)$. E can represent any desired association measure such as between Pearson correlation, partial correlation, mutual information, or lag- t correlation coefficient. Functional connectivity can generally be classified into two domains: (a) undirected measures of association and (b) directed measures of association. Throughout this dissertation, we focus on undirected functional connectivity whereby there is no associated direction of information flow between two regions. Two popular undirected measures are Pearson and partial correlations. Pearson correlation is easiest to calculate, but suffers from correlation transitivity, which is excess correlation be-

tween two ROIs due to associations with tertiary brain regions (Zalesky, Fornito and Bullmore, 2012). Subsequent edge sets are often overly dense, masking biologically and statistically meaningful functional connections.

Partial correlations quantify the association between any two regions conditional on all brain regions, which inherently reduces the influence of correlation transitivity. Friedman et al. (2008) were the first to estimate partial correlations via optimization of a penalized likelihood

$$\hat{\mathbf{\Omega}} = \arg \min_{\mathbf{\Omega}} -\log(\det(\mathbf{\Omega})) - \text{tr}(S\mathbf{\Omega}) + \lambda \sum_{i,j} |\omega_{ij}| \quad (1.1)$$

where $\mathbf{\Omega}$ is the precision matrix. Whittaker (1990) provide the explicit function for transforming the precision matrix into interregional partial correlation estimates. The tuning parameter, λ , controls the sparsity of $\hat{\mathbf{\Omega}}$ and controls the fit of the model to observed data. Partial correlation based approaches such as this have gained favor in the neuroimaging community due to the widely accepted hypothesis that the brain maintains a small set of interregional connections to increase overall network efficiency (Bassett et al., 2009). Although eq 1.1 has many merits, selection of an appropriate tuning parameter is notoriously difficult. Specifically, tuning parameters that optimize model fit criteria such as BIC and AIC may not produce biologically consistent network estimates. There is clearly a need to modify this powerful technique for brain imaging applications. We tackle this problem in chapter 3.

Networks have led to valuable insights on the brain’s functionality in health and disease. For example, irregular communication between regions in the default mode network have been implicated in depression populations, while impaired activity of the fronto-parietal task control network has been implicated in various disorders (Cole et al., 2014). Statistical methods for identifying connectivity based biomarkers of disease are very diverse. Broadly, the methods fall into one of two categories. In the first,

it is assumed that regional time series are stationary over time (static functional connectivity (sFC)). The second category assumes that the time series are non-stationary (dynamic functional connectivity (dFC)). While early work was devoted to analysis of static functional connectivity, increasing evidence suggests that functional connections dynamically transition on a small time scale. In fact, methods sensitive to dFC have highlighted the connections between stable dFC and cognitive performance (Cohen, 2018) as well as higher prediction accuracy of individuals suffering from post traumatic stress disorder (Jin et al., 2017).

1.2.1 Static Functional connectivity Methods

(i) Mass Univariate

Many methods for comparing brain networks across populations are classified as mass univariate tests. For this general approach, linear models are fit to each edge in the network, where the regressors are a disease indicator variable and confounding variables such as age or gender (Nichols and Holmes, 2002). These intuitively simple methods are in widespread use in the clinical investigations of functional connectivity (Athanasίου et al., 2018; Greicius et al., 2007a; Hasenkamp and Barsalou, 2012). Although intuitively simple, statistical issues arise in selecting edges with statistically different weights across the two populations. Post hoc methods which adjust for the large number of tests substantially reduce the power of mass univariate tests to detect meaningful edges (Zalesky, Cocchi, Fornito, Murray and Bullmore, 2012).

Methods such as the network-based statistic (Zalesky et al., 2010) attempt to circumvent the multiple testing problem by seeking the largest connected component of differentially weighted edges and assessing the statistical significance by permutation tests. Although this method is the most widely used network comparison approach, it performs very poorly if edge-wise differences are scattered across the network and do not form a connected component. Other approaches, such as the spatial pairwise

clustering (Hipp et al., 2011), sum of powered score (SPU) and adaptive sum of powered (aSPU) (Pan et al., 2014), attempt to leverage information at the edge level to determine the statistical significance of brain graphs across populations. These methods generally do not suffer from the multiplicity problem inherent in the mass univariate tests, but can be extremely sensitive to the selected tuning parameter.

(ii) Network summary statistics

Network summary statistics are a class of methods to detect differences in the topological structure of the brain network, \mathcal{G} . The primary motivation for these approaches is to circumvent the multiple comparison issue inherent in the edge-level statistics. Further, network topology has been directly related to neural efficiency, robustness, integration capabilities and communication capacity (Van Den Heuvel and Pol, 2010). Importantly, functional connectivity graphs have been shown to exhibit pronounced topological properties characteristic of small-world networks. As such, graphs can be summarized via a small set of network metrics, such as the average path length, clustering coefficient, or assortativity (see Wang et al. (2015a) for functions to calculate the metrics). With these metrics, one could compare the network topology across populations to detect difference in network architecture attributed to disease. Lord et al. (2012) found that the community structure, in addition to several other network metrics, accurately predicted whether an individual was healthy or depressed. Many network analyses study find altered local and global network topologies in other studies of depression (Leistedt et al., 2009) as well as in meditators (Ginestet et al., 2014). Simpson et al. (2013a) propose a general suite of permutation tests for assessing global differences in networks as well as differences in the nodal degree distribution.

The exponential random graph model (ERGM) is a model-based approach for estimating the conditional effect of a set of network metrics on an observed network structure. Simpson et al. (2011, 2012) investigate these models and demonstrate their effectiveness in representing brain graphs. Furthermore, they propose post hoc

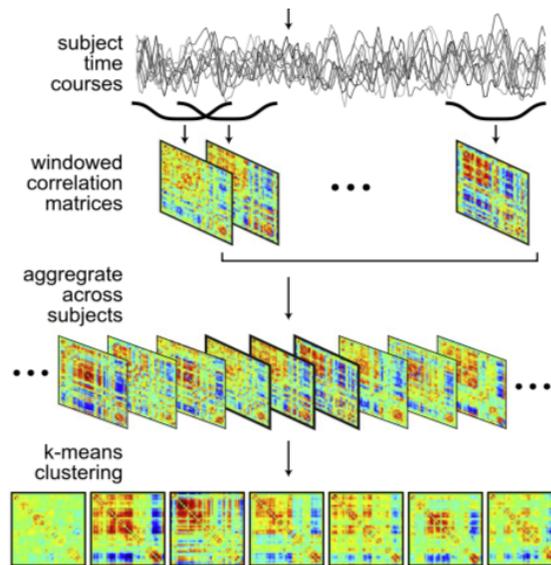
inference for the network features. The primary appeal of such an approach is that it jointly estimates the effect of graph-based features on the observed network. In spite of the advantages, the ERGM has several limitations. It is computationally demanding and prone to degenerate solutions if the network metrics are not carefully selected based on the application.

1.2.2 Dynamic functional connectivity Methods

Dynamic function connectivity methods are still in their infancy. Over the last five years, the dominant method for calculating dynamic functional connectivity is sliding window correlations (SWC). In SWC, functional networks are calculated over a window of fixed or varying width. These correlation networks provide a glimpse of fluctuations in functional connectivity throughout a scanning session. Although there are suggestions for selecting the window lengths (Leonardi and Van De Ville, 2015; Zalesky and Breakspear, 2015), this parameter has been shown to have substantial effects on the network estimates (Cohen, 2018). From the sliding window estimates, centroids are calculated using k-means clustering to detect the basis brain states common to all subjects. Many studies have show that disease classification becomes clear based upon the frequency in which a subject utilizes the “brain states” (Allen et al., 2014a).

Recently, more complex models have been used to detect dynamic changes in functional connectivity. Hidden Markov models (Jurafsky and Martin, 2014) are being implemented as a model-based approach to estimating the brain states shared across subjects as well as the probabilities of state transitions (Andersen et al., 2018; Vidau-re et al., 2017). The simplest implementation assumes that the observed regional time-series are generated by multivariate Gaussian distributions, where the covariance matrix of each Gaussian component represents a brain state. Unfortunately, these models do not scale well as the number of brain regions increases. An emerging

Figure 1.3: Estimation of sliding windows correlations.



(Hutchison et al., 2013)

method for estimating dFC are based on dictionary learning techniques, which were originally developed in the compressed sensing literature. Early explorations indicate that dictionary learners can successfully extract basis brain states from observed data (Li et al., 2014; Yaesoubi et al., 2018). Furthermore, the penalized optimization algorithms can detect the states contributing to the observed connectivity pattern at each time point throughout the scanning session.

1.2.3 Challenges in functional connectivity

Although network-based methods have led to important discoveries about cortical activity, various statistical and computational challenges persist. A large issue with many static connectivity analysis is how to threshold the edge set to detect meaningful connections. Network density has a non-trivial impact on many network metrics which can lead to ineffective group comparisons. As previously discussed, there are many different imaging modalities, each capturing different aspects of functional ac-

tivation in the brain. Hutchison et al. (2013) suggest that EEG, which has high temporal resolution, can be used to interrogate resting-state fMRI. Further, it is well known that brain anatomy constrains functional connectivity. Methods that flexibly incorporate complementary information are needed in to improve current functional connectivity methodologies. Dynamic functional connectivity also presents several challenges which will require advanced techniques. As Cohen (2018) suggest, dFC methods must disentangle neuronal activity from noise attributed to physiological, respiratory, and cardiovascular processes. Additionally, methods must be able to account for subject specific arousal, which has been suggested to reduce state transitions (Haimovici et al., 2017). Preprocessing can minimize most data artifacts, but methods must be developed which can adjust for such occurrences. Additionally, individual specific characteristics, such as intelligence, have been shown to modify co-activity patterns. In total, these effects may obscure dFC estimates, increasing the difficulty of distinguishing disease-related effects on functional connectivity. Available methods do not account for such effects. In the following sections, we present methods that overcome these limitations and apply them to populations exhibiting strong differences compared to healthy individuals. Fallani et al. (2014) and (Hutchison et al., 2013) provide excellent overviews of the challenges of static and dynamic functional connectivity.

1.3 Motivating Data

In recent years, network science tools have been increasingly applied in studies of brain function in health and disease. A large body of literature investigating diseases such as schizophrenia, autism, epilepsy, and depression has shown that most are identified by hyper- and hypo- connectivity between brain regions. Despite stark differences in how and where the diseases manifest in the brain, many have been

generically classified as “dysconnectivity syndromes” in the functional connectivity literature (Catani and Ffytche, 2005). In two chapters, we apply our methods to disorders where functional connectivity is a well-known and strong biomarker for disease. However, in chapter three we focus upon mentally healthy individuals and exclude populations in which disease potentially confounds the relationship between brain structure and function.

1.3.1 Major Depressive Disorder dataset

Major depressive disorder (MDD) is a severe mental condition and is the leading cause of disability in Americans aged 15-44. Although it is prevalent in every subpopulation, it disproportionately affects women and individuals 18-29 years old. MDD is often co-morbid with other serious conditions, such as anxiety disorders, obsessive-compulsive disorder, and psychotic disorder (Papan Thaipisuttikul et al., 2014), and is the second leading cause of death in 15-29 year olds in the United States. Despite the pervasiveness of the disorder, it is difficult to diagnose (Agrawal and Rickards, 2011), in part because the symptoms are similar to other neurological disorders. Various studies have indicated that major depressive disorder is typically associated with increased activity in the default mode network and induces cognitive deficits resulting from suppressed activity in the frontoparietal task control network (Brzezicka, 2013).

We acquire resting-state fMRI scans from twenty MDD subjects and nineteen healthy subjects from the Mayberg Lab in the Emory University Department of Psychiatry. MDD patients are on average 45.8 years old (SD: 9.6 years) and fifty percent male. The age and gender matched healthy participants are 47% male and 43 years old (SD: 8.9 years). MDD patients were evaluated with the 17-item, clinician-rated Hamilton Rating Scale for Depression and had an average score of 19, which corresponds to severe depression (Brown et al., 2008). Study participants were instructed to lie motionless and thoughtless in the scanner for the duration of the data acquis-

tion. Data were acquired on a 3T Tim Trio MRI scanner with a twelve-channel head array coil. fMRI images were captured with a z-saga sequence to minimize artifacts in the medial PFC and OFC due to sinus cavities (Heberlein and Hu, 2004). Z-saga images were acquired interleaved at 3.4x3.4x4 mm resolution in 30 4-mm thick axial slices with the parameters FOV=220x220 mm, TR=2920 ms, TE=30 ms for a total of 150 acquisitions and total duration 7.3 min. Several standard preprocessing steps were applied to the rs-fMRI data, including despiking, slice timing correction, motion correction, registration to MNI 2mm standard space, normalization to percent signal change, removal of linear trend, regressing out CSF, WM, and 6 movement parameters, bandpass filtering (0.009 to 0.08), and spatial smoothing with a 6mm FWHM Gaussian kernel.

1.3.2 Philadelphia Neurodevelopmental Cohort dataset

The Philadelphia Neurodevelopmental Cohort is a large study funded by the NIMH to investigate brain and behavior interaction with genetics. 9500 individuals in the greater Philadelphia area, ages 8-21, were neuropsychiatrically evaluated via structured interviews and completed the Computerized Neurocognitive Battery. A smaller subset of 1445 participants received imaging that includes functional imaging tasks of working memory and emotion identification, resting state imaging of functional activity, T1-weighted structural images, diffusion tensor imaging, and perfusion neuroimaging using arterial spin labeling. Our analysis relies upon the resting state fMRI and DTI images acquired on sixty-nine healthy, right handed children exhibiting normal cognitive development.

Resting-state fMRI scans were acquired on a single-shot, interleaved multi-slice, gradient-echo, echo planar imaging (GE-EPI) sequence (Satterthwaite, Elliott, Ruparel, Loughead, Prabhakaran, Calkins, Hopson, Jackson, Keefe, Riley et al., 2014). Nominal voxel size is 3x3x3mm with full brain coverage achieved with the following

parameters: TR/TE=3000/32 ms, flip=90°, FOV=200 × 220 mm, matrix= 64 × 64, 46 slices, slice thickness/gap=3 mm/0 mm for a total of 6.2 minutes. Participants were instructed to remain awake, motionless, and fixated on a crosshair throughout the duration of the data acquisition. Several standard preprocessing steps were applied to the rs-fMRI data, including despiking, slice timing correction, motion correction, registration to MNI 2mm standard space, normalization to percent signal change, removal of linear trend, regressing out CSF, WM, and 6 movement parameters, bandpass filtering (0.009 to 0.08), and spatial smoothing with a 6mm FWHM Gaussian kernel. Subsequent voxel level data is aggregated into 90 regions of interest (ROI) based on the Automated Anatomical Labelling atlas (Tzourio-Mazoyer et al., 2002). For each ROI, the average time series of all constituent voxels represents the region’s temporal BOLD signal.

Diffusion weighted images permit us to localize and orient white matter fiber bundles via the diffusion of water in the brain. Images were acquired on a twice-refocused spin-echo (TRSE) single-shot EPI sequence for a total of 64 diffusion-weighted directions with $b=1000$ s/mm² and 7 scans with $b=0$ s/mm² (Satterthwaite, Elliott, Ruparel, Loughead, Prabhakaran, Calkins, Hopson, Jackson, Keefe, Riley et al., 2014). Acquisition parameters were TR/TE=8100/82ms, matrix=128×128, FOV=240mm, slice thickness=2mm, GRAPPA factor=3. Due to gradient induced vibrations disturbing image quality, DWI images were acquired in two imaging runs to reduce the continuous duration in which subjects tolerate the scan. Standard pre-processing procedures, such as eddy current correction and bias-field correction are applied to the diffusion weighted data. Subsequently, we use the FSL functions `bedpostx` and `probtracx2` to estimate the distribution of fiber tensors at each voxel and the count of white matter fibers tracts connecting all pairs of brain regions, respectively. In order to obtain the SC scores, we compute $p_{jk} = \text{average}\{N_{jk}/N_{j*}, N_{kj}/N_{k*}\} \in (0, 1)$, where N_{jk} equals the number of permissible tracts initiated at region j that pass through

region k and N_{j*} refers to the total number of permissible tracts initiated at region j . Fiber tracks passing through gray matter or cerebrospinal fluid are discarded. These SC scores can be interpreted as the probability of structural connectivity between regions j and k , which we often refer to as the strength of SC.

1.3.3 Grady Trauma Project dataset

All images were acquired on a Siemens Tim Trio 3T scanner (z-saga interleaved echo-planar sequence). fMRI images were captured with a z-saga interleaved sequence. In total, 150 volumes are acquired with scanning parameters TR 2950 ms, TE 67.04 ms, flip angle 90°, FOV 220x220 mm, matrix 64 x 64, and voxel resolution 3.4375mm x 3.4375mm x 4mm. Raw images were preprocessed using the script released with the 1000 Functional Connectomes Project. Skull stripping was performed on the T1 image to remove extra-cranial tissues followed by removal of the first four volumes to stabilize the signal. The anatomical image was registered to the 8th volume of the rs-fMRI image sequence and normalized to the MNI standard brain space. The functional scans were registered to standard space with these normalization parameters and subsequently subsequently smoothed using a 6 mm FWHM Gaussian kernel. The functional images were further cleaned by regressing out nuisance signals including motion artifacts, global effects, white matter and cerebrospinal fluid signals. Finally, the functional time series data was band-pass filtered—.01 Hz to .1 Hz—to retain frequencies relevant to rs-fMRI.

We extract 266 regions from the preprocessed rs-fMRI data of which 264 are defined in the Power brain atlas (Power et al., 2011a). Although the Power atlas provides coverage of cortical and subcortical brain regions, it does not include the amygdala, which is critical in emotion processing and highly implicated in PTSD. Consequently, we include two subdivisions of the amygdala—basolateral (BLA) and centromedial (CEN) —as defined in the California Institute of Technology (CIT68)

probabilistic high-resolution in vivo atlas of the human amygdala (Tyszka and Pauli, 2016). Each brain region is a 10mm diameter sphere containing approximately 81 voxels. Regional time courses are the average of the constituent voxels' time course.

The Grady PTSD dataset consists of 99 subjects. We remove 41 subjects exhibiting movement in $\geq 55\%$ of the scans. All subsequent analyses are carried out on 42 healthy controls and 14 PTSD patients. PTSD status was determined by the modified PTSD Symptom Scale (PSS), which has high internal consistency and strong validity for diagnosing PTSD consistent with DSM-IV (Kilaru et al., 2016). All PTSD patients exhibit at least one symptom in the re-experiencing cluster, three symptoms in the avoidance cluster, and two symptoms in the hyperarousal cluster. Although all subjects have experienced at least one traumatic event, PTSD patients exhibit PTSD symptoms at the time of study. The study protocols were approved by the Institutional Review Board of Emory University School of Medicine and the Grady Health Systems Research Oversight Committee.

1.4 Proposed Research

This dissertation proposes a suite of network-based methods for elucidating complex functional activity in the brain and, subsequently, relating that to clinical outcomes. Methods presented in chapter two and four provide distinctly different approaches for addressing a fundamental, yet fleeting question in neuroimaging statistics: where and by which mechanisms do mental diseases produce irregular functional co-activity. Yet another fundamental question that is of particular interest in the neuroimaging community is the relationship between function and structure in the brain. Although various studies agree that brain anatomy constrains function (Honey et al., 2009, 2010; Messé et al., 2014), few methods exist that incorporate brain structure into the estimation of functional connectivity. We propose such an approach in chap-

ter three.

1.4.1 Topic 1: Difference degree test for comparing brain networks

In chapter two, we develop a test to detect brain regions exhibiting disease-related irregularities in local functional connectivity. Our method relies upon an appropriate set of null, random networks that serve as random realizations of the observed difference network. In our approach, the nulls permit isolation of true network topology from that induced by correlation transitivity. Extensive simulation studies highlight the advantages of our method over existing approaches. The test also allows us to detect brain regions implicated in major depressive disorder in the real data application.

1.4.2 Topic 2: Anatomically informed estimation of brain functional networks

In chapter three, we propose a structurally-informed Gaussian graphical model (siGGM) that incorporates the strength of structural connectivity between distinct brain regions into the procedure for estimating functional co-activity. Our hierarchical model flexibly incorporates structural connectivity information and produces a maximum a posteriori (MAP) solution of the partial correlations. The siGGM is applied to the publicly available Philadelphia Neurodevelopmental Cohort study, whereby we leverage available rs-fMRI and DTI data on sixty-nine cognitively healthy, right handed patients.

1.4.3 Topic 3: Semi-parametric Bayesian hierarchical dictionary learning

In chapter four, we propose a Bayesian hierarchical dictionary learning method that simultaneously estimates a set of brain network common to a population and classifies patients into subpopulations based on utilization of the states. Our model flexibly incorporates covariates that can confound dynamic functional connectivity estimates. We apply the method to a rs-fMRI acquired on healthy and PTSD patients.

Chapter 2

Comparison of nodal differential degree centrality

2.1 Introduction

In recent years, graph theoretical tools have become increasingly important in the analysis of brain imaging data. In particular, evaluations of the associations between spatially distinct regions has led to valuable insights into the brain’s organization in health and disease. Functional connectivity (FC), which measures the coherence between neurophysiological time series (Friston, 1994), has been extremely valuable in identifying disease-induced modifications to cortical and subcortical communication. In fact, altered cortical activity has been observed in major depressive disorder (MDD) (Craddock et al., 2009; Drysdale et al., 2017), Alzheimer’s disease (Stam et al., 2006), and schizophrenia (Liu et al., 2008; Rubinov et al., 2009). While Pearson correlation is a widely used FC measure, alternate association metrics such as partial correlations (Wang et al., 2016), mutual information (Salvador et al., 2005), and coherence (Bassett et al., 2011) are finding favor. Brain networks have become particularly important since the FC measures offer different perspectives on co-activation between brain regions, and many studies agree that psychiatric disorders and neurodegenerative diseases manifest as disruptions in local and global functional connectivity (Pandya et al., 2012).

Many methods exist for comparing brain networks and connectivity patterns across populations. The earliest approach tests for group differences at each edge in the network (Nichols and Holmes, 2002). For a network with N regions, this requires multiple testing corrections since $N(N-1)/2$ unique edges must be assessed. Unfortunately, controlling the family-wise error rate or false discovery rate leads to a reduction in power to detect group differences at the edge level. The sum of powered score (SPU) and adaptive sum of powered score (aSPU) tests (Pan et al., 2014) leverage edge level differences to assess overall deviation in the networks (Kim et al., 2015). While they can lead to high powered tests, there are practical difficulties in selecting the optimal tuning parameter. Furthermore, the tests do not specifically identify

edges, regions, or structures contributing to overall network differences, which leads to a loss in interpretability in the brain functional aspect.

Other approaches assume differences in brain connectivity result in large deviations in the network’s topology. The network-based statistic (NBS) (Zalesky et al., 2010) and the parsimonious differential brain connectivity network detection method (PARD) (Chen et al., 2015) are useful for identifying collections of differentially weighted edges (DWEs) forming interconnected subcomponents, but have limited exploratory value. A key assumption in these methods is that altered edges form connected subnetworks (Kim et al., 2015). The NBS is severely underpowered to detect differences in the networks if this assumption is violated (Zalesky et al., 2010). Furthermore, the NBS does not provide a principled approach for thresholding significant DWEs. As shown in Kim et al. (2014a), the NBS performance is reported across a grid of thresholds. In simulations, one can heuristically tune the threshold to produce the best performance. However, threshold selection becomes substantially more difficult in real data applications. In simulations and real data applications, we find that the NBS (extent) is unable to detect connected components spanning the difference network across a wide range of thresholds. For these reasons, we exclude the NBS results. We note that the NBS is a complementary method to the DDT. If the connected component assumption is valid, we expect the NBS to exceed the performance of the DDT. However, when the DWEs form small isolated clusters, we expect the DDT to perform much better.

Other methods (Rudie et al., 2013; Wang et al., 2015a) have focused on comparing graph metrics across networks, using two sample t-tests to test for differences. Unfortunately, these tests may often be underpowered to detect group differences (Kim et al., 2014a), and there are doubts on the suitability of two-sample t-tests to compare some network metrics (Fornito et al., 2010; Hayasaka and Laurienti, 2010). Alternatively, nonparametric approaches utilize permutation tests (Zalesky et al., 2010) or

generate random networks (Bullmore et al., 1999) in order to construct distributions for network metrics of interest under the null hypothesis and then use these reference distributions to evaluate the significance of the observed network features. However, generating the appropriate null network is non-trivial. Existing approaches attempt to randomly rewire edges while preserving the degree distribution and the clustering coefficient (Bansal et al., 2009; Maslov and Sneppen, 2002; Volz, 2004). Unfortunately, the network generation schemes are sensitive to the desired network measure (see Fornito et al. (2013) for an overview) and may not provide a complete picture of the network differences reflected by alternate summary measures.

In this paper, we propose a Differential Degree Test (DDT) to identify brain regions that demonstrate significant between-group differences in neural connections. Despite the naming convention, the DDT is not a comparison of the nodal degree across populations. Instead, it assesses the count of DWEs incident to each node. We incorporate the former into our simulations to verify that it leads to markedly inferior results compared to the DDT. The proposed testing approach facilitates the comparison of brain networks across populations while bypassing the limitations of current methods. Our method is based on the difference network, in which the edges represent the statistical significance of between-group differences. The observed difference network is compared to a set of null networks that are carefully constructed, via the Hirschberger-Qi-Steuer (HQS) algorithm, to maintain both the first and second moment characteristics of the observed difference network. Preservation of these two moments allows the null network to retain the nuisance topology present in the observed difference network, ignoring all intrinsic structures. Such a null network reflects the topology induced by the correlation measure encoding the brain connectivity associated with the observed network. In contrast, naively generated random networks, i.e. random edge rewiring and random edge sampling, can produce nulls which do not reflect the intrinsic nor nuisance topologies present in the observed network. This

will likely lead to the detection of spurious, meaningless features. Clearly, such random null networks potentially provide inappropriate assortments of random network configurations nor are they guaranteed to replicate the relevant structures embedded in the observed network. As such, traditional null network generation schemes can lead to falsely identified edges and artificial structures. Our approach overcomes these shortcomings and identifies intrinsic structures in the observed network by explicitly accounting for the nuisance topology embedded in the graph.

We adapt the null network generation scheme of Hirschberger et al. (2007) to replicate networks retaining only the nuisance structure in the difference network, which enables the separation of the network’s true topology from the nuisance topology. We note that while Zalesky, Fornito and Bullmore (2012) utilize HQS to examine the impact of nuisance topology on local network features, our approach is the first to use the HQS algorithm for the assessment of network differentiation across populations. We further propose an adaptive thresholding procedure to identify significant DWEs by comparing against the generated null difference networks using the HQS algorithm. Based on the thresholded difference adjacency matrix, the proposed Differential Degree Test (DDT) identifies nodes or brain regions that demonstrate a significantly higher number of DWEs as compared to the null distribution.

Through extensive simulations, we illustrate that the proposed method has greater power to detect differentially connected nodes across networks compared to standard multiple testing procedures, while also maintaining reasonable control over false positives. Furthermore, the adaptive threshold selection procedure leads to increased power to detect DWEs across the network as compared to Bonferroni and false discovery rate (FDR) correction procedures. Additionally, the adaptive threshold approach under the proposed method can automatically adapt to different network settings and hence is more generalizable compared to ‘hard’ thresholding approaches assuming a fixed threshold. Finally, we apply the proposed approach for the analysis of a ma-

for depressive disorder (MDD) dataset, which leads to meaningful findings regarding disrupted brain connectivity due to MDD.

The rest of the paper is organized as follows. In sections 2.2, we discuss the construction of null difference networks and the proposed DDT procedure. We present simulation results in Section 2.3, MDD data analyses in Section 2.4, and conclude with a discussion of the findings in Section 2.5.

2.2 Method

In the following, we discuss the construction of subject-specific functional brain networks, formulation of the difference network, and the details of the proposed DDT. Our analysis investigates connectivity disruptions across the entire brain, but is also amenable to hypothesis-driven investigations of functional connectivity containing a number of pre-selected brain regions.

2.2.1 Brain network construction

In network analysis of neuroimaging data, the brain can be represented as a graph defined by a finite set of nodes (brain regions) and edges showing the statistical association between pairs of nodes. For N nodes, the network is represented as a symmetric $N \times N$ connectivity matrix, \mathbf{G} , which can be thresholded to obtain the adjacency matrix \mathbf{A} , representing the edge set of the network. For selection of the node system, the naive approach is to treat each voxel as a putative region of interest. This approach results in an extremely high-dimensional connectivity matrix that not only poses challenges for subsequent analyses, but also tends to be unreliable and noisy. A more common approach is to define nodes based on anatomically defined brain structures, e.g. Automated Anatomical Labeling (AAL) atlas (Tzourio-Mazoyer et al., 2002) and Harvard-Oxford atlases (Fischl et al., 2004; Frazier et al., 2005).

When analyzing brain functional networks, it is suggested to parcellate the brain into putative functional areas based on clusters of voxels exhibiting similar signals in resting-state functional imaging data (Craddock et al., 2012). Some widely used examples of functionally defined node systems are the Power 264 node system (Power et al., 2011a), Yeo (Yeo et al., 2011), and Gordon (Gordon et al., 2014) atlases, among others.

For brain network based on functional magnetic resonance imaging (fMRI), the edges represents the coherence in the temporal dynamics between the blood oxygen-level dependent (BOLD) signal between node pairs. In this paper, we utilize undirected measures of connectivity such as Pearson and partial correlation, where Pearson correlation measures the marginal association between two regions and partial correlation measures their association conditioned on all other regions in the network. Given the heavy debate on the merits and disadvantages of each correlation measure in brain network analysis (Kim et al., 2015; Liang et al., 2012), we investigate both and compare the findings.

The resulting network, \mathbf{G} , is a weighted graph representing undirected statistical associations between all pairs of nodes. Often, a thresholding procedure is applied to produce a binary adjacency matrix, \mathbf{A} , where a value of 1 in the $(i, j)^{th}$ entry indicates a connection between the respective regions. This network formulation is particularly advantageous as it simplifies calculations of graph metrics and leads to intuitive metric definitions (see Bullmore and Bassett (2011); Rubinov and Sporns (2010a) for more details).

Since we are interested in between-group differences in functional networks, we consider a difference network which is defined on the same node system as the functional network but the edges represent the strength of between-group differences in the functional connections. Details of the difference network construction are presented in the following section. We focus on the number of thresholded edges incident to each

region in the difference network, which we call the differential degree. Similar to the interpretation of nodal degree in connectivity matrices, we focus upon this metric as it suggests regions contributing to local differences in the network architecture across diseases or conditions. We believe that a brain region incident to a large number of differentially weighted edges (DWEs) is potentially responsible for overall differences in brain network topology, without being sensitive to any particular network summary measure commonly used to capture connectome differences.

2.2.2 Differential Degree Test

In this section, we present the DDT method for identifying brain nodes whose connections demonstrate significant differences between group.

2.2.2.1 Difference Network Construction

Suppose we are comparing networks between two groups with n_r subjects in group r ($r = 1, 2$). Denote $\mathbf{G}^{k_r} = \{g_{r,ij}^{k_r}\}$ ($k_r = 1, \dots, n_r$) as the estimated brain connectivity matrices for the k_r th subject in the r th group ($r = 1, 2$) and $g_{r,ij}^{k_r}$ denotes the connectivity measure (such as the Pearson or partial correlation) between nodes i and j ($i, j \in \mathcal{N} = \{1, \dots, N\}$) for the k_r -th subject in the r th group. The first step of DDT is to construct a $N \times N$ difference network $\mathbf{D} = \{d_{ij} : i, j \in \mathcal{N}\}$, where d_{ij} represents the statistical significance of population-level differences in the connection strength between node i and j , i.e.

$$d_{ij} = 1 - p(\{g_{1,ij}\}, \{g_{2,ij}\}) \in [0, 1), \quad (2.1)$$

where $p(\{g_{1,ij}\}, \{g_{2,ij}\})$ is the p-value of a between-group difference test based on the estimated connectivity measures at edge (i, j) across subjects in the two groups. For example, one can obtain the p-value by applying two-sample t test to $\{g_{r,ij}^{k_1}\}$ and

$\{g_{r,ij}^{k_2}\}$. We will provide more detailed discussion on how to derive the p-values from various types of between-group tests in section 2.2.2.2. From (2.1), each element in the difference network d_{ij} serves as our measure of the difference of the edge connectivity $g_{r,ij}$ between the two groups, with larger values (i.e. smaller p-values) corresponding to larger group differences at the (i, j) th edge, and vice-versa. Note that $\mathbf{D} = \{d_{ij}\}$ is a symmetric matrix where $\forall i, j \in \mathcal{N} \ d_{ij} = d_{ji}$ and $d_{ij} = 0$ for $i = j$ given that we are not interested in the diagonal elements.

From the difference network $\mathbf{D} = \{d_{ij}\}$, we can derive the difference adjacency matrix $\mathbf{A} = \{a_{ij}\}$ where a_{ij} represent the presence of group differences in the connection between nodes i and j , i.e.

$$a_{ij} = I(d_{ij} > \tau), \quad (2.2)$$

where τ is a threshold for selecting edges which are differentially weighted. When d_{ij} exceeds the threshold τ , or equivalently the p-value for the group test is smaller than $1 - \tau$, we obtain $a_{ij} = 1$ indicating the presence of group difference at the edge (i, j) . Otherwise, $a_{ij} = 0$ represents no group difference at the edge (i, j) . In the following section, we will present a data-driven adaptive threshold selection method for finding τ .

Based on the difference adjacency matrix \mathbf{A} , we define the following differential degree measure for the i th node ($i = 1, \dots, N$),

$$d_i = \sum_{j \in \mathcal{N}, j \neq i} a_{ij} \quad (2.3)$$

The differential degree measure d_i represents the number of connections to node i that demonstrate significant differences between the two groups as captured by edge-wise p-values without multiplicity adjustments. In subsequent steps of the DDT, d_i will be used as the test statistic for investigating node i 's contribution to disrupted

communication in the brain. While the difference network provides edge-level information on between group differences, it is widely accepted that cognitive deficits in mental diseases are demarcated by disruptions in systems (Catani and Ffytche, 2005). Thus, collections of connected DWEs are more consistent with the systemwide disruption paradigm than evaluation of individual DWEs. The DWEs incident to each node form a locally connected component and indicate that irregular activity at the node of interest contributes to differentiated co-activation with adjacent regions. Investigation at the nodal level not only has biological justification, but also substantially improves the multiple testing problem. The number of statistical tests scales linearly with the network’s size rather than quadratically at the edge level. The notion of disruptions in sub-systems has also been used in previous work to mitigate the multiplicity problem common to network comparisons (Zalesky et al., 2010).

2.2.2.2 Deriving p-value from between-group tests

The p-value used to define the difference network in (2.1) can be derived based on various between-group testing procedures. The p-values fall into two categories: model-free and model-based. The model-free p-values are derived based on parametric or nonparametric tests between the two groups of subjects without accounting for the subjects’ biological or clinical characteristics. The common choices of such tests include the two-sample t test, the nonparametric Wilcoxon rank sum test or the permutation test. The model-based p-values are derived from regression models where the subject-specific connectivity measure (or some transformation) is modeled in terms of group membership and other relevant factors such as age and gender that may affect the brain connectivity. These p-values for between-group differences can then be derived based on the test of the parameter in the model associated with the group covariate. This model-based p-value reflects the degree of group differences while controlling for potential confounding effects. In many neuroimaging studies,

subjects' group memberships are not based on randomization but rather based on observed characteristics. In this case, the distribution of subjects' demographic and clinical variables tend to be unbalanced between the groups and there often exist some potential confounding factors in between group comparisons (Satterwaite et al., 2014a). For such studies, it may not be the case that the model-based p-values more accurately reflect group-induced variation in functional connectivity as compared with model-free p-values.

We note that when computing the difference network in (2.1), the proposed approach does not apply a multiple testing correction to the edge-wise between-group test p-values. Such multiplicity adjustment often reduces the power to detect DWEs. Additionally, since our goal is to detect differentially expressed nodes in the brain network, a multiplicity adjustment on the edge-wise tests is not crucial, provided the falsely identified DWEs are more or less uniformly distributed across the nodes without systematic differences. In such a case, the threshold τ in (2.2), which is chosen using an appropriately constructed null distribution as in Section 2.2.4, automatically adjusts for falsely identified DWEs occurring across nodes. Indeed, extensive simulation studies illustrated that the proposed method is able to control false positives at a nominal value.

2.2.2.3 Null distribution generation

After constructing the difference network $\mathbf{D} = \{d_{ij}\}$ and deriving the differential degree measure, d_i , for each node, the next step in the DDT procedure is to conduct a statistical test to evaluate whether there is significant group difference in the connections to the node. As a standard strategy in hypothesis testing, we will evaluate the test statistic, d_i , with respect to its null distribution under the hypothesis that there are no between-group differences. For this purpose, we first derive the null distribution by generating difference networks under the null hypothesis.

We present a procedure for generating null difference networks that maintain some of the fundamental characteristics of the observed difference networks but has a random pattern of between-group differences which is expected under the null hypothesis. Since the elements in the difference network lie within a restricted range, i.e. $(0, 1)$, we first apply a logit transformation, i.e.

$$\bar{\mathbf{D}} = \{\bar{d}_{ij} : \bar{d}_{ij} = \text{logit}(d_{ij}) \in (-\infty, \infty), i < j, i, j \in N\}. \quad (2.4)$$

We define the first and second moment characteristics for the observed difference network as follows,

$$\bar{e} = E[\bar{d}_{ij}] \text{ and } \bar{v} = \text{Var}[\bar{d}_{ij}] \text{ for } i < j, \quad e = E[\bar{d}_{ij}] \text{ for } i=j$$

where \bar{e} represents the mean of the off-diagonal elements, e represents the mean of the diagonal element and \bar{v} is the variance of the off-diagonal elements.

In the following, we present a procedure for generating a null difference network $\mathbf{C} \in \mathbb{R}^{N \times N}$ whose first and second moment characteristics matches that of the observed difference network, and preserves its true topology. Motivated by the Hirshberger Qi-Steuer (HQS) algorithm (Hirschberger et al., 2007), we propose to generate \mathbf{C} based on the multiplication of a random matrix and its conjugate transpose

$$\mathbf{C} = \mathbf{L} * \mathbf{L}^T, \quad (2.5)$$

where $\mathbf{L} \in \mathbb{R}^{N \times m}$. Based on the formulation of (Hirschberger et al., 2007), we generate $l_{ij} \sim N(\mu, \sigma^2)$ where $\mu = \sqrt{\frac{\bar{e}}{m}}$ and $\sigma^2 = -\mu^2 + \sqrt{\mu^4 + \frac{\bar{v}}{m}}$ and $m = \min\{2, \lfloor \frac{e^2 - \bar{e}^2}{\bar{v}} \rfloor\}$ where $\lfloor \cdot \rfloor$ is the floor function. Based on this specification, we can show that

$$E[c_{ij}] = \bar{e}, \text{Var}[c_{ij}] = \bar{v} \text{ and } E[c_{ii}] = e,$$

Please see equations A.1 for details. The generated null difference network $\bar{\mathbf{C}}$ main-

tains the first and second moment characteristics of the observed difference network $\bar{\mathbf{D}}$. Finally, we transform $\bar{\mathbf{C}}$ through the inverse logit function to obtain a null difference network \mathbf{C} such that $c_{ij} \in (0, 1)$.

The proposed generation procedure has several appealing features. First, it is a very fast algorithm for generating null networks. Second, the generated null difference network, $\bar{\mathbf{C}}$, preserves the first and second moment characteristics of the observed difference network $\bar{\mathbf{D}}$. An important advantage in maintaining these fundamental properties of the observed network is that it will help make the generated null network a meaningful reference for comparison with the observed network. For example, to perform meaningful comparison of the connectivity structure between two networks, a critical condition is that the two networks must have similar number of edges (Fallani et al., 2014). This condition would be violated if there exists a significant difference in the average connectivity measure between the two networks in the sense that the network with higher average connectivity is associated with larger number of edges. By generating null networks with the same first and second moment as the observed network, the proposed procedure makes sure the comparison between the observed network against the null networks would not be confounded by their differences in the fundamental characteristics. More importantly, replication of the first and second moments allows the null networks to preserve the nuisance topology of the observed difference network while annihilating intrinsic group structures of the observed network. As discussed in Zalesky, Fornito and Bullmore (2012), benchmarking against such null networks permits identification of the intrinsic topology in the observed network.

2.2.2.4 An adaptive threshold selection method

Recall that after obtaining the difference network $\bar{\mathbf{D}} = \{\bar{d}_{ij}\}$, we need to threshold it to derive the difference adjacency matrix $\mathbf{A} = \{\bar{a}_{ij}\}$. If $\bar{d}_{ij} > \gamma$, $\bar{a}_{ij} = 1$ indicating

the presence of a group difference at the edge (i, j) where $\gamma = \text{logit}(\tau)$. Otherwise, $\bar{a}_{ij} = 0$ represents no group difference at the edge (i, j) .

In the existing between-group network tests, the threshold value is typically selected by a multiple comparison method that controls the family-wise error rate or the false discovery rate. Others select a pre-specified cutoff or grid over a range of cutoffs (Zalesky et al., 2010). We propose to adaptively select the threshold based on the distribution of the between-group test statistic. Specifically, the \tilde{c}_{ij} are independent and identical samples from the mixture distribution, $H(\cdot)$,

$$H(\tilde{c}_{ij}) = \frac{2\sigma^2}{4}T - \frac{2\sigma^2}{4}Q, \quad (2.6)$$

where T and Q are non-central χ^2 and central χ^2 random variables, respectively. Each variable in the mixture distribution depends only on the mean and variance of the observed data. We propose two ways to select the threshold, γ as the 95th quantile: (1) aDDT which uses the theoretical critical value based on the parametric mixture distribution in (2.6), and (2) eDDT which uses the empirical critical value based on the empirical distribution. The numerical advantages and disadvantages of each of the two thresholding methods will be addressed in the simulation studies. Since the null difference network, $\bar{\mathbf{C}}$, is generated in a way that it matches the first and second moments of the observed difference network, $\bar{\mathbf{D}}$, the selected threshold value γ will automatically adapt to the properties of the observed difference network. Compared to hard thresholding approaches which use a fixed cut-off value, our threshold selection method can potentially provide an adaptive and general approach for choosing suitable threshold values for different studies. Once the threshold value γ is computed as above, one can apply it to the generated null difference networks $\bar{\mathbf{C}}$ to obtain difference adjacency matrices $\bar{\mathbf{A}} = \{\bar{a}_{ij}\}$ such that $\bar{a}_{ij} = 1$ if $\bar{c}_{ij} > \gamma$ and 0 otherwise.

The proposed threshold selection procedure controls the selection of false positive

edges, while circumventing the loss of power inherent in existing multiplicity corrections methods. This is achieved by adaptively selecting the threshold based on the distribution of the elements in the difference network. Similar approaches (Newton et al., 2004; Kundu et al., 2018) have effectively controlled type I error by using the empirical distribution of edge probabilities to select a threshold in order to detect important connections. We do note that this thresholding procedure does not guarantee a universally optimal choice. The selected threshold is only utilized to control the type I error in the weak sense at the edge level and may be suboptimal with respect to other manual or automatic thresholding procedures. However, simulations show that it leads to superior performance in detecting the nodes of interest.

2.2.2.5 The DDT Test

In this section, we present a statistical test for the difference degree measure, d_i , for node i based on the generated null difference networks. Recall that $d_i = \sum_{j \in \mathcal{N}, j \neq i} a_{ij}$ where a_{ij} is a binary variable indicating the presence of group difference at the connection between node i and j . \bar{d}_i essentially is a count variable representing the number of connections out of a total of $N - 1$ connections of node i that show between group difference. Therefore, we can model d_i with a binomial distribution. Under the null, $d_i \sim f^{null} = \text{Binomial}(N - 1, p_i^{null})$ where p_i^{null} is the the expected probability for each connection of node i to demonstrate between group difference under the null hypothesis. We can estimate the null probability p_i^{null} based on the generated null difference networks, that is,

$$\hat{p}_i^{null} = \frac{1}{M(N - 1)} \sum_{m=1}^M \sum_{j \in \mathcal{N}, j \neq i} \bar{a}_{ij}^{(m)},$$

where M is the total number of null networks and $\bar{a}_{ij}^{(m)}$ are elements of m^{th} thresholded null network, $\bar{\mathbf{A}}^{(m)}$. By comparing the observed d_i against the null distribution, we

identify all regions incident to more DWEs than is expected by chance.

In the following algorithm, we summarize the procedure of the DDT.

Algorithm 1 DDT Procedure

- 1: Construct the difference network $\bar{\mathbf{D}} = \{\bar{d}_{ij}\}$ where $\bar{d}_{ij} = \text{logit}(1 - p(\{g_{ij}^{k1}\}, g_{ij}^{k2}))$
 - 2: Obtain the first and second moment of $\bar{\mathbf{D}}$, i.e. $\bar{e} = E(\bar{d}_{ij})$ and $\sigma^2 = \text{Var}(\bar{d}_{ij})$
 - 3: Generate M null Difference Networks $\bar{\mathbf{C}}^m (m = 1, \dots, M)$ based on \bar{e} and σ^2 using the proposed procedure such that $\bar{\mathbf{C}}^m$ preserves the first and second moment characteristics of $\bar{\mathbf{D}}$
 - 4: Apply the adaptive threshold selection method to find γ based on the percentile of the non-central chi-square distribution of the null difference networks. Specifically, aDDT uses parametric percentiles and eDDT uses empirical percentiles in the adaptive thresholding.
 - 5: Apply threshold γ to $\bar{\mathbf{D}}$ to obtain the difference adjacency matrix $\mathbf{A} = \{\bar{a}_{ij}\}$ and the difference degree measure d_i for node $i \in \mathcal{N}$
 - 6: Derive the null distribution for d_i , i.e. f^{null} based on the generated null difference networks. Specifically, $f^{null} = \text{Binomial}(N - 1, p_i^{null})$, where p_i^{null} is obtained from $\{\bar{A}^m\}$ based on equation 2.2.2.5.
 - 7: Compare the observed difference degree measure, d_i , against its null distribution, f^{null} , to evaluate the significance of between-group differences in the connections of node i .
-

2.3 Simulation

We conduct extensive simulation studies to assess the proposed method’s ability to detect regions with significantly different connections between two groups of subjects. Unless otherwise noted, the generated networks contain $N = 35$ nodes, and we consider sample sizes of 20 and 40 for each of the two groups. For the first set of simulations, we consider the case where there is only one node in the network incident to a specified number of DWEs. Without loss of generality, we refer to it as node 1, and assess whether the proposed DDT can accurately identify this node. We consider both DDT methods, i.e. aDDT based on parametric percentiles and eDDT based on empirical percentiles in the adaptive thresholding step. Three network structures are considered in the simulation: (1) random (2) small world (3) hybrid. Random

networks contain edges that are equally likely to be positive or negative for all connections. We generate this structure by sampling edge weights independently from a $N(0, .04)$ distribution, which produces a connectivity matrix with no structural zeros. The small world network retains the cliquishness of the regular lattice and the short path length of the random network. This structure retains small world properties observed in functional and structural brain networks (Bassett and Bullmore, 2006; Hilgetag and Goulas, 2016; Salvador et al., 2005). The hybrid network seeks to fuse the block diagonal structure observed in real brain networks, while maintaining the small world-ness inherent to human brains. The “blocks” correspond to functional modules observed in the brain such as the default mode and visual networks.

In order to evaluate the performance of our method, we simulate data similar to that of Chen et al. (2015) and Zalesky et al. (2010). All subjects share a common base brain network, \mathbf{B} , which is a correlation matrix generated according to the random, small world and hybrid network structure. We perturb the edge weights in \mathbf{B} to induce subject-level correlation network while controlling the distribution of DWEs across the populations. For subjects $i_1 = 1, \dots, n_1$ and $i_2 = 1, \dots, n_2$ in the two groups, we generate the subject-level networks, \mathbf{H}_{i_1} and \mathbf{H}_{i_2} , as follows: for n_1 subjects in group 1, $\mathbf{H}_{i_1} = \mathbf{B} + \mathbf{W}_{i_1}$, where $\mathbf{W}_{i_1} \in R^{N \times N}$, $w_{ij,i_1} \sim N(0, .02)$ for $1 \leq i < j \leq N = 35$ and $w_{ij,i_1} = 0 \forall i = j$; for n_2 subjects in population two, $\mathbf{H}_{i_2} = \mathbf{B} + \mathbf{W}_{i_2}$ where $\mathbf{W}_{i_2} \in R^{N \times N}$. Let I be the set of differentially connected nodes where $I = \{1\}$ for the first set of simulation. For $i \in I$, we generate q off-diagonal elements corresponding to the DWEs in the i -th row and column edges connected with i from $N(.1, .02)$ and other edges of i from $N(0, .02)$. For $i \notin I$, we have $w_{ij,i_2} \sim N(0, .02)$. We consider $q = 4, 7$ and 11 to assess our method’s power to detect differentially connected region(s) when the number of DWEs increases. We construct the difference network with model-free p-values, where we conduct a two sample t-test and record one minus the p-value as the weight for each edge.

We compare the performance of DDT to that of two other tests. The first comparison method ($T_{(10\%)}$) is a standard two sample t-test of local nodal degree. For this test, we threshold the subject-specific correlation matrices to attain 10% density, evaluate the subject-level degree measure at each node and then perform a two sample t-test to compare the nodal degree across groups. We also investigated but did not include the results obtained from 15% density and 1% network density, which were less powerful in detecting differentially connected regions than 10% density. We also consider two binomial tests which are similar to DDT in that they directly assess the number of differentially weighted edges incident to a node but differ from DDT in that they apply some multiple comparison corrections to detect the DWEs. Specifically, the first binomial test, Bin_B , applies a Bonferroni correction to detect the DWEs (Tyszka et al., 2013) and the second binomial test, Bin_F , implements a less stringent FDR multiple testing correction. For both binomial tests, each node’s differential degree is the sum of all DWEs incident to it. We do not report the performance of the NBS since it is fundamentally a different test. Consequently, it performs poorly under the desired simulation settings.

In the second set of simulations, we assess the methods’ performances when there are 3 differentially connected regions. We consider two scenarios in this setting. First, the network size is fixed while the number of DWEs varies with $q = 4, 7$ and 11 . Second, we fix the proportion of DWEs for the differentially connected nodes to be 30% while increasing the size of the network. We report various metrics to quantify the methods’ accuracy in detecting differentially connected nodes across the simulations. The false positive rate (FPR) is calculated as $\sum_{s=1}^S \sum_{n=1}^N I(\hat{R}_{n,s} = 1, R_n = 0)/(S * N)$ and quantifies the chance that each method incorrectly identifies a differentially connected region. The true positive rate (TPR) is calculated as $\sum_{s=1}^S \sum_{n=1}^N I(\hat{R}_{n,s} = 1, R_n = 1)/(S * N)$ and measures the correct identification. Here, S is the total number of simulations. $\hat{R}_{n,s}$ takes the value 1 if region n in simulation s is selected

as differentially connected and 0 otherwise. R_n is a binary indicator of whether region n is differentially connected in the ground truth. We compare accuracy in selecting truly differentially connected regions by Matthews correlation coefficient (MCC) (Johnstone et al., 2012), which is a popular measure for accessing the correspondence between predicted and true class labels. MCC, which is computed as $\frac{TP \times TN - FP \times FN}{\sqrt{(TP+FP)(TP+FN)(TN+FP)(TN+FN)}}$, takes values in $[-1, 1]$ where 1 indicates perfect agreement between the predicted and true class labels, 0 no agreement, and -1 inverse agreement. In this formula, TP, TN, FP, FN denote the number of nodes that are true positives, true negatives, false positives and false negatives, respectively. In a supplementary analysis of the simulation results, we assess the performance of the adaptive thresholding procedures presented in section 2.2.2.4 in correctly detecting DWEs. We compare the MCC in selecting the true DWEs based on the proposed aDDT and eDDT thresholding procedures with that based on two hard thresholds at .95 and .99 as well as based on multiple comparison corrections thresholds using the Bonferroni and FDR methods.

2.3.1 Results

Table 2.1 displays accuracy measures for identifying one differentially connected node across two populations in the first set of simulations. Generally, the proposed DDT methods, i.e. aDDT and eDDT, exhibit larger TPR than the T-tests and the Binomial tests across various sample sizes and network structures. The Binomial tests achieve the lowest FPR, which is attributed to the Bonferroni and FDR multiple testing corrections. However, the multiplicity corrections reduce the power to detect the correct region. The T-test attains the nominal type I error rate ($\alpha = .05$). For all methods, the TPR improves when the sample size increases and the number of differentially connected edges increase. Overall, the two proposed DDT approaches exhibit superior performance as compared to the other tests. Among the two DDT

methods, eDDT typically exhibits higher TPR, but the latter has a slightly higher FPR, although the FPR under both approaches is less than the nominal level of .05.

Table 2.1: False positive and true positive rates for the Random, Small World, and Hybrid network structures considered.

		$n_1 = n_2 = 20$					$n_1 = n_2 = 40$					
		DDT		Binomial		T-test	DDT		Binomial		T-test	
Network Structure		DWE*	aDDT	eDDT	Bin _B	Bin _F	10%	aDDT	eDDT	Bin _B	Bin _F	10%
FALSE POSITIVE RATE	Random	4	.021	.046	.002	.001	.052	.019	.046	.002	.002	.052
		7	.022	.047	.002	.002	.051	.017	.044	.002	.002	.054
		11	.020	.046	.002	.002	.052	.014	.044	.003	.002	.055
		20	.017	.045	.0004	.003	.057	.008	.033	.0008	.004	.060
	Small world	4	.022	.046	.001	.001	.050	.019	.045	.002	.002	.054
		7	.023	.046	.002	.002	.051	.016	.044	.002	.002	.051
		11	.020	.045	.002	.002	.052	.012	.042	.003	.002	.055
		20	.019	.045	.0004	.003	.054	.008	.033	.0005	.003	.058
	Hybrid	4	.024	.042	.001	.002	.054	.019	.045	.001	.001	.055
		7	.022	.045	.002	.001	.055	.018	.045	.002	.002	.059
		11	.020	.046	.002	.002	.054	.015	.046	.002	.002	.064
		20	.017	.045	.0005	.003	.058	.009	.034	.0006	.003	.063
TRUE POSITIVE RATE	Random	4	.370	.458	.036	.049	.111	.710	.891	.123	.112	.203
		7	.631	.767	.240	.292	.230	.977	.991	.747	.738	.379
		11	.893	.885	.694	.686	.450	1.00	.999	.999	.998	.619
		20	.994	.991	.981	.999	.791	1.00	1.00	1.00	1.00	.962
	Small world	4	.287	.505	.040	.034	.131	.699	.913	.113	.123	.227
		7	.639	.738	.274	.226	.155	.982	.994	.784	.764	.307
		11	.895	.908	.696	.692	.551	1.00	.999	.998	.999	.840
		20	.994	.991	.981	.999	.791	1.00	1.00	1.00	1.00	.962
	Hybrid	4	.284	.423	.042	.052	.255	.650	.884	.136	.114	.435
		7	.571	.675	.225	.222	.524	.974	.995	.707	.719	.650
		11	.874	.888	.693	.693	.681	.999	1.00	.997	.994	.739
		20	.996	.988	.975	.993	.820	1.00	1.00	1.00	1.00	.978

*Number of Differentially Weighted Edges incident to node 1

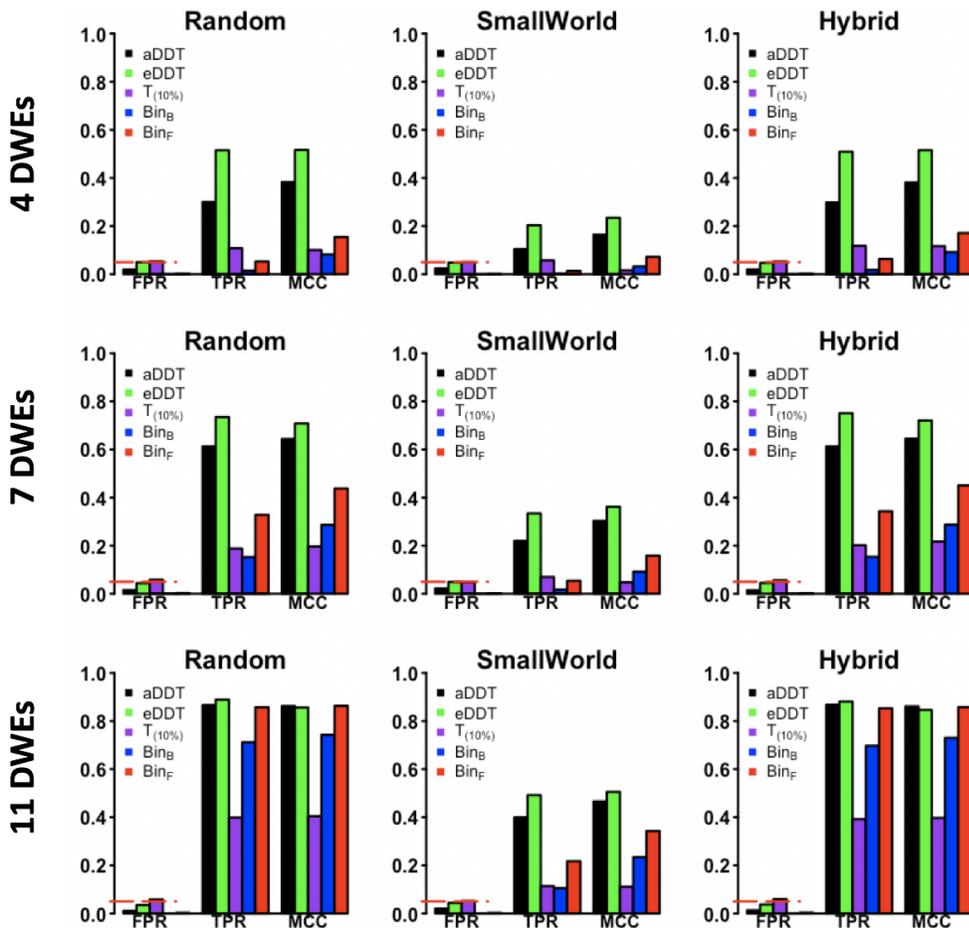
Bin_B = Binomial, bonferroni correction; **Bin_F** = Binomial, FDR correction

aDDT = DDT, theoretical threshold; **eDDT** = DDT, empirical threshold

The advantages of proposed aDDT and eDDT over the alternative methods persist in the second set of simulations where three regions are differentially connected. In Figure 2.1, across all network structures with a fixed number of nodes ($N=35$) and with four, seven, or eleven DWEs incident to each of the three nodes of interest, the DDT methods have the highest power to detect the regions of interest while attaining FPR comparable to that of $T_{(10\%)}$. We note that our method is superior to the multiplicity corrected Binomial tests when the differentially connected regions are incident to a small to moderate number of DWEs and is comparably powered to

detect differentially connected nodes as the FDR corrected tests when the number of DWEs is large. Furthermore, as in the first simulation setting, eDDT typically exhibits higher TPR than aDDT, but the former has slightly higher FPR compared to aDDT. Notably, both methods exhibit FPR values close to the nominal level of 0.05.

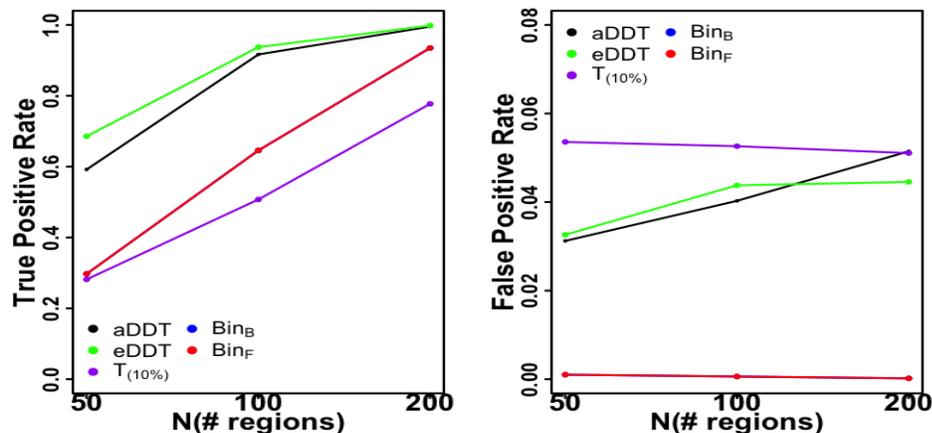
Figure 2.1: Comparison of eDDT, aDDT, t -test ($T_{(10\%)}$), binomial tests (Bin_F , Bin_B) in the second set of simulations with three differentially connected nodes incident to four (*first row*), seven (*second row*), and eleven (*third row*) DWEs. The true positive rate (TPR), false positive rate (FPR) and Matthew’s correlation coefficient (MCC) is presented for all methods across the three network structures considered and the red dashed line demarcates the nominal significance level (.05). DDT exhibits superior performance in detecting the differentially connected nodes while not exceeding the allowable type I error rate.



We also examine at the performance of the approaches as the number of nodes increases, while keeping the proportion of DWEs incident to the region of interest

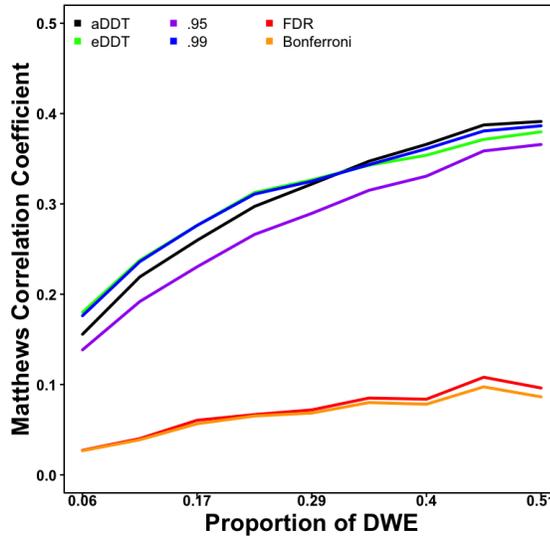
fixed at 30%. Figure 2.2 clearly illustrates the advantages of DDT for detecting regions incident to DWEs, while having a comparable or lower FPR as the network’s size increases. Consistent with Table 2.1 and Figure 2.1, eDDT exhibits the best TPR while the multiplicity corrected binomial tests have the smallest FPR, although the FPR levels under the DDT approaches are less than or equal to the nominal level across varying numbers of nodes. However, the TPR for eDDT and aDDT becomes increasingly similar as the number of regions is increased.

Figure 2.2: Performance of aDDT, eDDT, t -test($T_{10\%}$), and binomial (Bin_B , Bin_F) tests in identifying the differentially connected node as the network size increases and the proportion of DWEs is fixed at 30%. (Note, the results of the Bin_B , Bin_F tests are very close and hence the two lines overlap)



Although detection of regions incident to a significant number of DWEs is our primary focus, we also investigate the performance of the thresholding procedure for detecting DWEs in a supplementary analysis in terms of the MCC values. Figure 2.3 indicates that aDDT’s and eDDT’s adaptive thresholding procedures outperform the Bonferroni and FDR multiplicity corrections over varying proportion of DWEs. Moreover, our method also exhibits superior MCC than the arbitrary hard threshold of 0.95, and at least one of the aDDT and eDDT approaches perform as well as the conservative hard threshold set at 0.99 as the proportion of DWEs across the network increases.

Figure 2.3: Comparison of thresholding procedures implemented in aDDT and eDDT as well as four competitors (.95 and .99 hard threshold; FDR and Bonferroni multiplicity corrections) in detecting differentially weighted edges.



2.4 Data Application

Existing literature has identified multiple brain regions implicated in major depressive disorder (MDD). For example, MDD patients experience reduced connectivity in the fronto-parietal network as well as modified activity in areas such as the insula (Deen et al., 2010), amygdala (Sheline et al., 1998), hippocampus (Lorenzetti et al., 2009; Schweitzer et al., 2001), dorsomedial thalamus (Fu et al., 2004; Kumari et al., 2003), subgenual and dorsal anterior cingulate cortex (Mayberg et al., 1999). We apply the DDT to a MDD resting-state fMRI study (Dunlop et al., 2017) to investigate brain regions contributing to differences in overall functional network organization in the affected population.

To construct the brain network, we choose the 264-node system defined by (Power et al., 2011a). Each node is a 10mm diameter sphere in standard MNI space representing a putative functional area consistently observed in task-based and resting-state fMRI meta-analysis. We focus upon 259 nodes located in cortical and subcortical regions, excluding a few nodes lying in the cerebellum. Each node is assigned to one of

twelve functional modules defined in Power et al. (2011a): sensor/somatomotor (SM), cingulo-opercular task control (CIO), auditory (AUD), default mode (DMN), memory retrieval (MEM), visual (VIS), fronto-parietal task control (FPN), salience (SAL), subcortical (SUB), ventral attention (VAN), dorsal attention (DAN), and uncertain (UNC).

We measure the functional association between all pairs of brain regions with Pearson correlation and partial correlation. Partial correlations are estimated using the *DensParcorr* R package (Wang et al., 2016). For both correlation measures, we conduct the between-group tests on the Fisher Z-transformed correlation coefficient at each edge and derive both the unadjusted p-values when confounding variables are not accounted for and also the adjusted p-values when they are accounted for.

We construct the difference networks based on the model-free and model-based between-group test p-values for connectivity measured by both Pearson correlation and partial correlation. We apply the proposed eDDT method to identify brain regions incident to a statistically significant number of DWEs. Subsequently, we investigate the distribution of the DWEs across the networks as well as between and within functional modules. We also comment on the differences identified by the NBS. We apply a threshold of three as used in the seminal paper on the method (Zalesky et al., 2010).

2.4.1 Data

Our data consists of resting state fMRI scans for twenty MDD patients and nineteen healthy subjects. Details on the data acquisition and preprocessing prior to analysis can be found in 1.3.1.

2.4.2 Results

Table 2.2(A) and 2.2(B) list the top twenty differentially connected nodes for model-based Pearson and partial correlations. Pearson correlation generally leads to more DWEs incident to nodes. Thirty percent of the regions identified in Table 2.2(A) are located in the SM module while twenty percent are in the DMN. Similarly, Table 2.2(B) shows DMN nodes are extremely prominent (35%) as well as the FPN and CIO which compose the task control system. Similar trends are observed for Bin_B and Bin_F . Additionally, these two competing methods detect five regions that are also found by our method (*Temporal_Sup_L*, *Insula_R*, *Frontal_Inf_Orb_L*, *Supp_Motor_Area_R*, *Frontal_Sup_Medial_L*). Jointly, these results suggest that altered connectivity in the DMN differentiates the brain networks in the MDD population from healthy controls.

Figure 2.4 displays the distribution of DWEs across the respective difference network. Here, we group the nodes based on the functional module assignment provided in Power et al. (2011a). The diagonal blocks represent within-module connections while the off-diagonal blocks represent between-module connections. For Pearson model-free and model-based analyses, we identified 793 and 776 DWEs, respectively. For partial correlations, we identified 458 DWEs based on model-free p-values and 772 DWEs for model-based p-values. The Pearson correlation derived difference networks exhibit spatial clustering of DWEs, specifically within the SM and between the SUB and VIS functional modules. Table 2.3 reports the consistently and inconsistently detected DWEs when comparing the four difference networks investigated, i.e. model-free/model-based Pearson correlation networks and model-free/model-based partial correlation networks. Insignificant edges persist across all the difference networks considered and account for as much as 90% of the edges in the networks. Generally, the findings are more consistent between the model-free and model-based p-values within the same correlation measure and less consistent across correlation

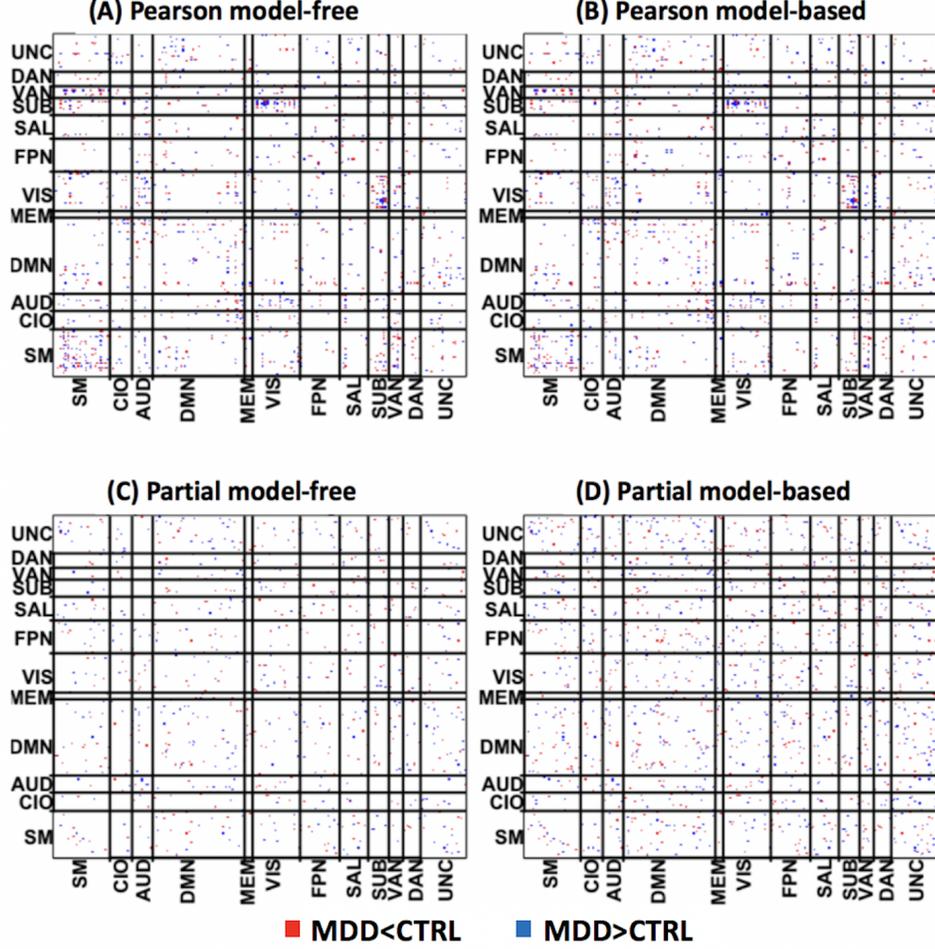
Table 2.2: Top twenty differentially connected nodes in the major depressive disorder study based on (a) model-based Pearson correlations and (b) model-based partial correlations.

(A) model-based Pearson correlations						
X	Y	Z	Name	Module	#DWE	
-53	-22	23	SupraMarginal_L (aal)	Auditory	29	
52	7	-30	Temporal_Pole_Mid_R (aal)	Default mode	27	
46	16	-30	Temporal_Pole_Mid_R (aal)	Default mode	26	
29	1	4	Putamen_R (aal)	Subcortical	26	
47	-30	49	Postcentral_R (aal)	Sensory/somatomotor Hand	23	
10	-46	73	Precuneus_R (aal)	Sensory/somatomotor Hand	19	
-24	-91	19	Occipital_Mid_L	Visual	19	
-54	-23	43	Parietal_Inf_L	Sensory/somatomotor Hand	18	
31	33	26	Frontal_Mid_R	Saliency	18	
51	-29	-4	Temporal_Mid_R (aal)	Ventral attention	18	
13	-33	75	Postcentral_R (aal)	Sensory/somatomotor Hand	16	
-46	31	-13	Frontal_Inf_Orb_L (aal)	Default mode	16	
23	10	1	Putamen_R (aal)	Subcortical	15	
-44	12	-34	Temporal_Pole_Mid_L (aal)	Default mode	15	
31	-14	2	Putamen_R (aal)	Subcortical	15	
-38	-27	69	Postcentral_L (aal)	Sensory/somatomotor Hand	14	
-60	-25	14	Temporal_Sup_L (aal)	Auditory	14	
27	16	-17	Insula_R (aal)	Uncertain	14	
52	-2	-16	Temporal_Mid_R (aal)	Default mode	14	
50	-20	42	Postcentral_R (aal)	Sensory/somatomotor Hand	13	
(B) model-based partial correlations						
-31	19	-19	Frontal_Inf_Orb_L (aal)	Uncertain	15	
24	32	-18	Frontal_Sup_Orb_R (aal)	Uncertain	13	
-38	-15	69	undefined	Sensory/somatomotor Hand	13	
-26	-40	-8	ParaHippocampal_L (aal)	Default mode	13	
-31	-10	-36	Fusiform_L (aal)	Uncertain	13	
17	-91	-14	Lingual_R (aal)	Uncertain	12	
-16	-46	73	Parietal_Sup_L (aal)	Sensory/somatomotor Hand	11	
23	33	48	Frontal_Sup_R (aal)	Default mode	11	
-28	-79	19	Occipital_Mid_L (aal)	Visual	11	
37	-81	1	Occipital_Mid_R (aal)	Visual	11	
-42	-55	45	Parietal_Inf_L (aal)	Fronto-parietal Task Control	11	
-54	-23	43	Parietal_Inf_L (aal)	Sensory/somatomotor Hand	10	
7	8	51	Supp_Motor_Area_R (aal)	Cingulo-opercular Task Control	10	
-45	0	9	Rolandic_Oper_L (aal)	Cingulo-opercular Task Control	10	
-60	-25	14	Temporal_Sup_L (aal)	Auditory	10	
-13	-40	1	Precuneus_L (aal)	Default mode	10	
-68	-23	-16	Temporal_Mid_L (aal)	Default mode	10	
-10	39	52	Frontal_Sup_Medial_L (aal)	Default mode	10	
22	39	39	Frontal_Sup_R (aal)	Default mode	10	
-8	48	23	Frontal_Sup_Medial_L (aal)	Default mode	10	

measure.

The distribution of DWEs within and between functional modules provides insight into disrupted communication among functionally segregated sub-systems in the brain. We conduct analysis to identify functional modules that are associated

Figure 2.4: Differentially weighted edges detected by eDDT in the major depressive disorder study under the four difference network configurations: (a) model-free Pearson and (b) model-based Pearson (c) model-free partial (d) model-based partial. Red edges indicate the average edge weight in the MDD population is statistically smaller than in healthy adults whereas blue edges demarcate the edge is statistically larger. The network is decomposed into 12 functional modules: Sensor/somatomotor (SM), Cingulo-opercular Task Control (CO), Auditory(Aud), Default Mode (DMN), Memory Retrieval (Mem), Visual (Vis), Fronto-parietal Task Control (FPN), Salience (SN), Subcortical (Sub), Ventral attention (VAN), Dorsal attention (DAN), Uncertain (Un).



with higher number of DWEs as compared with other modules. Specifically, we propose the following chi-square statistic to help identify functional module pairs for which there are unusually high number of DWEs than what is expected by chance,

$$X_{g_1, g_2}^2 = \frac{(Q_{(g_1, g_2)} - E_{(g_1, g_2)})^2}{E_{(g_1, g_2)}}, \quad (2.7)$$

Table 2.3: Consistency of DWEs identified based on four difference networks. Values presented in the table are the number of network edges for each results consistency classification.

Difference Network		Sig in I	Insig in I	Sig in I	Insig in I
I	II	Sig in II	Insig in II	Insig in II	Sig in II
Pearson model-free	Pearson model-based	662	32504	131	114
Pearson model-free	Partial model-free	51	32392	691	406
Pearson model-based	Partial model-based	73	31936	703	699
Partial model-based	Partial model-free	454	32936	4	318

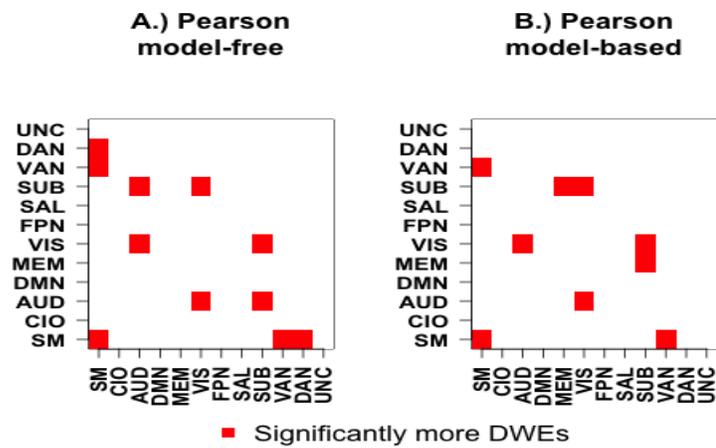
The total number of edges in the network is 33411.

where $g_1 \in \{1, \dots, \mathcal{G}\}$ and $g_2 \in \{1, \dots, \mathcal{G}\}$ are indices corresponding to one of the $\mathcal{G} = 12$ functional modules. When $g_1 = g_2$, (g_1, g_2) represents a within module block, whereas it represents a between-module block when $g_1 \neq g_2$. $Q_{(g_1, g_2)}$ represents the observed number of DWEs in the (g_1, g_2) block and $E_{(g_1, g_2)}$ represents the expected number of DWEs in the (g_1, g_2) block when the edges distribute randomly across the module blocks in the network. Let $|g|$ represent the total number of nodes within the g th module, and p^* represent the proportion of DWEs among all the edges across the network. It is straightforward to see that $E_{g_1, g_2} = p^* * \lfloor \frac{|g_1| * (|g_2| - 1)}{2} \rfloor$ for within module blocks, i.e. $g_1 = g_2$, and $E_{g_1, g_2} = p^* * [|g_1| * |g_2|]$ for between-module blocks.

Figure 2.5 displays functional modules and module pairs exhibiting a significantly high number of DWEs based on the thresholded chi-square test statistic. The results are derived from the model-free Pearson correlations (Figure 2.5(A)) and model-based Pearson correlations (Figure 2.5(B)), respectively. Based on model-free Pearson correlations (Figure 2.5(A)), there are significantly high number of DWEs within the sensorimotor module and between the module pairs of sensorimotor-ventral attention, sensorimotor-dorsal attention, visual-auditory, subcortical-auditory and subcortical-visual. After accounting for age and gender, the model-based Pearson correlations (Figure 2.5(B)) also exhibit a large number of DWEs within the sensorimotor module and between the module pairs of sensorimotor-ventral attention, visual-auditory and subcortical-visual. However, the model-based Pearson correlations no longer show

significantly high number of DWEs between the sensorimotor-dorsal attention and subcortical-auditory module pairs. Instead, the model-based correlations find significant number of DWEs between the subcortical-memory module pair which is not identified by the model-free Pearson correlations.

Figure 2.5: Heat map of the $X^2_{(g1,g2)}$ statistic for (A) model-free Pearson correlations and (B) model-based Pearson correlations. Red squares indicate modules with more statistically significant DWEs than that consistent with the null hypothesis. We control the overall false discovery rate by only selecting module pairs with a multiplicity corrected p-value $< .05$.



2.5 Discussion

While the estimation of brain networks is gaining increasing attention in the neuroimaging literature, the fundamental question of how brains differ in functional organization across disease populations is not yet resolved. Our proposed method exhibits two strengths. First, our automated threshold selection permits identification of DWEs without sacrificing power as is the case with many methods dependent upon multiplicity corrections. Second, we use the generated null networks to test if each brain region is incident to more DWEs than would be expected by random chance.

We hypothesize that network wide disconnectivity is driven by brain regions that irregularly communicate with other regions. The results from the real data analysis suggest that the DDT appropriately identifies problematic brain regions in major depressive disorder. The existence of differential connectivity between nodes in the auditory and visual networks (Figure 2.5) has previously been observed (Eyre et al., 2016). Further, multivariate pattern analyses have suggested that the most discriminative functional connectivity patterns lie within and across the visual network, DMN, and affective network (Zeng et al., 2012). The parahippocampal gyrus, which we detect as a problematic region (Table 2.2), has also suggested as a region with differentiated connectivity patterns in depressed populations.

Our simulation results demonstrate superior performance of the proposed DDT tests. Although the binomial test’s false positive rate is smaller than DDT, its ability to detect differentially connected nodes is severely attenuated. DDT maintains the desired false positive rate while achieving higher true positive rates than the t -test across all network structures and sample sizes considered. Simulations indicate DDT’s adaptive threshold selection is superior to conservative FDR and Bonferroni adjustments.

An obvious limitation to this work is the *i.i.d.* assumption on edge weights in the null networks. Although the independence assumption did not severely impact the

method's performance relative to suitable competitors, it is likely that incorporation of inter-edge dependence structures will lead to better power to detect differentially connected nodes.

Chapter 3

Anatomically Informed Estimation of Functional Brain Networks

3.1 Introduction

In the previous chapter, we developed a method to compare functional brain networks. Although many methods compare functional connections across populations, it is well known that brain functional connectivity (FC) exhibits substantial within subject variability. This variability, which is typically unaccounted for, can contribute to erroneous edge detection and negatively impact network comparisons. In this chapter, we propose a new method for estimating functional connectivity when structural connectivity (SC) knowledge is available. Our motivation for such an approach is two fold: (1) evidence that brain structure constrains function and (2) potential increases in the reliability of functional connectivity estimates in the presence of structural knowledge. Various studies have demonstrated the suggestive relationship between brain function and structure.

Despite strong evidence regarding the role of white matter fiber tracts in regulating FC (Damoiseaux and Greicius, 2009; Honey et al., 2010; Sporns, 2013) and considerable progress in separately estimating FC and SC, there have been comparatively limited advances in FC approaches which are guided by underlying anatomical knowledge. Incorporating anatomical knowledge in estimating FC is clearly desirable since it is expected to produce more accurate estimates of the network, which translates to greater reproducibility of the findings as illustrated via our fMRI data analysis. However, several considerations need to be taken into account, such as the complexity of the structure-function relationship (Hermundstad et al., 2013), heterogeneity in FC for a given SC strength, which presumably is attributed to the fact that FC is only partially dependent on SC (Damoiseaux and Greicius, 2009; Messé et al., 2014) and regulated by unobservable dynamics in underlying neuronal activity (Bressler and Tognoli, 2006).

Recently, Venkataraman et al. (2012) and Xue et al. (2015) proposed approaches to jointly model the probability of co-activation based on fMRI data while

incorporating direct structural connections. They provide measures of functional co-activation deviating from standard measures of FC such as Pearson or partial correlations. Hinne et al. (2014) proposed a Bayesian approach which uses fMRI data to model the distribution of partial correlations for edges determined by the given SC information. The assumption that FC only exists between anatomically connected regions ignores the contributions of indirect anatomical pathways and does not capture the complexity of the relationship between brain structure and function (Honey et al., 2009, 2010; Messé et al., 2014). Moreover, the above approaches use multi-subject data which requires registration of images to a shared template under the assumption that the volumes are similar and can be matched. Unfortunately, this assumption has limitations for human brain images considering the substantial variability in cortical anatomy and function (Zhu et al., 2012). This variability is especially pronounced during the developmental phases of childhood and adolescence as is the case with our motivating Philadelphia Neurodevelopmental Cohort (PNC) study. Ng et al. (2012) and Pineda et al. (2014) proposed approaches for estimating sparse functional networks for individual subjects via an adaptive graphical lasso. In both models, edge specific shrinkage parameters are deterministic functions of the SC information. Under these approaches, FC with less anatomical support are more heavily penalized, and vice-versa. However, the parametric form of the shrinkage parameters may not adequately capture the complex underlying structure-function relationships and does not account for heterogeneity in FC for a given SC strength resulting from non-anatomical sources of variation such as BOLD signal sensitivity to proximity to blood vessels (Zhang et al., 2016) as well as cardiac, acquisition and pre-processing effects (Skudlarski et al., 2008). Moreover, such a parametric relationship may lead to network estimates which are not robust and more sensitive to the misspecification of anatomical knowledge. Figure 3.1 illustrates a representative FC-SC relationship in the PNC study (refer to section 3.3.2.1 for details on the calculation

of SC strengths).

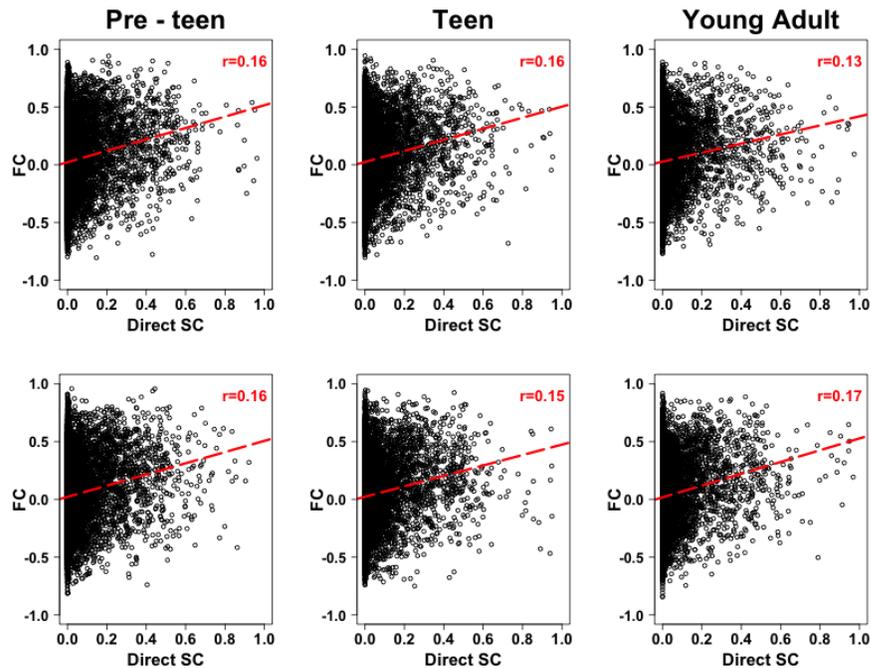


Figure 3.1: Plots of the associations between partial correlation (FC) and direct structural connectivity (SC) for all males (*top row*) and females (*bottom row*) in our study. While mild positive correlations—approximately .16—are observed between SC and FC, FC exists between regions with little to no direct structural connections. We also observe large variation in FC for a given SC level (median standard deviation $\approx .275$). The red dashed line is the line of best fit and r is the Spearman correlation coefficient between direct SC and FC.

The above discussions highlight a serious need for integrative modeling approaches which adaptively estimate FC by incorporating structural knowledge in an appropriate manner. In designing such an approach, our primary goals for the method are that it will (a) correctly identify true interregional functional connections and identify which connections are absent in the true network; (b) lead to reproducible estimates of FC such that the estimated brain network can be replicated across multiple scanning sessions, which is a topic of great importance in current literature (Varoquaux et al., 2010); and (c) specify a flexible structure-function relationship which is robust to misspecification of SC information (arising from limitations in existing image acquisition technology) and can accommodate heterogeneity in FC for a given SC. We propose

a novel hierarchical Bayesian Gaussian graphical modeling approach for estimating FC based on single subject fMRI data which incorporates given SC information in a manner that addresses the aforementioned objectives. The FC is computed via sparse precision matrices whose elements are estimated under Laplace type priors having edge specific shrinkage parameters that are random variables modeled using SC information and an independent baseline component. The prior encourages stronger FC given a large SC (and vice-versa), but also accounts for edge specific variations in FC unrelated to the brain anatomy, via the baseline component. Thus, the approach is flexible in accounting for anatomical knowledge with the FC being guided by, but not completely determined by, the SC information. Our method is motivated by the variable selection approach in Chang et al. (2016), which incorporates prior graph knowledge in a linear regression setting, but is distinct in addressing graphical model selection and precision matrix estimation, as well as the manner in which the prior knowledge is incorporated. Under certain choices of model parameters, the proposed approach reduces to an adaptive shrinkage approach specifying a parametric relationship between the shrinkage parameters and the anatomical information, similar to Ng et al. (2012) and Pineda et al. (2014).

While Markov chain Monte Carlo (MCMC) can be used to implement the proposed approach, it is not scalable to large networks needed for whole brain connectome analysis as in our application. Moreover, an additional thresholding step is often needed for model selection. We propose an optimization algorithm to obtain the *maximum a posteriori* (MAP) estimate, which is computationally efficient, scales to a large dimensions, and does not require post-hoc thresholding. Under various simulation studies, we observe superior performance of the proposed method as compared to alternative approaches with or without SC information. The advantages of our approach become more evident as the degree of misspecification of anatomical knowledge increases, and/or the number of nodes grows larger which is particularly

relevant for whole brain connectome analysis.

Our efforts are motivated by data from the PNC study, a large-scale, NIMH funded initiative to understand the developmental trajectory of the brain from childhood to adolescence (Satterthwaite et al., 2014). The PNC data contains both DTI and resting state fMRI measurements from boys and girls ages 8-21 years, with suggestive but unclear SC-FC relationships (Figure 3.1). We fuse multimodal brain imaging data to examine gender differences in brain networks across three age brackets—pre-teens (ages 8-12), teens (ages 13-17), and young adults (ages 18-21)—and discover several gender based differences in FC within and across the age brackets. We also assess the reliability of computed network metrics across scanning sessions and find that the proposed approach yields strong reproducibility in the estimation of network metrics, which is almost always higher than alternative approaches. While other studies have separately examined the reproducibility of functional and structural brain networks (see Welton et al. (2015) for a review), ours is one of the first to examine the reproducibility of anatomically informed functional networks to the best of our knowledge.

Section 2 describes the proposed methodology and the optimization routine for estimating networks, while sections 3 presents numerical studies and application of our method to PNC data. We conclude with a brief discussion in Section 4.

3.2 Materials and methods

3.2.1 Gaussian graphical model for brain networks

While early work on brain network estimation utilized Pearson correlation to measure undirected interregional dependencies, recent literature has focused on Gaussian Graphical Models (GGMs) which are parametrized by the inverse covariance matrix and measure functional connectivity via partial correlations. The precision matrix

has often been reported to result in more accurate and robust estimation of underlying brain network structures as compared to other methods such as thresholding of the covariance matrix (Smith et al., 2011). Compared to the covariance approach, the precision matrix can distinguish a true, direct functional connection between two regions from those that exist because of confounding with other nodes in the network.

GGMs assume observations are normally distributed and that zeros in the inverse covariance matrix correspond to absent edges in the network, \mathcal{G} . A standard GGM specifies $\mathbf{y}_t \sim N_p(0, \mathbf{\Omega}^{-1})$ where \mathbf{y}_t is a vector containing the fMRI signal at each of the p ROIs for the t^{th} time-index, where $t=1, \dots, T$ and T is the total number of image volumes acquired during the scanning session. In everything that follows, $\mathbf{Y}_{T \times p} = \{y_1, \dots, y_T\}$ is the data matrix where each row contains the fMRI signals across all brain regions at time index t . $\mathbf{\Omega}$ is the inverse covariance matrix (or precision matrix) which belongs to the cone of $p \times p$ symmetric positive definite matrices, denoted by M_p^+ . Under this framework, estimating \mathcal{G} is equivalent to estimating structural zeros in the positive definite precision matrix $\mathbf{\Omega}$.

Due to tradeoffs between the cost and efficiency of information transfer, it is typically assumed that the brain seeks efficient organization favoring a sparse set of active connections at any point in time (Eavani et al., 2015). The GGM approach is well equipped to handle such sparse networks by imposing penalties that shrink sufficiently weak functional connections to zero, where the L_1 penalty under the graphical lasso (Friedman et al., 2008) is a popular choice (Ng et. al, 2012; Monti et. al, 2014 ; Pineda-Pardo et. al, 2014). The graphical lasso can be thought of as an extension of the Lasso approach in regression settings and penalizes the full data likelihood to estimate the inverse covariance matrix as

$$\hat{\mathbf{\Omega}} = \arg \max_{\mathbf{\Omega} \in M_p^+} \log \det(\mathbf{\Omega}) - \text{tr}(\mathbf{S}\mathbf{\Omega}) - \lambda \sum_{j \leq k} |\omega_{jk}|, \quad (3.1)$$

where $\mathbf{S} = \sum_{t=1}^T (\mathbf{y}_t \mathbf{y}_t^T) / T$ is the sample covariance matrix, $|x|$ denotes absolute value of x , $\det(\cdot)$ is the determinant operator, $\text{tr}(\cdot)$ is the matrix trace operator, and $\lambda > 0$

is the penalty parameter controlling overall network sparsity. If $\lambda = 0$, one obtains the maximum likelihood estimate, while large values of λ shrinks an increasing number of off-diagonal elements to zero. The typical graphical lasso approach fits a series of graphs under various choices of the tuning parameter λ and chooses the optimal network as the one minimizing some goodness of fit criteria (Yuan and Lin, 2006).

3.2.2 Structurally informed Bayesian Gaussian graphical model

Bayesian GGM approaches have been successfully used for estimating brain networks (see Mumford and Ramsey (2014) for a review). One such approach is the Bayesian graphical lasso (Wang, 2012), which has similarities with the graphical lasso approach in the sense that the *maximum a posteriori* (MAP) estimator is equivalent to the penalized likelihood estimate. This approach assumes that the p dimensional fMRI signal at time index t is distributed as $\mathbf{y}_t \sim N_p(0, \mathbf{\Omega}^{-1}), t = 1, \dots, T$, with the prior on the inverse covariance as

$$\pi(\mathbf{\Omega} \mid \lambda) = C_{\lambda}^{-1} \prod_{k=1}^p \text{Exp}(\omega_{kk}; \lambda) \prod_{j < k} \text{DE}(\omega_{jk}; \lambda) I(\mathbf{\Omega} \in M_p^+), \quad (3.2)$$

where $\pi(\cdot)$ represents the prior distribution and $I(x)$ is an indicator function that takes the value one when condition x is true. The diagonal element ω_{kk} is modeled under an exponential prior distribution $\text{Exp}(\lambda)$, the off-diagonal element ω_{jk} is modeled with double exponential or Laplace prior distribution $\text{DE}(\lambda)$, and λ is the shrinkage parameter. In a fully Bayesian paradigm, λ is typically assigned a prior distribution, and is thus learned from the data, resulting in an adaptive shrinkage of the elements in $\mathbf{\Omega}$.

In order to incorporate anatomical knowledge in functional connectivity estimation, we propose a hierarchical Bayesian structurally informed Gaussian graphical model (siGGM). It is based on the generic Bayesian GGM in eq. (3.2), but involves

edge specific shrinkage parameters which are modeled using anatomical knowledge. Throughout this article, we will denote the structural connectivity metric as p_{jk} for edge (j, k) , where a larger value denotes a stronger anatomical connection and vice-versa. For example, in our data application, p_{jk} corresponds to the probability of SC obtained via probabilistic tractography (please see section 3.3.2.1 for details). The proposed approach to estimating the brain functional network incorporating anatomical knowledge is defined as follows

$$\begin{aligned}\pi(\boldsymbol{\Omega} | \boldsymbol{\lambda}) &= C_{\boldsymbol{\lambda}, \nu}^{-1} \prod_{k=1}^p \text{Exp}(\omega_{kk}; \frac{\nu}{2}) \prod_{j < k} \text{DE}(\omega_{jk}; \nu \lambda_{jk}) \mathbf{I}(\boldsymbol{\Omega} \in \mathbf{M}_p^+), \\ \pi(\boldsymbol{\lambda} | \boldsymbol{\mu}, \eta) &= C_{\boldsymbol{\lambda}, \nu} \prod_{j < k} \text{LN}(\mu_{jk} - \eta p_{jk}, \sigma_{\lambda}^2),\end{aligned}\tag{3.3}$$

where (i) $\boldsymbol{\lambda} = \{\lambda_{jk}, j < k, j, k = 1, \dots, p\}$ denotes the collection of edge specific shrinkage parameters having a log-normal (LN) type distribution which restricts the shrinkage parameters to non-negative values; (ii) ν refers to the tuning parameter controlling the network's overall sparsity and also corresponds to the scale parameter for the exponential prior on the diagonals; (iii) η is a positive random variable which controls the average effect of SC on FC; (iv) μ_{jk} denotes the random edge specific baseline component representative of non-anatomical sources of variations regulating FC; and (v) $C_{\boldsymbol{\lambda}, \nu}$ is the intractable normalizing constant for the prior on the precision matrix depending on $\boldsymbol{\lambda}$ and ν . This constant assures a proper prior distribution on $\boldsymbol{\Omega}$; however, it is not possible to analytically evaluate it due to the constraint that $\boldsymbol{\Omega} \in \mathbf{M}_p^+$, as described in Wang (2012). We utilize the trick introduced in that paper where the intractable constant is included in the prior on $\boldsymbol{\lambda}$ such that it cancels with the term in $\pi(\boldsymbol{\Omega} | \boldsymbol{\lambda})$, leading to a closed form full posterior that facilitates computation. We note that in the extreme case when $\log(\lambda_{jk}) = \mu_{jk} - \eta p_{jk}$, the model specifies a parametric relationship which has similarities with Ng et. al (2012) and Pineda-Pardo

et. al (2014).

The anatomically informed prior on the shrinkage parameters in (3.3) specifies a probabilistic relationship between the edge specific shrinkage parameters and the given SC knowledge via η . In particular, increasing positive values of η implies an increasing dependence on the given SC, potentially resulting in a functional connection even for small SC weights. Figure 3.2 illustrates that for large η and increasing SC, the Laplace prior has heavier tails and less mass around zero, which is interpreted as increased probability of strong FC. In contrast, small values of η do not result in a noticeable change in the prior distribution under varying SC strengths, implying a negligible relationship between SC and FC. Moreover, the shrinkage parameters are stochastically monotonically decreasing with respect to the SC strength, under the restriction $\eta > 0$. This implies that as the SC strength (p_{jk}) for the edge (j, k) is increased, the corresponding shrinkage parameter λ_{jk} will take smaller values in probability, resulting in values of ω_{jk} which are away from zero. Hence, the presence of FC at edge (j, k) is encouraged for large values of p_{jk} , and similarly small values of p_{jk} will encourage greater shrinkage for ω_{jk} resulting in the absence of FC at edge (j, k) .

Additionally, the baseline effect, μ_{jk} , corresponds to variations in underlying neuronal activity that are independent of the brain anatomy. This formulation enables (a) more flexibility in the FC-SC relationship by allowing the possibility of strong FC when an anatomical connection is not obvious, and vice-versa; and (b) heterogeneity in FC across edges which possesses similar SC strength that is often encountered in practice. Overall, increasing(decreasing) absolute values of the baseline effect discourages(encourages) the presence of an edge in a manner that is independent of the anatomical information. Although one could also accommodate variations in FC for a given SC strength via σ_λ while keeping μ fixed, our formulation incorporating an edge specific baseline effect confers several advantages. First, it permits the accom-

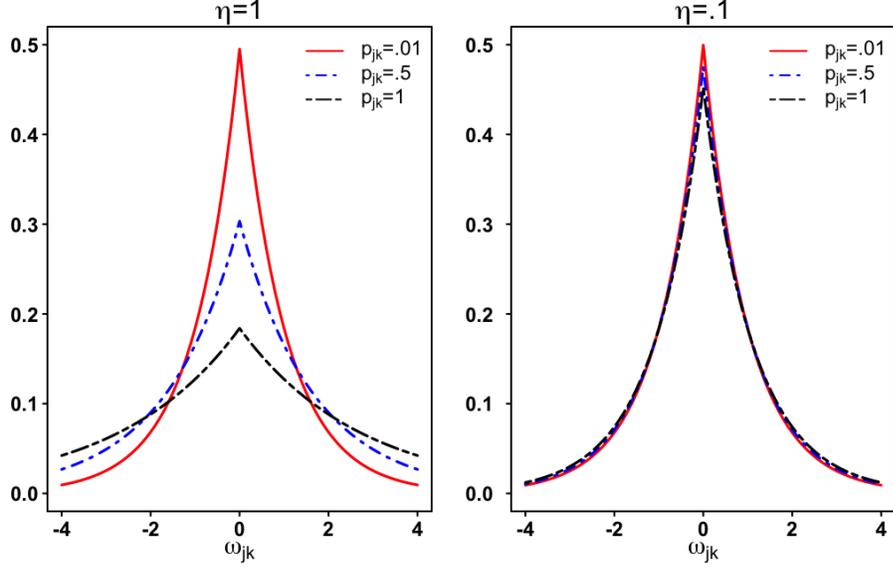


Figure 3.2: Prior distribution of ω_{jk} when $\lambda_{jk} = \mu_{jk} - \eta p_{jk}$. μ_{jk} is fixed at zero with varying values of η and SC. Solid, dashed and dotted lines correspond to $p_{jk} = 0.01, 0.5$, and 1, respectively. Left panel (a): for large η values ($\eta = 1$), the prior places increasing mass at the tails, which encourage stronger functional connectivity; right panel (b): for small values ($\eta = 0.1$), the prior on ω does not change noticeably with the change in SC information.

modation of edge specific tuning of the shrinkage parameters while σ_λ controls global variation in FC. Second, μ_{jk} can be used to differentiate edge specific variation in FC that is not attributed to direct SC, and can be used to characterize edges with enhanced or trivial SC influences. The hyperparameters $\boldsymbol{\mu}$ and η are unknown and are learned in a data-adaptive manner under the following priors

$$\begin{aligned}\pi(\mu_{jk}) &= N(\mu_0, \sigma_\mu^2) \quad \text{for } j < k, j, k = 1, \dots, p \\ \pi(\eta) &= Ga(a_\eta, b_\eta)\end{aligned}\tag{3.4}$$

where (μ_0, σ_μ^2) , and (a_η, b_η) , are typically pre-specified. In equation (3.4), $Ga(a_\eta, b_\eta)$ is the gamma distribution with scale parameter a_η , rate parameter b_η , and expected value a_η/b_η . We select these distributions for μ_{jk} and η because they are conjugate and provide adequate performance under the proposed method. Hyperparameters governing these priors are discussed in Appendix C.1. We note that the scale parameter ν

controls the overall network sparsity and is treated as a tuning parameter, enabling the estimation of a series of networks with varying sparsity levels. The optimal network is chosen as the point estimate corresponding to the value of ν minimizing the Bayesian Information Criteria.

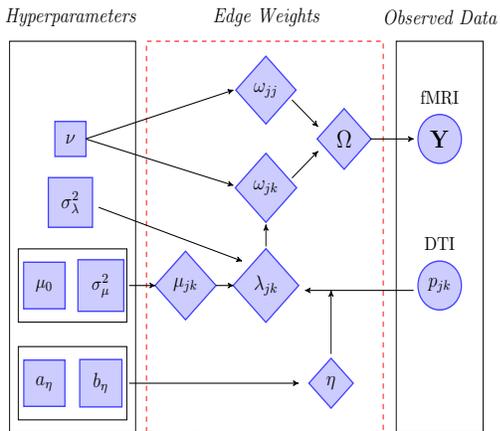


Figure 3.3: Graphical illustration of model parameters and their contribution to estimation of anatomically informed functional connectivity based on resting state fMRI data \mathbf{Y} . Circles represent observed data, diamonds represent parameters to be updated, and squares represent fixed values. Hyperparameters contained in rectangles jointly inform the distribution of the associated edge weight parameter. Parameters in the red, dashed edge weights pane are estimated outputs of the siGGM.

3.2.3 Model Estimation

Although the proposed model can be implemented using MCMC, it is not scalable to high dimensional settings involving a large number of nodes. Moreover, MCMC samples require a post hoc thresholding step to select important edges since estimates cannot take exact zeros under a Laplace prior. We bypass these limitations by computing a MAP estimate for the parameters of interest. Our iterative optimization approach employs an existing graphical lasso algorithm to sample the precision matrix given all other parameters, coupled with additional optimization steps to sample the shrinkage parameters and associated hyperparameters inherent in the Bayesian specification (3.3). In order to fit the proposed model, we esti-

mate $\Theta = (\Omega, \alpha, \eta, \mu)$ by maximizing the log-posterior distribution in (3.6), where $\alpha = \log(\lambda) = \{\log(\lambda_{jk}), j < k\}$ and $\mu = \{\mu_{jk}, j < k\}$ denotes the vector of edge specific log-shrinkage parameters and baseline effects in (3.3), respectively. Note that α is normally distributed due to the log-normal prior placed on λ . The posterior distribution can be written as

$$\begin{aligned}
P(\Theta|y_1, \dots, y_T) &\propto P(\Theta)P(y_1, \dots, y_T|\Theta) \\
&= P(\Omega, \alpha, \mu, \eta) \prod_{t=1}^T P(y_t|\Theta) \\
&= P(\Omega|\alpha, \mu, \eta)P(\alpha|\mu, \eta)P(\mu, \eta) \prod_{t=1}^T P(y_t|\Omega^{-1}) \\
&= P(\Omega|\lambda)P(\alpha|\mu, \eta)P(\mu)P(\eta) \prod_{t=1}^T P(y_t|\Omega^{-1}) \\
&= P(\eta) \prod_{t=1}^T P(y_t|\Omega^{-1}) \prod_{j,k} P(\Omega|\lambda)P(\alpha|\mu, \eta)P(\mu) \\
&= \text{Ga}(\eta; a_\eta, b_\eta) \prod_{t=1}^T \mathcal{N}_p(y_t; 0, \Sigma) \prod_{k=1}^p \text{Exp}(\omega_{kk}; \nu/2) \\
&\quad \prod_{j < k} \text{DE}(\omega_{jk}; \nu \lambda_{jk}) \mathcal{N}(\alpha_{jk}; \mu_{jk} - \eta p_{jk}, \sigma_\lambda^2) \mathcal{N}(\mu_{jk}; \mu_0, \sigma_\mu^2) \quad (3.5)
\end{aligned}$$

We find the MAP solution for the model parameters by maximizing over the the posterior log-likelihood as $\hat{\Theta} = \arg \max_{\Theta} \dot{l}(\Theta)$, where

$$\begin{aligned}
\dot{l}(\Theta) &= -\frac{T}{2} \log|\Omega| + \frac{1}{2} \text{tr}(S|\Omega|) + \nu \sum_{j < k} e^{\alpha_{jk}} |\omega_{jk}| + \sum_{j < k} \frac{(\alpha_{jk} - (\mu_{jk} - \eta p_{jk}))^2}{2\sigma_\lambda^2} \\
&\quad - (a_\eta - 1) \log(\eta) + b_\eta \eta + \sum_{j < k} \frac{(\mu_{jk} - \mu_0)^2}{2\sigma_\mu^2} - p \log\left(\frac{1}{2}\nu\right) + \frac{1}{2}\nu \sum_{k=1}^p \omega_{kk}. \quad (3.6)
\end{aligned}$$

All parameters in the posterior distribution are updated iteratively until convergence. The precision matrix is updated given other parameters using the existing

graphical lasso algorithm, whereas μ_{jk} and η are updated via closed form expressions and α is updated via a Newton-Raphson step since a closed form solution does not exist. The iterative updates continue until $|\dot{l}(\Theta^{(m)}) - \dot{l}(\Theta^{(m+1)})| < \epsilon |\dot{l}(\Theta^{(m+1)})|$ for $\epsilon = 10^{-4}$. At convergence, $\hat{\Theta} = (\hat{\Omega}, \hat{\alpha}, \hat{\eta}, \hat{\mu})$ is the solution, where $\hat{\Omega}$ is the anatomically informed functional brain network based on single subject data. In general, the method makes its largest improvements within the first three iterations (see Figure 3.6) and converges rapidly. We note that one could alternatively treat μ and η as tuning parameters and compute a range of networks over a grid of (μ, η) values, and then select the optimal network as the one minimizing some goodness of fit criteria. However, this strategy did not result in adequate numerical performance, highlighting the advantages of specifying suitable priors on hyperparameters in order to estimate them in a data-adaptive manner. The computational steps for updating the model parameters are detailed in Appendix A. The method was developed in R version 3.3.0 and is available on github (<https://github.com/IxavierHiggins/siGGMrepo>).

3.3 Results and Discussion

3.3.1 Simulations

3.3.1.1 Simulation Setting

We conduct numerical studies to assess the performance of siGGM relative to SC naive and SC informed competitors. SC naive approaches are representative of methods that do not incorporate auxiliary information, and includes the graphical lasso or *Glasso* (Friedman et al., 2008), the partial correlation approach *Space* proposed by Peng et al. (2009), and the proposed approach in (3.3) without structural information obtained by setting $\eta = 0$, denoted *siGGM*($\eta = 0$). SC informed approaches incorporate anatomical information into the estimation routine. We consider the Bayesian

G-Wishart approach by Hinne et. al (2014) which treats FC as completely determined by SC and is denoted by G-Wishart as well as the adaptive graphical lasso approach by Pineda-Pardo et. al (2014) which specifies a parametric relationship between the shrinkage parameters and SC, and is denoted by aGlasso. All of the above approaches, except *Space* which estimates partial correlations, calculate sparse inverse precision matrices, where a zero off-diagonal entry implies the absence of an edge. The *Glasso* and *Space* approaches are implemented via the R packages *glasso* and *space*, respectively. We estimate the precision matrix under the G-Wishart approach using Matlab scripts available on the author’s website and incorporate adaptive weights in the *glasso* algorithm to implement aGlasso.

Data Generation: In order to assess the performance of our approach, we simulate data under three assumed network structures and consider various relationships between SC and FC. The network structures are (a) Erdos-Renyi (ER) networks consisting of edges randomly generated with probability 0.15; (b) small-world (SM) networks generated under the Watts-Strogatz model (Watts and Strogatz, 1998) in which most nodes may not be directly connected, but can reach other nodes via a small number of steps, and (c) scale-free (SF) networks generated using the preferential attachment model (Barabási and Albert, 1999), in which nodes are more likely to link to a highly connected node than to a node with few connections, resulting in a hub network. For each network, we consider varying the number of nodes corresponding to $p = 100, 200$. The data were generated using a Gaussian graphical model $\mathbf{y}_t \sim N_p(0, \Omega_{\mathcal{G}}^{-1})$ with $T = 200$ time points ($t = 1, \dots, 200$), where $\Omega_{\mathcal{G}}$ is the precision matrix with zero off-diagonal elements corresponding to absent edges in the binary network \mathcal{G} , and non-zero off-diagonal elements otherwise. Conditional on \mathcal{G} , $\Omega_{\mathcal{G}}$ is constructed as follows: the non-zero off-diagonal elements corresponding to important edges are generated from a Uniform(-1,1) distribution, and the diagonal elements were fixed to be one. In order to ensure positive definiteness, we subtracted

the minimum of the eigenvalues from each diagonal element of the generated precision matrix.

We also assess the methods’ performance for non-Gaussian data which are generated using realistic fMRI time series under spatiotemporal separability assumptions of independent component analysis (ICA), using the framework in the SimTB (Allen et al., 2011) Matlab toolbox. We generate $Y_{T \times p} = A_{T \times q} S_{q \times p} + E_{T \times p}$ where the columns of A represent the temporal dynamics of functional networks and contain realistic BOLD-type resting state fMRI time courses sampled using the neuRosim R package (Welvaert et al., 2011); S is the source map representing the spatial distribution of functional subnetworks constructed under small-world networks replicating properties of realistic brain networks; and E is a matrix of zero mean white noise. The true precision matrix, Ω_0 , can be computed theoretically as $(S' \text{cov}(A') S + \epsilon^2 I)^{-1}$, where $\epsilon = .5$ and I is the $p \times p$ identity matrix. Although this simulation approach produces realistic time series data, it results in dense precision matrices with no off-diagonal zeros. Given strong support in the literature of the sparsity of brain networks (Eavani et al., 2015; Kim et al., 2015; Power et al., 2013), we calculate the true network as having 10% density by retaining only those edges for which the corresponding partial correlation was in the top 10th percentile. However, we note that the results under the proposed method are stable over varying true network densities.

Prior SC Knowledge: Conditional on Ω_G , we construct several types of SC information according to the following schemes, where the SC strengths were generated randomly between (0, 1) and FC is measured by partial correlations. For scenario *MI*, we specify that 50% of those edges with strong FC (partial correlation $> .06$) also have strong SC (> 0.7), while 25% of those edges with strong FC have moderate SC (0.3 – 0.7), and the remaining 25% have weak SC (0 – 0.3). For scenario *MII*, the proportion of edges that have strong FC, coupled with strong SC, moderate SC, and weak SC, are 30%, 35% and 35% respectively. For each of the scenarios

MI and MII , we also consider two levels of misspecification of the SC information, which are denoted as $MI(a), MI(b)$, and $MII(a), MII(b)$, respectively. For $MI(a)/MII(a)$, we specify that 10% of those edges with zero FC have non-zero SC, while for $MI(b)/MII(b)$, we fix 20% of those edges with zero FC to have non-zero SC. All the other edges with zero or weak FC are assumed to have small SC, while remaining edges with moderate FC have non-zero SC strengths. We note that edges having zero FC but non-zero SC represent potential misspecification of anatomical knowledge, based on the notion that strong SC typically underlies robust non-zero FC (Shen et al., 2015; Kemmer et al., 2017). For the non-Gaussian data based on the ICA model, we generate the SC information as in scenario MI, and consider varying levels of misspecification.

Comparison Metrics: To assess the performance under different approaches, we compute the area under the curve (AUC), which is a measure of the estimated sensitivity versus specificity over different network sparsity levels. Sensitivity is computed as $TP/(TP + FN)$, while specificity is defined as $TN/(TN + FP)$, where TP, TN, FP, FN denote the number of edges that are true positives, true negatives, false positives and false negatives, respectively. These measures are derived by comparing the true binary network structure, \mathcal{G} , to $\hat{\mathcal{G}}$ which contains edges corresponding to non-zero elements in the estimate $\hat{\Omega}$. To evaluate the point estimate of the network obtained under the Bayesian information criteria (BIC), we compute the Matthews Correlation Coefficient (MCC), which is a scalar measure combining sensitivity and specificity and is defined as $MCC = \frac{TP \times TN - FP \times FN}{\sqrt{(TP + FP)(TP + FN)(TN + FP)(TN + FN)}}$ (Matthews, 1975; Wang, 2012). We compute the relative L_1 norm error, $(|\hat{\Omega} - \Omega|_1)/|\Omega|_1$, to assess accuracy in estimating the precision matrix encapsulating the FC strengths. Since brain networks are also often evaluated in terms of summary statistics reflecting network organization, we evaluate the accuracy in estimating the global efficiency which is a commonly used measure for global integration of brain connectivity.

Since the true precision matrix is dense for the ICA generated data, we consider two alternate measures of performance under this scenario—the inverse error (Padmanabhan et al., 2016) and Kullback-Leibler divergence (Hinne et al., 2014). The inverse error captures the accuracy of the estimated precision matrix, and is defined as $\|\Omega_0^{-1}\hat{\Omega} - I\|_F$ where Ω_0 is the true inverse covariance matrix, $\hat{\Omega}$ is the estimated precision matrix, I is a $p \times p$ identity matrix, and $\|\cdot\|_F$ is the Frobenius norm. On the other hand, the Kullback-Leibler divergence is defined as $1/\{2(\log(2))\}[\log(\det(\Omega_0)/\det(\hat{\Omega})) + \text{tr}(\hat{\Omega}\Omega_0^{-1}) - p]$, with a larger divergence corresponding to a poorer fit.

3.3.1.2 Results

The results under siGGM and SC naive approaches under $MI(a)$ are presented in Table B.1, and Table 3.1 displays results for SC informed approaches under various levels of misspecification. Table B.1 illustrates that either the proposed siGGM approach, or the variant of the proposed approach with no prior knowledge ($\eta = 0$), have the lowest bias in estimating the global efficiency (Eglob) across all network sizes ($p=100, 200$). Moreover, the proposed approach always has higher MCC and AUC values, and lower L_1 error norm compared to all other SC naive approaches. These results demonstrate the advantages of using structural knowledge to guide network estimation.

When the misspecification levels are varied, Table 3.1 illustrates that the proposed method has a consistently lower bias in estimating the global efficiency for both 10% and 20% misspecification levels, compared to alternative SC informed approaches. Moreover, while the G-Wishart approach may have a higher MCC for $p = 100$ when the misspecification level is 10% (cases $MI(a)$ and $MII(a)$), the proposed method has a comparable or higher MCC for 20% misspecification levels (cases $MI(b)$ and $MII(b)$). Moreover under $p = 200$, the MCC under the proposed approach is the highest for small-world and scale-free networks, and comparable to the G-Wishart

method for the Erdos-Renyi network. We also note that while aGlasso often has the lowest MCC values, it may sometimes yield a higher AUC under small-world and scale-free networks for $p = 100$. However, the proposed approach is shown to have the highest AUC for $p = 200$ for all scenarios, highlighting the advantages of incorporating prior knowledge in a flexible manner in higher dimensions. Finally, siGGM consistently has the lowest L_1 error in estimating the precision matrix across all networks and dimensions. The above results illustrate a robust performance of the proposed method for $p = 100$ and a superior performance for $p = 200$ under the small-world and scale-free networks, which closely resemble brain networks encountered in practical applications.

Although Table 3.1 provides some idea about the relative performance under misspecification, it is of interest to look at the effects of misspecification in more detail. Hence, we examined the AUC and L_1 error values as the misspecification level was gradually increased from 4% to 50%, under different networks for $p = 100$. The results, presented in Figure 3.4, illustrate that the proposed method has a significantly higher AUC under the Erdos-Renyi network across all misspecification levels, and registers a significantly higher AUC for larger misspecification levels under the small-world network. The differences in AUC between siGGM and aGlasso are not significant for $p = 100$, but we note that siGGM has a higher AUC for larger dimensions ($p = 200$) as in Table 3.1. For all networks, the proposed method is seen to have a significantly lower L_1 error across all misspecification levels, while the error under the G-Wishart increases sharply as the misspecification level is increased.

The superior performance of siGGM relative to SC informed and SC naive approaches is also observed in the ICA based simulated data exhibiting the spatiotemporal dynamics of the BOLD signals. In Figure 3.5, siGGM exhibits the largest or comparable AUC across all levels of SC misspecification, with significant improvements in AUC for $p = 100$. Our method also has significantly lower inverse error

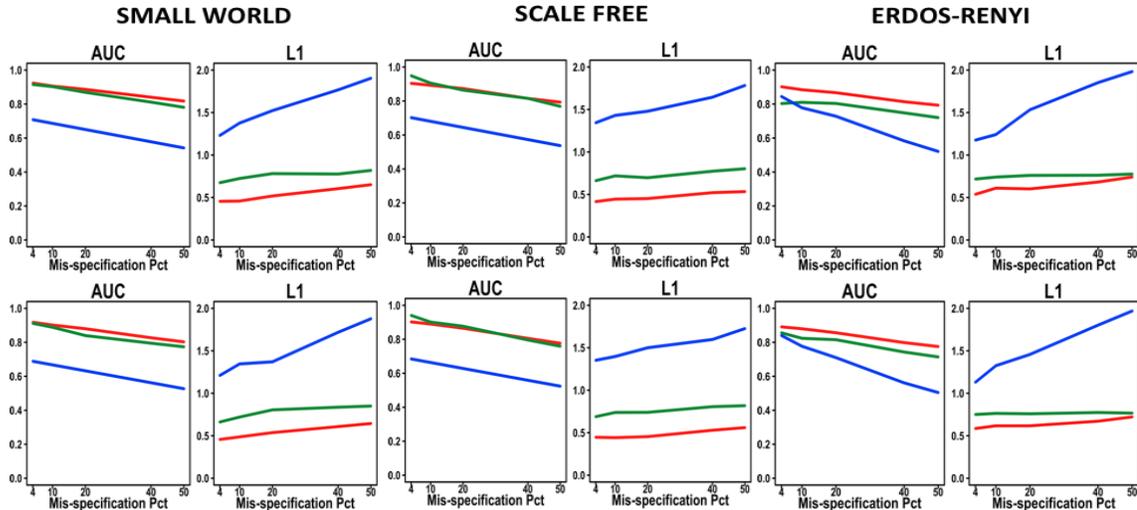


Figure 3.4: Comparison of siGGM (red), G-Wishart (blue), and aGlasso (green) simulation results for different network structures with $p=100$ regions under scenario MI (top row) and MII (bottom row). Each panel displays the AUC or L1 relative error as the percentage of conditionally independent edges with non-zero anatomical connectivity increases. The AUC is significantly higher under siGGM for higher misspecification levels under the small-world network and for all misspecification levels for the Erdos-Renyi network. The L1 error is significantly lower under siGGM for non-trivial misspecification levels and all networks.

and Kullback-Leibler divergence scores across misspecification levels. On the other hand, the AUC for aGlasso drops sharply for $p = 100$ for higher misspecification levels. Moreover, G-Wishart's strict adherence to the SC information contributes to poor detection and estimation of edges, with the AUC declining sharply and the inverse error and Kullback-Leibler divergence increasing steeply as the misspecification proportion increases. The above discussions clearly illustrate the ability of siGGM to recover the true network of connection strengths via flexible incorporation of anatomical information under both Gaussian and non-Gaussian settings, with a robust performance under varying SC misspecification levels.

Finally, we note that the siGGM can be implemented fairly quickly. On a 2.5Gz Intel Core i5 processor, the procedure estimates the optimal graph structure in approximately three seconds for $p=40$, twenty seconds for $p=100$, and approximately four minutes for $p=200$. While these computation times are slightly slower compared to generic graphical modeling approaches naive to anatomical knowledge, the overall

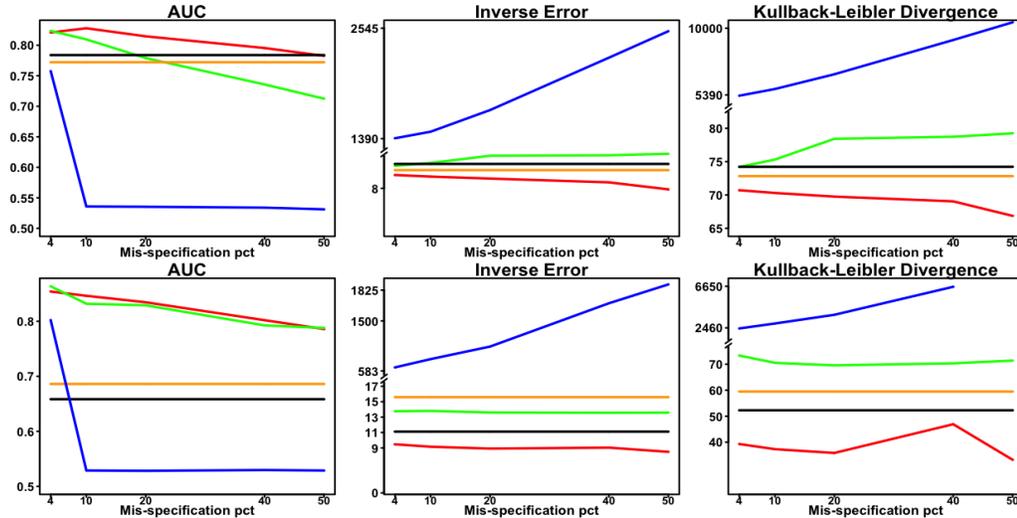


Figure 3.5: Results from the simulation studies for non-Gaussian ICA data for 100 (*top row*) and 200 (*bottom row*) brain regions. We compare the performance of siGGM(**red**), aGlasso (**green**), G-Wishart (**blue**), glasso(**black**), and Space (**orange**) with respect to AUC, inverse error, and Kullback-Leibler divergence. The AUC under siGGM is significantly higher for most misspecification levels under $p = 100$, while the inverse error and Kullback-Leibler divergence is significantly lower under siGGM under all misspecification levels, for $p = 100$ and $p = 200$.

computation is sufficiently quick and feasible for practical implementation in whole brain connectome analysis. Moreover, the siGGM approach converges fairly quickly over a wide range of simulation scenarios, as illustrated in Figure 3.6.

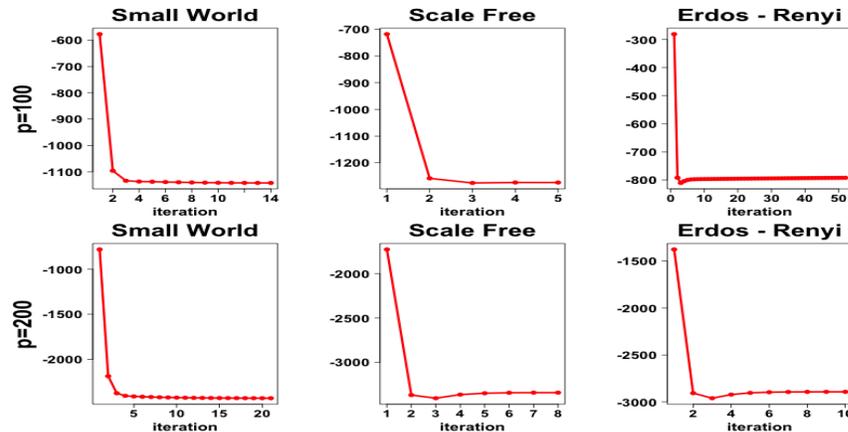


Figure 3.6: Convergence of the log posterior likelihood for the structurally informed Gaussian graphical model (siGGM) under the three network structures (small-world, scale-free, Erdos-Renyi) and sizes ($p=100, 200$) investigated. The method typically makes the most substantial improvements within the first three iterations.

Table 3.1: Performance of SC informed methods on simulated network data with $p=100$ and 200 nodes. Eglob is the bias in global efficiency.

	p=100				p=200			
	Eglob	MCC	AUC	L1	Eglob	MCC	AUC	L1
Small World								
G-Wishart MI(a)	0.120	0.592	0.698	1.345	0.175	0.468	0.865	1.510
G-Wishart MI(b)	0.173	0.447	0.574	1.498	0.224	0.337	0.819	1.801
G-Wishart MII(a)	0.117	0.567	0.676	1.319	0.174	0.447	0.842	1.506
G-Wishart MII(b)	0.171	0.424	0.650	1.492	0.224	0.318	0.797	1.788
aGlasso MI(a)	-0.145	0.522	0.903	0.724	-0.234	0.454	0.800	0.773
aGlasso MI(b)	-0.240	0.407	0.869	0.782	-0.258	0.396	0.788	0.801
aGlasso MII(a)	0.183	0.477	0.889	0.720	-0.221	0.439	0.735	0.770
aGlasso MII(b)	-0.279	0.364	0.840	0.805	-0.272	0.376	0.775	0.814
siGGM MI(a)	0.078	0.590	0.889	0.478	0.121	0.526	0.906	0.532
siGGM MI(b)	0.112	0.500	0.880	0.531	0.166	0.419	0.875	0.603
siGGM MII(a)	0.075	0.576	0.879	0.486	0.125	0.514	0.901	0.563
siGGM MII(b)	0.122	0.490	0.846	0.547	0.169	0.406	0.870	0.633
Scale Free								
G-Wishart MI(a)	0.104	0.590	0.695	1.410	0.146	0.471	0.864	1.503
G-Wishart MI(b)	0.159	0.446	0.671	1.483	0.196	0.339	0.822	1.583
G-Wishart MII(a)	0.102	0.568	0.675	1.397	0.145	0.452	0.846	1.442
G-Wishart MII(b)	0.156	0.424	0.650	1.481	0.196	0.323	0.801	1.604
aGlasso MI(a)	0.223	0.509	0.905	0.719	-0.311	0.375	0.701	0.730
aGlasso MI(b)	-0.195	0.460	0.864	0.697	-0.312	0.336	0.685	0.736
aGlasso MII(a)	-0.256	0.442	0.901	0.739	-0.333	0.351	0.690	0.739
aGlasso MII(b)	-0.252	0.404	0.877	0.740	-0.303	0.320	0.658	0.746
siGGM MI(a)	0.054	0.562	0.853	0.428	-0.075	0.473	0.868	0.442
siGGM MI(b)	0.093	0.467	0.822	0.457	0.131	0.359	0.843	0.492
siGGM MII(a)	0.061	0.552	0.845	0.447	0.078	0.459	0.865	0.457
siGGM MII(b)	0.099	0.451	0.812	0.469	0.132	0.346	0.839	0.523
Erdos-Renyi								
G-Wishart MI(a)	0.174	0.505	0.821	1.300	0.140	0.519	0.860	1.572
G-Wishart MI(b)	0.239	0.368	0.765	1.491	0.187	0.380	0.821	1.976
G-Wishart MII(a)	0.171	0.483	0.807	1.277	0.139	0.501	0.838	1.560
G-Wishart MII(b)	0.237	0.349	0.747	1.464	0.186	0.365	0.805	1.956
aGlasso MI(a)	-0.335	0.404	0.810	0.741	-0.424	0.171	0.596	0.712
aGlasso MI(b)	-0.339	0.353	0.803	0.761	-0.423	0.176	0.600	0.708
aGlasso MII(a)	-0.362	0.333	0.824	0.764	-0.424	0.162	0.631	0.707
aGlasso MII(b)	-0.360	0.318	0.815	0.759	-0.424	0.147	0.648	0.709
siGGM MI(a)	0.124	0.442	0.861	0.624	0.049	0.514	0.862	0.689
siGGM MI(b)	0.186	0.363	0.838	0.646	0.110	0.380	0.826	0.697
siGGM MII(a)	0.122	0.421	0.852	0.638	0.050	0.499	0.846	0.690
siGGM MII(b)	0.171	0.344	0.825	0.665	0.106	0.367	0.810	0.706

3.3.2 PNC Data Application

Existing literature has examined various neural substrates for age related changes using structural and functional neuroimaging (Gur et al., 2012; Shaw et al., 2008; Raznahan et al., 2011). Moreover, gender differences have been extensively documented in behavioral measures (Halpern et al., 2007; Hines, 2010), structural neuroimaging (Lenroot et al., 2007), and functional imaging measures (Lenroot and Giedd, 2010). However, gender related differences in the developmental trajectory of the brain functional network from childhood to adolescence are still not understood well (Gur et al., 2012), and further, limited attempts have been made to investigate such differences by fusing functional and structural neuroimaging data. We use resting state fMRI and DTI data from the Philadelphia Neurodevelopment Cohort (PNC) study to obtain preliminary answers to these questions. After estimating brain functional connectivity based on SC knowledge, we examine FC differences between boys and girls across different age groups.

We perform the analysis separately for each gender within the three age groups 8-12 (pre-teen), 13-17 (teen), 18-21 (young adult), where each age group contains approximately 9 to 12 individuals, and is constructed as in Ingahalikar et al. (2014). All subjects are right-handed, and physically and mentally healthy, enabling a fair comparison between the groups. In addition to assessing gender based network differences, we also perform a secondary analysis to assess our method’s ability to reliably estimate functional networks. For this analysis, we split each subjects’ resting state fMRI time series into two equally sized scanning sessions (60 scans each) and calculate the intraclass correlation coefficient or ICC (refer to equation (B.4.0.1)) for seven network metrics which are widely used to summarize brain networks. The network metrics include clustering coefficient, characteristic path length, local efficiency, global efficiency, modularity, hierarchy, and degree, and they were calculated with the Matlab toolboxes Brain Connectivity Toolbox (Rubinov and Sporns, 2010b) and

GRETNA (Wang et al., 2015b). Mathematical definitions for each metric are presented in Appendix B.6. We note that the ICC is a commonly used measure designed to assess the similarity of network estimates across scanning sessions (Braun et al., 2012; Choe et al., 2017; Niu et al., 2013), and it is usually derived by calculating the proportion of the total variation attributed to variability across scanning sessions. Thus, small variation across sessions relative to variation between individuals produces high ICC values, indicating strong reproducibility.

3.3.2.1 Data preprocessing

For details on data cleaning and acquisition, please see 1.3.2

3.3.2.2 Results

In Figure 3.7, we see that for males (top left panel) and females (top right panel), the estimated association between FC and SC along structurally connected regions is largest for aGlasso, indicating close adherence with the anatomical information. However, the literature suggests that FC is not fully explained by direct structural connections and thus a large association is not realistic. This strong dependence on SC contributes to aGlasso’s inferior correlation with the empirical FC (*bottom row* in Figure 3.7). On the other hand, siGGM’s flexible incorporation of SC while accounting for non-anatomical sources of variation produces desirable results since it adheres to the SC information while maintaining an association with the empirical FC observed in SC-naive approaches. Additionally, we also discover that siGGM leads to larger shrinkage and smaller variance for conditional dependencies between anatomically isolated brain regions compared to the generic graphical lasso without prior knowledge. This yields a smaller number of functional connections between anatomically disconnected ROIs.

For network analysis, we classify each ROI into one of eight functional modules

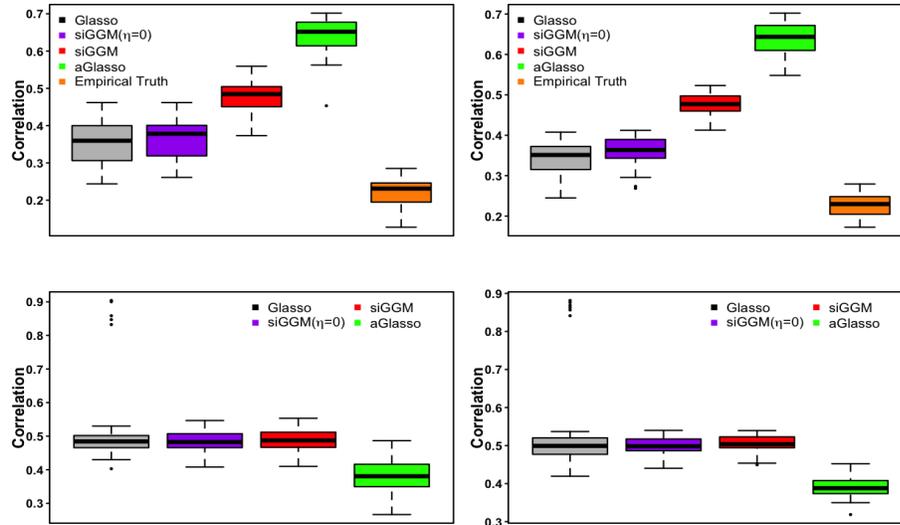


Figure 3.7: Correlation between structural and functional connectivity males(left column) and females (right column). The top row displays the correlation between FC estimates and SC for regions structurally connected, and the bottom row displays the correlation between the estimated FC and the empirical FC for structurally connected regions. We note that aGlasso closely adheres to the SC information due to its shrinkage parameter specification.

corresponding to resting state networks as defined in Smith et. al (2009) . These functional modules include a medial visual network, an occipital pole and lateral visual network (“VIS”, 18 nodes), the default mode network (“DMN”, 8 nodes), a sensorimotor network (“SM”, 9 nodes), an auditory network (“AUD”, 10 nodes), an executive control network (“EC”, 19 nodes), right and left frontoparietal modules (“FPR” and “FPL”, 11 and 10 nodes, respectively) and an unknown module containing unassigned nodes (“UNK”, 5 nodes). Figure 3.8 shows that males and females have similar connectivity patterns with primarily positive connections within functional modules. Further comparisons of male and female brain networks within each age group reveals that consistent connections across age groups persist within module while inconsistent connections mainly exist between modules. After standardizing by the number of nodes in each module, the SM and AUD were found to be the two most highly connected functional modules in males and females across all age groups. Figure 3.9 (A) illustrates the similarity in network architecture for males and females

with all metrics having non-significant differences across genders (except for local efficiency for teens and young adults), which implies shared patterns in brain organization across gender and age groups. Figure 3.9 (B) illustrates that males exhibit greater (but non-significant) between module but smaller within-module connectivity differences in teens and young adults. These findings are supported by previous work and has been linked to variations in emotional identification and spatial cognitive tasks (Satterthwaite et al., 2016).

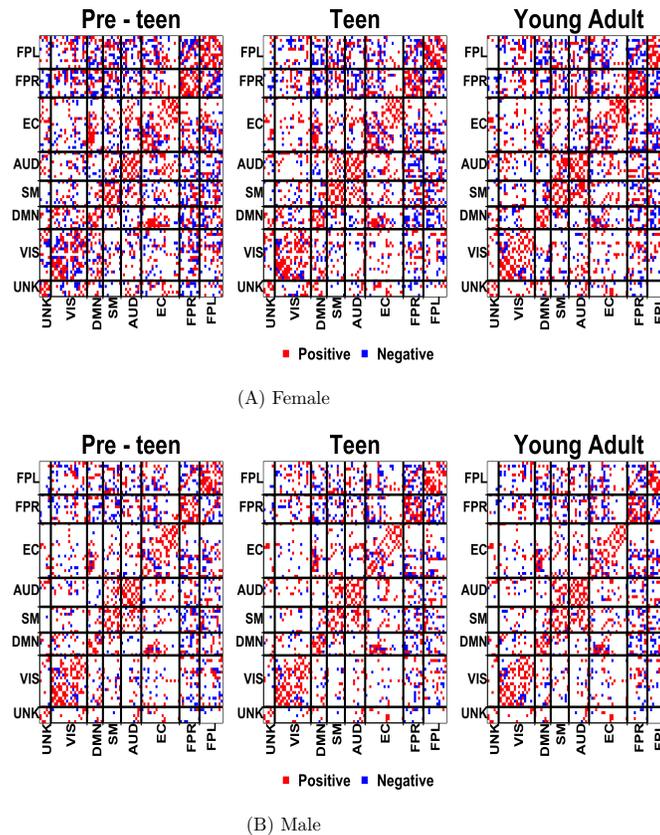


Figure 3.8: Network estimates for females (*top row*) and males (*bottom row*) in each of the age ranges, illustrating those connections corresponding to absolute partial correlations > 0.005 . While both sexes have similar network structures across the three age groups and have network densities close to 13%, female networks exhibit slightly increased connectivity relative to the networks of males.

As a second level of the analysis, we are also interested in the distribution of differentially weighted edges between males and females within each of the eight functional modules. Differentially weighted edges were identified as connections for which the

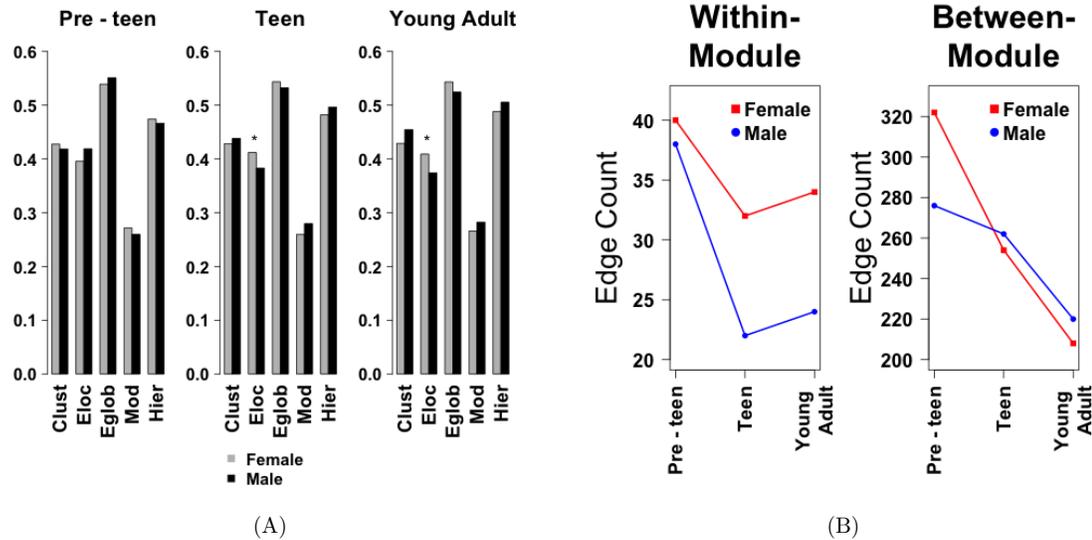


Figure 3.9: Topological features of estimated networks in males and females across the three age groups. (A) displays five network properties—Clust (clustering coefficient), Eloc (local efficiency), Eglob (global efficiency), Mod (modularity), Hier (hierarchy)—averaged over the respective gender and age group; (B) displays differentially weighted edges within- and between- module stratified by gender. In teens and young adults, females have more within module connections and fewer between module connections than males. In (A), the local efficiency is statistically significantly different between males and females in teens and young adults ($p < .05$). In (B), there are no significant results at the .05 level of significance.

FC strength was significantly different between genders under a permutation test. To evaluate if the number of differentially weighted edges within and between modules occur more often than allowed by chance, we define a goodness of fit measure (equation (B.3.0.1) in Appendix A) which represents the deviation between observed and expected numbers of differentially weighted edges for each module block, standardized by the expected number. This measure captures whether a given module block has unusually high or low occurrence of differentially weighted edges and enables us to identify modules with the most pronounced differences across gender. From the results presented in Table 3.2, we discover statistically significant differences in the number of differentially weighted edges occurring in the executive control (EC) module in pre-teens and young adults, which is supported by previous results on gender related differences in the EC (Hyde, 1981; Mansouri et al., 2016). Table 3.2 also suggests that gender based differences attenuate with development, with the largest

number of differentially weighted edges in the pre-teen group (377) and the smallest in the young adult group (272). We also find the differentially weighted edge between the `cingulum_ant_L` in the EC and `parietal_inf_L` in the DMN exists in pre-teens, teens, and young adults, which suggests consistent gender based differences during the developmental phase. These regions are known to have brain volume differences between males and females which may point to subtle cognitive variations (Frederikse et al., 1999; Ruigrok et al., 2014).

A major challenge in resting state connectivity studies is to ensure reproducibility of the findings (Griffanti et al., 2016). We demonstrate that appropriately incorporating anatomical connectivity information leads to stable topological features of estimated networks across scanning sessions. Figure 3.10 displays the ICC of seven network metrics under different approaches, where the details for computing the ICC are outlined in equation (B.4.0.1) in Appendix A. It is clear that the proposed siGGM produces estimates that have notably larger ICC measures for all the network metrics compared to all the other approaches considered. The reproducibility under the proposed approach is substantial for the clustering coefficient, global efficiency, and degree, and is moderate for all the other metrics. Moreover, it is reassuring to see that these three metrics with the highest ICC values under the proposed approach have been shown to be the most reproducible network metrics in independent studies (Niu et al., 2013; Telesford et al., 2010; Wang et al., 2011). In contrast, reproducibility is barely moderate under aGlasso for most metrics and weak under SC naive approaches. We note that siGGM has significantly higher ICC ($p < .001$) than Glasso and Space with respect to all network metrics, and a significantly higher ICC compared to aGlasso for all network metrics ($p < .004$) except local efficiency and degree ($p > .05$). These findings highlight the benefits of incorporating anatomical information in a flexible manner. Although not presented, we note that the G-Wishart approach leads to an unrealistic ICC value of one in all cases, which is starkly different than the reliability

Table 3.2: Within- and between- module differences in functional connectivity between males and females. Bolded values with an asterisk indicate statistically significant modules at the .05 level of significance (FDR correction for multiplicity) and values within parenthesis are the number of differentially weighted edges where the average edge weight is larger in males than females. The total number of differentially weighted edges decrease across the age groups (pre-teen 377 DWE, teen 312 DWE, young adult 272 DWE).

Pre-Teen								
	Unknown	Visual	DMN	SM	Aud	EC	FP left	FP right
Unknown	0(0)							
Visual	8(3)	24(12)						
DMN	7(3)	15(7)	0					
SM	3(1)	14(7)	3(1)	2(2)				
Aud	2(1)	17(8)	13(8)	5(2)	4(0)			
EC	6(5)	32(17)	9(4)	20(9)	21(13)	42*(26)		
FP left	2(1)	14(7)	8(5)	4(4)	10(5)	13(6)	0(0)	
FP right	5(2)	18(9)	9(6)	11(8)	11(8)	16(10)	3(1)	6(0)
Teen								
	Unknown	Visual	DMN	SM	Aud	EC	FP left	FP right
Unknown	2(2)							
Visual	6(3)	22(14)						
DMN	3(2)	8(3)	2(0)					
SM	1(1)	11(8)	9(4)	4(2)				
Aud	3(0)	7(3)	7(4)	13(6)	4(2)			
EC	8(5)	22(11)	13(6)	12(8)	12(5)	18(10)		
FP left	3(1)	16(9)	7(3)	6(4)	9(4)	13(3)	0(0)	
FP right	2(1)	12(6)	9(4)	5(5)	15(5)	18(10)	8(3)	2(2)
Young Adult								
	Unknown	Visual	DMN	SM	Aud	EC	FP left	FP right
Unknown	2(0)							
Visual	1(1)	10(4)						
DMN	1(1)	14(4)	2(0)					
SM	1(0)	10(3)	9(4)	4(4)				
Aud	1(1)	15(5)	5(2)	5(3)	8(4)			
EC	5(4)	16(8)	10(3)	12(4)	12(6)	26*(16)		
FP left	1(1)	7(4)	5(3)	7(4)	11(8)	9(5)	6(6)	
FP right	1(1)	8(4)	7(2)	9(5)	7(6)	17(7)	8(5)	0(0)

values reported in previous studies (Welton et al., 2015). The perfect reliability is due to the fact that G-Wishart relies entirely on the SC information for specifying

the functional network structure, resulting in the exact same network for both the sessions. Hence the reproducibility results under G-Wishart are not comparable.

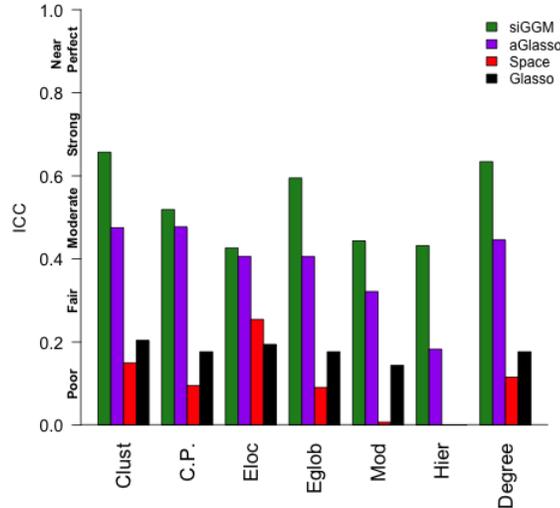


Figure 3.10: Reliability of network metrics across scanning sessions for siGGM, aGlasso, *Space*, and *glasso*. We estimate seven network attributes (clustering coefficient (Clus.), characteristic path length (C.P.), local efficiency (Eloc), global efficiency (Eglob), modularity (Mod), hierarchy (Hier), and degree (Degree)) and report $ICC(3,1)$ for all subjects. ICC values are classified according to the agreement scale $0 < ICC \leq .2$ (poor), $.2 < ICC \leq .4$ (fair), $.4 < ICC \leq .6$ (moderate), $.6 < ICC \leq .8$ (strong), and $.8 < ICC \leq 1$ (near perfect) as suggested by Telesford et. al (2010). siGGM has significantly higher ICC ($p_i.001$) than Glasso and Space with respect to all network metrics, and a significantly higher ICC compared to aGlasso for all network metrics ($p_i.004$) except local efficiency and degree ($p_i.05$). Generally, anatomically informed FC estimates produce more reliable networks than SC naive methods.

3.4 Conclusion

In this chapter, we propose a novel Bayesian method that allows brain structure to guide the estimation of functional connectivity. Our siGGM flexibly incorporates *a priori* known anatomical connectivity information, bypassing the limitations of existing approaches by accommodating complex structure-function relationships while also permitting unknown sources of variation independent of underlying anatomical structure. Our method is more biologically plausible than existing methods and

outperforms alternative SC-informed and -naive approaches as illustrated via extensive numerical studies. In regard to SC-informed approaches, the advantages of the siGGM become more evident as the misspecification levels for the anatomical knowledge and/or the number of nodes is increased, which is important given limitations of current imaging technologies. In the real data application, our method estimates networks with more reproducible topological features compared to popular competitors. Incorporating structural information not only has the advantages discussed in this chapter, but also can have tremendous influences on network comparison methods presented in the previous chapter as well as methods to predict clinical outcomes from functional connectivity biomarkers. Original source material can be found in *NeuroImage*.

Chapter 4

Semi-parametric Bayesian hierarchical dictionary learning

4.1 Introduction

In chapters two and three, we proposed methods that assume that cortical co-activation is fixed at rest. In this chapter, we develop a method that allows brain coupling patterns to vary across time in the resting state. Critically, we assume that the functional activity patterns are a linear mixture of a finite basis of elements. For this work, we work with resting-state fMRI data, which quantifies cortical activity via fluctuations in the blood oxygen level dependent (BOLD) signal when subjects are at rest. From these signals, networks can be inferred by cross correlating the temporal activity patterns between all pairs of brain regions. In fact, this very simple approach has led to substantial insights on the functional architecture of the human brain (Rogers et al., 2007; Smith et al., 2009b). Further, functional connectivity analyses have elucidated how diseases such as depression (Cheng et al., 2018; Greicius et al., 2007b) and schizophrenia (Lynall et al., 2010) are characterized by irregular functional connectivity in the brain.

Early studies (Biswal et al., 1997; Cordes et al., 2000; Lord et al., 2012) focused exclusively on static functional connectivity, whereby it is assumed that the regions exhibit fixed co-activity. More recently, researchers have observed temporally evolving co-activation patterns between brain regions, indicating the brain does not idle in a fixed state as assumed in static functional connectivity analyses. Consequently, there has been rapid development of methods for estimating dynamic functional connectivity (dFC) in the human brain. Dynamic functional connectivity, defined as “rapid, time-varying changes in functional activation” (Cohen, 2018), relaxes the stationarity assumption and has led to new insights into the architecture of the brain. Furthermore, analyses based on dFC have highlighted how many diseases and disorders can be identified by overuse of specific co-activation patterns. For example, (Damaraju et al., 2014a; Sun et al., n.d.) have shown how schizophrenic populations frequently exhibit globally segregated functional co-activations. The authors explain

that this functional organization is also present in healthy individuals, although at dramatically reduced rates. Other studies have shown that dFC produces significantly higher classification accuracy of disorders as compared to static functional connectivity measures (Cohen, 2018; Jin et al., 2017).

Various methods exist for measuring dFC. Model-free approaches such as sliding window correlations (SWC) (Allen et al., 2014b, 2012) dominate the networking literature. In SWC, pairwise associations between all regional time-series are estimated within a window of fixed length. This window is repeatedly slid forward one time step until the entire scan length has been covered. Despite its widespread use, SWC is very sensitive to noise and physiological artifacts (Handwerker et al., 2012). Additionally, the width of window is difficult to select and often produces non-zero connectivity between uncorrelated regions (Lindquist et al., 2014). Further, SWC is ill-equipped to handle abrupt changes in functional connectivity and applies equal weight to all time points in the window (Lindquist et al., 2014). In total, these limitations can lead to poor estimates of dFC. Co-activation patterns (Liu and Duyn, 2013) and multiplication of temporal derivatives (Shine et al., 2015) are two other model-free approaches that cluster regions based on similar changes to the BOLD signal across consecutive time points.

Model-based approaches, such as psychophysiological interaction (Friston et al., 1997) and dynamical conditional correlations (DCC) (Allen et al., 2014b, 2012), do a better job of minimizing the influence of noise artifacts as well as relating time-varying changes in functional connectivity to behavioral outcomes. DCC, proposed by Engle (2002), is a multivariate volatility model that treats observed time series as a GARCH process, where the “conditional variance at time t is a linear combination of past values of the conditional variance and of the square of the process” (Lindquist et al., 2014). Lebo and Box-Steffensmeier (2008) show that DCC effectively estimates time-varying variances and Choe et al. (2017) demonstrate its superior performance in

separating true signal from noise artifacts. Unfortunately, DCC is prohibitively slow to use in practice. Hidden Markov models (Vidaurre et al., 2017; Woolrich et al., 2013) are another approach for estimating the dynamics in BOLD signals. While this approach circumvents well documented limitations of SWC (i.e. selection of the window length), it is computationally demanding and scales poorly as the number of brain regions increases (Rydén et al., 2008).

Recent findings suggest that dynamics in functional activity detected by these methods are the result of the brain traversing a fixed set of cognitive states (Cohen, 2018; Damaraju et al., 2014b). These cognitive states are distinguished by functional organization of the brain regions. Smith et al. (2012a) show that the latent cognitive states are spatially overlapping but temporally independent. Hutchison and Morton (2015) and Shine et al. (2016) show that a relatively small number of states adequately captures dFC in resting-state and task-based fMRI studies as well as across age and mental health condition. Cohen (2018) suggest that it is not the form of the states that differ across populations, but rather the frequency and form of transitions between them that differentiates healthy and affected populations. For many of the methods previously discussed, k-means clustering is regularly applied to dFC estimates to uncover latent states producing the dynamics. Additionally, principal component analysis has been applied to SWC to detect the “eigenconnectivities” which are analogous to the latent brain states (Leonardi et al., 2013). Hidden Markov models estimate the latent network structures as the covariance matrices governing the latent emission distributions and also provide information about the transitions between states.

Critically, these approaches assume the brain occupies exactly one state at any point in time. Despite the appeal of this simplification, recent investigations suggest that the human brain supports the concurrent activity of multiple states. Leonardi et al. (2014) show that FC patterns are a combination of those observed during

individual tasks rather than rapid changes between the individual states. Although at rest, the brain actively recruits multiple cognitive states in order to generate and maintain predictions about forthcoming external stimuli (Deco and Corbetta, 2011) as well as moderate internal activities such as day dreaming, free association, stream of consciousness and inner rehearsal (Ghosh et al., 2008). Further, there is evidence that the latent states evolve continuously throughout the scan duration, rather than discretely as enforced in hitherto discussed methods (Andersen et al., 2018; Smith et al., 2012b).

We pursue dictionary learning (DL) methods to simultaneously explore the dynamics of rs-fMRI and classify individuals based on usage of the latent states. Originally proposed by Olshausen and Field (1997) to examine the neurobiological implications of sparse coding, DL has been widely applied in the engineering community for reconstructing and classifying images (Aharon et al., 2006; Naumova and Schnass, 2018; Zhou et al., 2012). Recently, DL methods have successfully decomposed complex spatio-temporal relationships in neuroimaging data into linear combinations of a finite set of basis elements (Dohmatob et al., 2016; Eavani et al., 2012; Yuan et al., 2017). Furthermore, the framework permits investigation of subpopulations in heterogeneous diseases such as posttraumatic stress disorder (PTSD). While DL has been used in the neuroimaging community to confirm the dynamics of BOLD signals, this is the first attempt to use this information to classify subjects to the best of our knowledge.

We propose a hierarchical Bayesian dictionary learning method that permits joint activation of latent brain states. A Gaussian process prior on the mixing coefficients imposes smoothly varying activation of states as a function of time. Our goal is twofold: (1) recover the common set of latent states and (2) classify the subjects based on utilization of the states. A mixture of Gaussian processes on the mixing coefficients imposes disease-specific smoothness on the states' time courses for each

subject. The mixture distribution allows us to exploit findings in the literature that mental disorders can be detected based upon utilization of brain states that are common across healthy and diseased populations.

In what follows, we elaborate on the proposed method. In section 4.2, we discuss the model and provide details on parameter estimation. In section 4.3, we report the performance of our method on synthetic data and resting-state fMRI from a study of posttraumatic stress disorder (PTSD). We conclude with a brief discussion of the model and suggest future directions in section 4.4.

4.2 Methodology

4.2.1 Dictionary learning

We pursue dictionary learning methods to reconstruct observed signals using a finite basis of elements (Mairal et al., 2009). Dictionary learning methods typically solve

$$(\hat{\mathbf{D}}, \hat{\mathbf{A}}) = \arg \min_{(\mathbf{D}, \mathbf{A})} \|\mathbf{X} - \mathbf{D}\mathbf{A}\|_F + \lambda \|\mathbf{A}\|_q \quad \text{s.t. } \|\mathbf{d}_k\|_2 = 1 \quad \forall k \quad (4.1)$$

where \mathbf{X} is the observed signal; the columns of the over-complete dictionary, \mathbf{D} , are the dictionary elements (basis signals); \mathbf{A} is the matrix of mixing coefficients (encoders); $\|\cdot\|_F$ denotes the Frobenius norm; $\|\cdot\|_q$ is the ℓ -0 ($q = 0$) or ℓ -1 ($q = 1$) norm which imposes sparsity on the mixing coefficient matrix; and λ is a regularization parameter. Sparsity on \mathbf{A} ensures that only the most relevant basis signals in the over-complete dictionary are recruited to uniquely reconstruct the observed signals. In image processing, eq. 4.1, and its variants, have successfully denoised and reconstructed images (Sadeghi et al., 2014). Dictionary learning has been successfully applied to neuroimaging data, specifically fMRI where its been shown that DL can detect time-varying changes in functional connectivity (Yaesoubi et al., 2018). Addi-

tionally, Yaesoubi et al. (2018) show that DL has higher accuracy in recovering the latent connectivity states, represented by the basis signals, compared to SWC.

While Bayesian analogues to eq 4.1 exist in the image processing literature (Yang et al., 2016; Zhou et al., 2012), we are not aware of any methods proposed for neuroimaging applications. Although computationally demanding, Bayesian approaches circumvent the difficulty in selecting an appropriate regularization parameter. In theory, knowledge of the sparsity level or noise/residual variance are necessary for selecting the optimal parameter (Yang et al., 2016); in practice, this is rarely ever known and poor “guesses” at the quantity will lead to suboptimal estimates of the dictionary and matrix of mixing coefficients. Furthermore, Bayesian approaches can easily account for variability in the basis dictionary elements. Practically, that allows for variation in the latent brain states observed in individuals. We build upon the sparse Gaussian hierarchical DL model proposed by Yang et al. (2016) where training signals are reconstructed as a linear combination of a fixed dictionary.

4.2.2 Bayesian Dictionary Learning (BayDiL)

Our proposed Bayesian Dictionary Learning (BayDiL) method extends the method of Yang et al. (2016) to incorporate covariates while imposing temporal smoothness on the mixing coefficients. This restriction is critical to our hypothesis that the latent states transition smoothly across time rather than discrete state switching as commonly implemented in the literature.

Further, our BayDiL model can incorporate covariates such as IQ (Song et al., 2008) and personality (Adelstein et al., 2011) that potentially modulate spontaneous functional activity. Our model is

$$\begin{aligned}
\mathbf{X}^i; \mathbf{c}_t^i \Big| \mathbf{D}, \mathbf{A}^i, \mathbf{B}, \sigma_\epsilon &\sim \prod_{t=1}^T N(\mathbf{x}_t^i; \mathbf{D}\mathbf{a}_t^i + \mathbf{B}\mathbf{c}_t^i, \sigma_\epsilon \mathbf{I}_M) \\
\mathbf{D} \Big| \sigma_d &\sim \prod_{k=1}^K N(\mathbf{d}_k; 0, \sigma_d \mathbf{I}_M) \\
\mathbf{B} \Big| \sigma_\beta &\sim \prod_{c=1}^C N(\boldsymbol{\beta}_c; 0, \sigma_\beta \mathbf{I}_M) \\
\mathbf{A}^i \Big| \boldsymbol{\rho} &\sim \sum_{g=1}^G \rho_g \prod_{k=1}^K N(\mathbf{a}_k^i; 0, \mathbf{R}_{\ell_k^g}) \\
\boldsymbol{\rho} &\sim \text{Dirichlet}(1, \dots, 1) \\
\log(\ell_k^g) &\sim N(0, 10) \\
\sigma_\epsilon &\sim \text{InvGamma}(\alpha_{\sigma_\epsilon}, b_{\sigma_\epsilon}), \quad \sigma_\beta \sim \text{InvGamma}(\alpha_{\sigma_\beta}, b_{\sigma_\beta}) \\
\sigma_d &\sim \text{InvGamma}(\alpha_d, b_d)
\end{aligned} \tag{4.2}$$

where $i=1, \dots, n$ subjects, $k=1, \dots, K$ dictionary elements, $t=1, \dots, T$ time points, and $g=1, \dots, G$ groups. For the i^{th} subject, $\mathbf{X}^i \in R^{m \times T}$ is the observed signals at m locations across T timepoints, $\mathbf{A}^i \in R^{K \times T}$ is the matrix of mixing coefficients that linearly combines the dictionary elements to reconstruct the covariate-adjusted regional time series, and \mathbf{c}_t^i is a p -dimensional vector of covariates. As formulated, covariate effects are time-independent (i.e. $\mathbf{c}_t^i = \mathbf{c}^i \forall t$). $\mathbf{D} \in R^{p \times K}$ is a dictionary shared by all subjects, where the k th column, \mathbf{d}_k , is a basis dictionary element and $\mathbf{B} \in R^{p \times T}$ is a matrix of covariate effects. Other parameters in the model are the vector of group inclusion probabilities, $\boldsymbol{\rho} \in R^G$ and the variance parameters σ_μ , σ_ϵ , σ_d , and σ_β .

We impose temporal structure on the mixing coefficient in each component via parametric covariance kernel functions. We make the assumption that observed signals close in the time domain should be strongly correlated in the signal domain. The covariance function for the k th dictionary element, $\mathbf{R}_{\ell_k^g}$ which is based on the set of parameters $(\ell_k^{g'}, \tau_k^{g'})$; $\ell_k > 0$ is the characteristic length scale parameter for the

kth component in the gth, and $\tau_k^{g'} > 0$ is the signal variance. In this work, we set $\tau_k^{g'} = 1 \forall k, g'$ to help uniquely identify the dictionary and mixing coefficients. In all that follows, we utilize the squared exponential kernel function to define $\mathbf{R}_{\ell_k^{g'}}$ where,

$$\mathbf{R}_{\ell_k^{g'}}(t, t') = \tau_k^{g'} \exp\left(-\frac{(t - t')^2}{2\ell_k^{g'^2}}\right) = \exp\left(-\frac{(t - t')^2}{2\ell_k^{g'^2}}\right). \quad (4.3)$$

As formulated, the prior on $\mathbf{A}^i \forall i$ poses various computational issues. If the ith subject's group membership were known, this prior would simplify to product of multivariate normal distributions for the gth group. We exploit this idea by augmenting the data to include group membership as missing information. Thus, if the ith individual is in the g' group, i.e. $z_i = g'$, then the prior on \mathbf{A}^i and the group membership is

$$\begin{aligned} \mathbf{A}^i & \left| \begin{array}{l} z_i = g' \sim \prod_{k=1}^K N(\mathbf{a}_k^i; 0, \mathbf{R}_{\ell_k^{g'}}) \\ z_i \sim \text{Multinomial}(1; \rho_1, \dots, \rho_G) \end{array} \right. \end{aligned} \quad (4.4)$$

4.2.3 Posterior Sampling

Due to the conjugacy of the priors, most parameters can be easily sampled from posterior distributions via a Gibbs sampling routine. For the covariance kernel, we sample the set of hyperparameters, $\{\theta_k\}_{k=1}^K$, and the residual noise of the mixing coefficients, σ_a , using a gradient-based Monte Carlo method. Note that $P(\omega| -)$ is the posterior of ω conditional on the data and all other parameters. The posterior distributions are defined as follows:

- *Sample \mathbf{D}*

We sequentially update the columns of the dictionary, \mathbf{d}_k for $k=1, \dots, K$. Yang et al. (2016) suggest that sequential updates lead to faster convergence compared to simultaneous sampling of all dictionary elements. For the rth iteration,

we update the k th column as follows:

$$P(\mathbf{d}_k | -) \sim N(\Sigma_{d_k} (\sum_{i=1}^n \tilde{\mathbf{X}}^i (a_k^i)'), \Sigma_{d_k}) \text{ for } k=1, \dots, K$$

$$\text{where } \Sigma_{d_k} = (\sum_{i=1}^n \frac{a_k^i (a_k^i)'}{\sigma_\epsilon} + \frac{1}{\sigma_d})^{-1} \mathbf{I}_m.$$

Also note that $\tilde{\mathbf{X}}^i = \mathbf{X}^i - \mathbf{B}(\mathbf{C}_i \otimes \mathbf{1}'_T) - \mathbf{D}_{/k} \mathbf{A}^i$, where $\mathbf{D}_{/k} = [\mathbf{d}_1, \dots, \mathbf{d}_{k-1}, \mathbf{0}, \mathbf{d}_{k+1}, \dots, \mathbf{d}_K]$

- *Sample $\mathbf{A}^i \forall i$*

Since we want to incorporate temporality into the mixing weights for each of the K network components, we update the k^{th} row of \mathbf{A}^i as follows

$$P(a_k^i | -) \sim N(\Sigma_{a_k^i} ((\tilde{\mathbf{X}}^i)' \mathbf{d}_k), \Sigma_{a_k^i})$$

where $\Sigma_{a_k^i} = (\frac{1}{\sigma_\epsilon} \mathbf{d}'_k \mathbf{d}_k \mathbf{I}_T + \mathbf{R}_{\rho^g})^{-1}$, $\tilde{\mathbf{X}}^i = \mathbf{X}^i - \mathbf{B}(\mathbf{C}_i \otimes \mathbf{1}'_T) - \mathbf{D} \mathbf{A}^i$, and the k^{th} row of \mathbf{A}^i is replaced by $\mathbf{0} \in \mathbf{R}^T$.

- *Sample \mathbf{B}*

We jointly update the rows of the coefficient matrix. For row w of \mathbf{B} , the posterior is

$$P(\mathbf{b}_w | -) \sim N(\mathbf{L}_w \mathbf{W}, \mathbf{W}) \text{ where}$$

$$\begin{aligned} - \mathbf{L} &= \sum_{i=1}^n \sum_{t=1}^T \tilde{x}_t^i c'_i \\ - \mathbf{W} &= (T \sum_{i=1}^n \frac{1}{\sigma_\epsilon} c_i c'_i + \frac{1}{\sigma_\beta} \mathbf{I}_p)^{-1} \\ &\text{where } \tilde{x}_t^i = x_t^i - D a_t^i. \end{aligned}$$

- *Sample ρ*

$$P(\rho | -) \sim \text{Dirichlet}(\delta_1^*, \dots, \delta_G^*) \text{ where } \delta_g^* = 1 + n_g \text{ for } n_g = \sum_{i=1}^n \mathbf{I}(z_i = g)$$

- *Sample σ_ϵ*

$$P(\sigma_\epsilon | -) \sim \Gamma^{-1}(\alpha_\epsilon + \frac{nmT}{2}, b_\epsilon + \frac{1}{2} \sum_{i=1}^n \sum_{t=1}^T (x_t^i - D a_t^i - B c_i)' (x_t^i - D a_t^i - B c_i))$$

- *Sample σ_d*

$$P(\sigma_d | -) \sim \Gamma^{-1}(\alpha_d + \frac{pK}{2}, b_d + \frac{1}{2} \sum_{k=1}^K d'_k d_k)$$

- *Sample σ_β*

$$P(\sigma_\beta | -) \sim \Gamma^{-1}(\alpha_\beta + \frac{pC}{2}, b_\beta + \frac{1}{2} \sum_{j=1}^p \beta'_j \beta_j)$$

- Sample $\Theta = \{\ell_1^1, \dots, \ell_1^G, \ell_2^1, \dots, \ell_K^G\}$

We jointly update the covariance kernel parameters using Metropolis-adjusted Langevin algorithm (Besag et al., 1995). We fix the step size as .001 and choose a diagonal matrix for the mass matrix. Partial derivatives with respect to each parameter are derived from the log posterior,

$$g(\Theta) = \log \left(\prod_{i=1}^n \prod_{k=1}^K N(\mathbf{a}_k^i; 0, \mathbf{R}_{\ell_k^{g'}} | z_i = g') \prod_{g=1}^G \prod_{k=1}^K N(\log(\ell_k^g); 0, 10) \right)$$

We sample Θ^* from the proposal distribution, $q(\Theta^*|\Theta) \sim N(\Theta + \frac{h}{2}\nabla g(\Theta), h\mathbf{I})$ for stepsize $h > 0$ and accept it with probability

$$r = \min \left\{ 1, \frac{\pi(\Theta^*)q(\Theta|\Theta^*)}{\pi(\Theta)q(\Theta^*|\Theta)} \right\},$$

where $\pi(\Theta) = \exp(g(\Theta))$ is the target distribution.

- *Sample $z_i \forall i$*

$$P(z_i = g' | -) = \frac{\rho_{g'} \prod_{k=1}^K GP(\mathbf{a}_k^i; 0, \mathbf{R}_{\ell_k^{g'}}) \prod_{t=1}^T N(X_t^i | D, A^i, z_i = g')}{\sum_{g=1}^G \rho_g \prod_{k=1}^K GP(\mathbf{a}_k^i; 0, \mathbf{R}_{\ell_k^g}) \prod_{t=1}^T N(X_t^i | D, A^i, z_i = g)}$$

A full description of our sampling procedure is available in **algorithm 2**, where the posterior distributions are as previously defined.

4.2.4 Subcase: Regional time series as Observed Signal

The flexibility of our approach allows the observed data, \mathbf{X}^i , to take various forms. If sliding windows correlations or DCC are calculated, the observed signal would be composed of the upper diagonal elements of each network estimate over the time domain. For p brain regions, $\mathbf{X}^i \in R^{m \times H}$ where $m = p(p-1)/2$ and H is the number of windows estimated. Due to the limited availability of robust dFC methods, we use the regional time series as the observed information. Critically, we assume that the signal observed at time t is a linear combination of basis signals of dimension p . The

Algorithm 2 Parameter Sampling

Let $\mathbf{X} = [\mathbf{X}^1, \dots, \mathbf{X}^n]$ and $\mathbf{A}^{(r)} = [\mathbf{A}^{1(r)}, \dots, \mathbf{A}^{n(r)}]$ be the collection of observed data and mixing coefficients, respectively, for all subjects in the r th MCMC sample. Similarly, let $\Theta^{(r)} = \{\ell_1^{1(r)}, \dots, \ell_1^{G(r)}, \ell_2^{1(r)}, \dots, \ell_K^{G(r)}\}$.

Assign initial values for all model parameters, $\mathbf{A}^{(0)}$, $\mathbf{D}^{(0)}$, $\mathbf{B}^{(0)}$, $\Theta^{(0)}$, $\sigma_{err}^{(0)}$, $\sigma_d^{(0)}$, $\sigma_\beta^{(0)}$, $\{z_i^{(0)}\}_{i=1}^n$.

For $r=1, \dots, 10000$ repeat the following steps

- 1: Block Gibbs sampling for $\mathbf{B}^{(r)}$

$$\mathbf{b}_\omega^{(r)} \sim P(\mathbf{b}_\omega | \mathbf{X}, \mathbf{A}^{(r-1)}, \mathbf{D}^{(r-1)}, \sigma_\beta^{(r-1)}) \quad \forall \omega = 1, \dots, m$$

- 2: Gibbs sampling for $\sigma_\beta^{(r)}$

$$\sigma_\beta^{(r)} \sim P(\sigma_\beta | \mathbf{B}^{(r)})$$

- 3: Iterative Gibbs sampling for $\mathbf{D}^{(r)}$

$$\mathbf{d}_k^{(r)} \sim P(\mathbf{d}_k | \mathbf{X}, \mathbf{A}^{(r-1)}, \mathbf{D}_{/k}^{(r)}, \mathbf{B}^{(r)}, \sigma_d^{(r-1)}), \quad k = 1, \dots, K$$

where $\mathbf{D}_{/k}^{(r)} = [\mathbf{d}_1^{(r)}, \dots, \mathbf{d}_{k-1}^{(r)}, \mathbf{0}, \mathbf{d}_{k+1}^{(r)}, \dots, \mathbf{d}_K^{(r-1)}]$

- 4: Gibbs sampling for $\sigma_d^{(r)}$

$$\sigma_d^{(r)} \sim P(\sigma_d | \mathbf{D}^{(r)})$$

- 5: Block Gibbs sampling for $\mathbf{A}^{i(r)} \quad \forall i = 1, \dots, n$

$$\mathbf{a}_k^{i(r)} \sim P(\mathbf{a}_k^i | \mathbf{X}, \mathbf{D}^{(r)}, \mathbf{B}^{(r)}, z_i^{(r-1)}, \Theta^{(r-1)}), \quad k = 1, \dots, K$$

- 6: Gibbs update for $\sigma_{err}^{(r)}$

$$\sigma_{err}^{(r)} \sim P(\sigma_{err} | \mathbf{X}, \mathbf{A}^{(r)}, \mathbf{D}^{(r)}, \mathbf{B}^{(r)})$$

- 7: MALA update for $\Theta^{(r)}$

- Use all available updated information: $\mathbf{A}^{(r)}$, $\mathbf{D}^{(r)}$, $\Theta^{(r-1)}$

- 8: Posterior sampling for $z_i = g$

$$z_i^{(r)} \sim P(z_i = g | \mathbf{X}, \mathbf{A}^{(r)}, \mathbf{D}^{(r)}, \mathbf{B}^{(r)}, \sigma_{err}^{(r)}, \sigma_d^{(r)}, \sigma_d^{(r)}, \sigma_\beta^{(r)}, \boldsymbol{\rho}^{(r-1)})$$

- 9: Gibbs sampling for $\boldsymbol{\rho}^{(r)}$

$$\boldsymbol{\rho}^{(r)} \sim P(\boldsymbol{\rho} | \{z_i^{(r)}\}_i)$$

network connectivity associated with this basis vector is proportional to its self-outer product. As argued in Yaesoubi et al. (2018), the outer product adequately captures the variability observed in the data when the time dimension is ignored. Although the authors formulate a model that considers only one basis vector to reconstruct the signal at time t , we consider a linear combination of multiple basis vectors. Since we have developed a Bayesian method, we can use the R MCMC samples to estimate the covariance of the k^{th} dictionary atom as $(1/(R-1)) \sum_{r=1}^R (\mathbf{d}_k^{(r)} - \bar{\mathbf{d}}_k)(\mathbf{d}_k^{(r)} - \bar{\mathbf{d}}_k)'$. Under the assumption that the dictionary elements are independent, the instantaneous network connectivity at time t is the sum of the weighted covariance estimate for each dictionary component. Thus, the instantaneous network connectivity generating the observed signal at time t is at most a rank K , symmetric positive definite network. Our model scales well since we will only need to consider p -dimensional basis vectors rather than m -dimensional basis vectors. We also ignore the effect of covariates, \mathbf{B} , on the observed activity within each region.

4.3 Results and Discussion

4.3.1 Simulations

4.3.1.1 Simulation Setting

We conduct extensive simulations studies to assess the performance of the BayDiL method relative to popular competitors. As previously discussed, sliding windows correlations dominates the literature on dynamic functional connectivity. Consequently, we select a SWC method (Jin et al., 2017) that is representative of the class of methods. The first method, fixed sliding window correlation (fSWC), calculates the network connectivity at time t as the Pearson correlation between all pairs of regions over the previous m timepoints. fSWC represents standard SWC methods

with fixed window widths. We consider window widths from 5 to 55 in intervals of 10. Kmeans clustering is applied to the window estimates to determine the centroids representative of the basis network states. We use the jump statistic to select the optimal number of centers. Finally, we consider an online dictionary learning method (Mairal et al., 2009) which is implemented in the DICTOL Matlab toolbox (<https://github.com/tiepvupsu/DICTOL>). The functions in this toolbox solve eq. (4.1), where we specify the tuning parameter, λ , over a grid of values, $[0, .01, .1, .2, .3, .4, .5, .6, .7, .8, .9, 1]$.

Data generation: We simulate data under various settings to assess the methods' performance. Under the hypothesis that utilization of a common set of networks differentiates disease populations, we simulate the regional time series for each subject as a linear mixture of three basis network components. We enforce group differences by setting by sampling the mixing coefficients for the third basis network from normal distributions where the characteristic scale length parameter is very small in group one and large in group two. Formally, we have $X_{i,t} = \sum_{k=1}^K a_{i,t,k,g} d_k + E_{i,t}$, where $X_{i,t} \in \mathbb{R}^p$ is the observed signal at time t for subject i ; $a_{i,t,k,g} \in \mathbb{R}$ is the mixing weight on the k^{th} dictionary element in the g th group; $d_k \in \mathbb{R}^p$ is the the k th basis signal; and $E_{i,t} \in \mathbb{R}^p$ is random Gaussian noise. For these simulations, we consider $p=20$ regions, $t=1,\dots,120$ timepoints, $k=1,\dots,3$ dictionary elements, $g \in \{1, 2\}$, and $i=1,\dots,100$ subjects. We are interested in the methods' behavior in four scenarios: (1) correct model specification; (2) mis-specified distribution on the mixing coefficients; (3) incorrect number of *a priori* dictionary elements; and (4) spatially overlapping and non-overlapping dictionary elements. We also consider the performance when the mixing coefficients are sampled from four zero-mean Gaussian distributions with the covariance function defined as: (a) squared exponential, (b) rational quadratic, (c) squared exponential times periodic with fixed periodic parameter (SQ), and (d) squared exponential times periodic with fixed square exponential parameter (PER).

We also consider nonparametric mixing coefficients which are represented by phase shifted sinusoidal curves with group constrained periodic effects.

Metrics: BayDiL detects basis atoms producing the observed signal and clusters individuals based upon the utilization of those elements. In order to measure the accuracy of estimated dictionary atoms, we measure the distance between each estimated dictionary atom (\hat{d}_j for $j=1,\dots,K$) and each true atom (d_k for $j=1,\dots,K$) as $\text{dist}(\hat{d}_j, d_k) = 1 - |\hat{d}_j d_k| / (|\hat{d}_j|_2 \times \|d_k\|_2) \in [0, 1]$. As in Yang et al. (2016), we consider \mathbf{d}_k to be fully recovered if the distance from it to any estimated atom is less than .1. We report the recovery success rate as the proportion of true dictionary atoms that are fully recovered. We also seek the clustering performance and measure it with the adjusted rand index available in the *mclust* R package (Scrucca et al., 2016). The adjusted rand index is a measure of the similarity of two partitions of subjects into non-overlapping groups. It is invariant to the well known label switching problem in Bayesian mixture models. For two partitions, $X = (x_1, \dots, x_G)$ and $Y = (y_1, \dots, y_G)$, of the subjects into exactly G disjoint classes, the adjusted rand index is calculated as $\frac{\sum_{i,j} \binom{n_{ij}}{2} - [\sum_i \binom{a_i}{2} \sum_j \binom{b_j}{2}] / \binom{n}{2}}{\frac{1}{2}(\sum_i \binom{a_i}{2} + \sum_j \binom{b_j}{2}) - [\sum_i \binom{a_i}{2} \sum_j \binom{b_j}{2}]}$ where $n = \sum_{i=1}^G |x_i| = \sum_{j=1}^G |y_j|$, $n_{ij} = |x_i \cap y_j|$, $a_i = |x_i|$ for $i=1,\dots,G$, and $b_j = |y_j|$ for $j=1,\dots,G$ for $|v|$ =cardinality of set v . Note that the metric adjusts for the expected number of random pairings via the term $[\sum_i \binom{a_i}{2} \sum_j \binom{b_j}{2}] / \binom{n}{2}$ and takes a value of 1 if the two partitions are exactly the same and zero if the pairings are no better than random chance. It is possible to calculate negative valued-indices, which indicates that the clustering results are worse than random group assignments. In this work, we set negative adjusted rand indexes to zero.

4.3.1.2 Results

We first evaluate the methods' abilities to recover the true basis signals from the regional time series. As seen in figure C.2.1 , BayDiL does an adequate job esti-

mating the basis signals. The network estimates from the fixed sliding windows and dictionary learner do not perform as well. Specifically, fSWC estimate similar basis networks that each contain all true basis signals. This is reflective of SWC network estimates presented in the literature that regularly exhibit spatially consistent basis networks. As alluded to in the introduction, SWC averages the signals within the respective window. Alterations in network structure on shorter timescales than the window length will not be detected and can produce network estimates that are an amalgamation of active subnetworks in the respective window. The traditional dictionary learning method also struggles to detect the basis signals. The DL method’s performance is closely tied to the prespecified tuning parameter, λ . For $\lambda \in [.2, .7]$, the DL recovers on average recovers two-thirds of the dictionary whereas recovery is poor outside this range.

Table 4.1: Dictionary atom recovery success rate for the BayDiL, fixed SWC, and dictionary learner. We compare the recovery success when the mixing coefficients are generated by four parametrically defined covariance functions (Gaussian distributed) and a phase shifted sinusoidal curve.

	S.E.	R.Q.	S.E.×Per(SE)	S.E.×Per(PER)	Sin.
BayDiL	1.00	1.00	.999	1.00	1.00
DL	.306	.055	.556	.139	.000
fSWC	.000	.000	.000	0.00	.000

Given the suggestions that networks are spatially overlapping, we next investigate the performance of BayDiL as basis states increasing share activated regions. In figure 4.1 we see that the BayDiL is able to adequately recover the true dictionary atoms despite the overlap. The DL also exhibits consistent performance. Next, we evaluate the performance of BayDiL when the mixing coefficients are generated from four covariance functions and a class of phase-shifted sinusoidal curves. Figure 4.1 shows the spatial consistency of the estimated basis signals relative to the truth under these four specifications. For the first three parametric kernel specifications, BayDiL classifies subjects with high accuracy even as the dictionary atoms increasingly overlap

(Table 4.2). Our method exhibits poor classification if the group differences arise from changes in the periodicity of the mixing coefficients, evidenced by the results for Squared exponential xPeriodic (PER) and Sin.

Figure 4.1: Comparison of the the BayDiL and DL methods in recovering the basis signals when the signals share 1, 2 and 3 active regions in common.

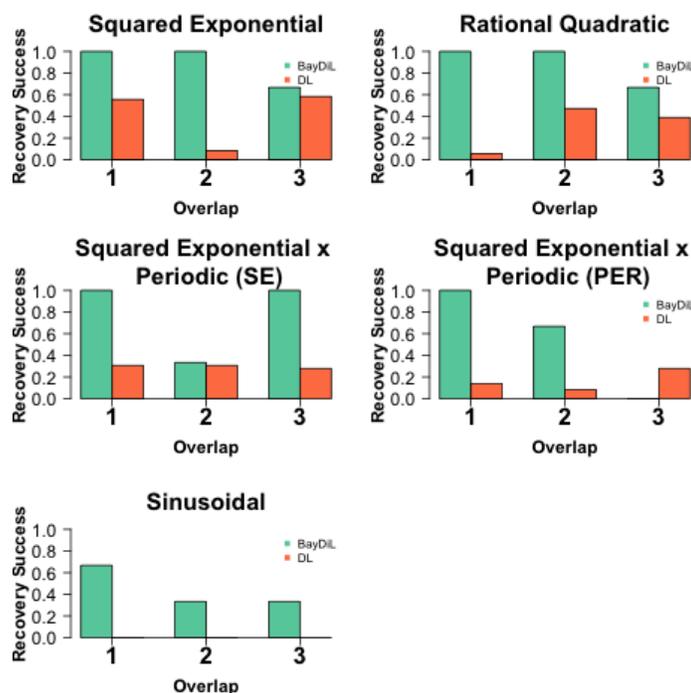


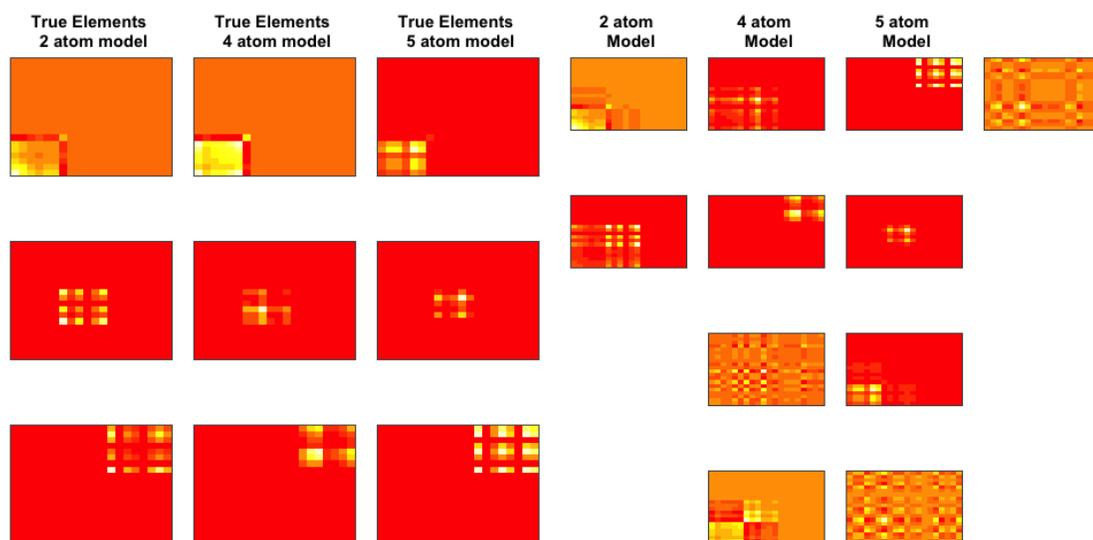
Table 4.2: Average classification accuracy as measured by the adjusted rand index for the BayDiL method.

Overlap	S.E.	R.Q.	S.E. \times Per(SE)	S.E. \times Per(PER)	Sin.
1	1.000(0)	1.00(0)	.999(.005)	0(0)	.167(.059)
2	.996(.043)	1.00(0)	.991(.018)	0(0)	0(0)
3	1.000(0)	1.00(0)	.999(.005)	0(0)	0(0)

The BayDiL exhibits encouraging behavior when the true number of basis signals is misspecified. Specifically, when the true number of atoms is three but we initialize BayDiL with four or five latent components, our method is able to detect the dictionary atoms with high success (.67 and 1, respectively) as well as classify individuals with high probability (1 and 1, respectively). However, the performance is severely

affected if the specified number of atoms underestimates the true number in the data. Figure 4.2 depicts the spatial patterns of the estimated atoms relative to the true atoms. In particular, we note that the four and five component fits produce spatially consistent atoms with the addition of atoms recovering noise in the data. In total, this emphasizes the necessity of allowing the data to inform the BayDiL about the appropriate number of components. Specifying too few components impacts dictionary recovery success as well as classification performance with providing too many produces satisfactory performance at the expense of unnecessary computation time.

Figure 4.2: Basis elements estimates when the BayDiL is initialized with the incorrect number of components.



(A) True atoms (outer product) (B) Estimated atoms (outer product)

BayDiL exhibits other pathological behavior related to the noise level and sample size. As the sample size decreases, the dictionary elements are recovered less successfully. Similarly, as the noise level increases, we note declines in the classification accuracy and successfully recovered basis elements. Finally, we note that the BayDiL successfully classifies patients as the length of the dictionary atoms grows. In the most extreme setting, we assessed performance when the dictionary atoms each contain 5,000 elements. Even in this high dimensional setting, the subjects are classified

with high accuracy and the atoms are all fully recovered.

Finally, we evaluate the performance of the BayDiL when the input are sliding windows estimates. Similar to the k-means clustering spatial maps, we find that the BayDiL merges all true signals into one dictionary atom. All other estimated dictionary atoms are noise artifacts associated with the window length. Interestingly, the BayDiL is unable to linearly separate the signals in each sliding window despite the linearity of the Pearson correlation. We think that this is likely due to the fact that SWC averages the effect of active states within the respective window and the sliding window technique introduces nonlinearities such that the true signals cannot be linearly separated. Although intuitively simple and easy to calculate, we find that SWC are a poor choice of observed signals.

4.3.2 Posttraumatic Stress Disorder Data Application

Please see study participant information and preprocessing details for the rs-fMRI in Chapter 1. Figure C.2.3 presents the spatial locations of the twenty seven brain regions investigated in this analysis.

4.3.2.1 Results

We investigate the rs-fMRI data for three, four, five, six and seven dictionary atoms, assuming there are two subgroups. Allen et al. (2014a) and Yaesoubi et al. (2018) have suggested that this number of latent networks sufficiently explain brain dynamics at rest. We also find that five components are sufficient to reconstruct the observed regional time series and produce interesting clustering results. In figure 4.3, we see that the five atoms are spatially overlapping but exhibit unique patterns. The BayDiL and DL estimates of the covariance of the dictionary atoms are very similar. Note that our estimates are more noisy than those of the DL because we do not impose any constraints on the norm of the dictionary atoms. However, our

covariance structures are more informative because we can account for variation in each atom, unlike the optimization routine of the DL that produces point estimates only. Components two, three, and four demonstrate different connectivity patterns between the subcortical and amygdala regions. Component one exhibits connectivity primarily restricted to the mPFC while component five exhibits integration across the subcortical and amygdala regions but a functional segregated mPFC. The DL method's estimates are spatially consistent with the BayDiL (figure 4.4) and exhibit stronger within module co-activity patterns as well as more inter-modular activity. Further, we note that the DL atoms exhibit consistent patterns across all grid values investigated. Figure 4.5 displays the covariance of the first (A) and fourth (B) dictionary atoms. In figure 4.5(A), we note the strong covariances between amygdala and subcortical regions while figure 4.5(B) shows strong connections between within the mPFC. As we discuss later, the utilization of these structures defines two interesting subpopulations.

Figure 4.3: The covariance matrices of the five basis components estimated by the BayDiL method that linearly mix to produce the observed regional time series for 56 subjects in the PTSD study.

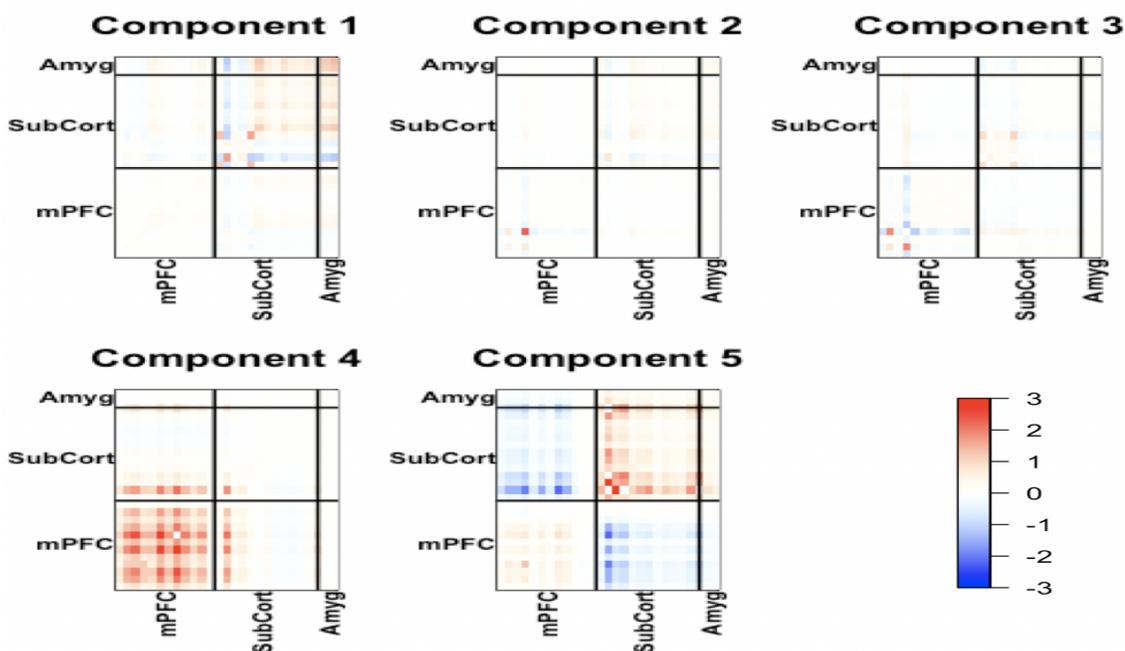


Figure 4.4: The covariance matrices of the five basis components estimated by the DL method ($\lambda = .3$) that linearly mix to produce the observed regional time series for 56 subjects in the PTSD study.

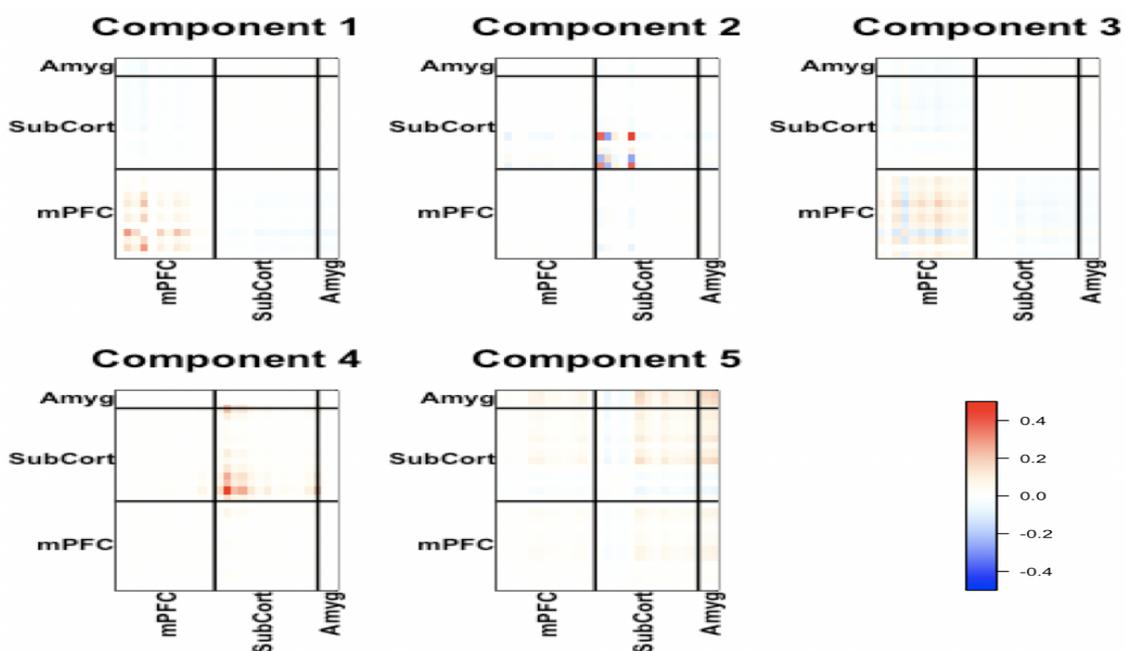
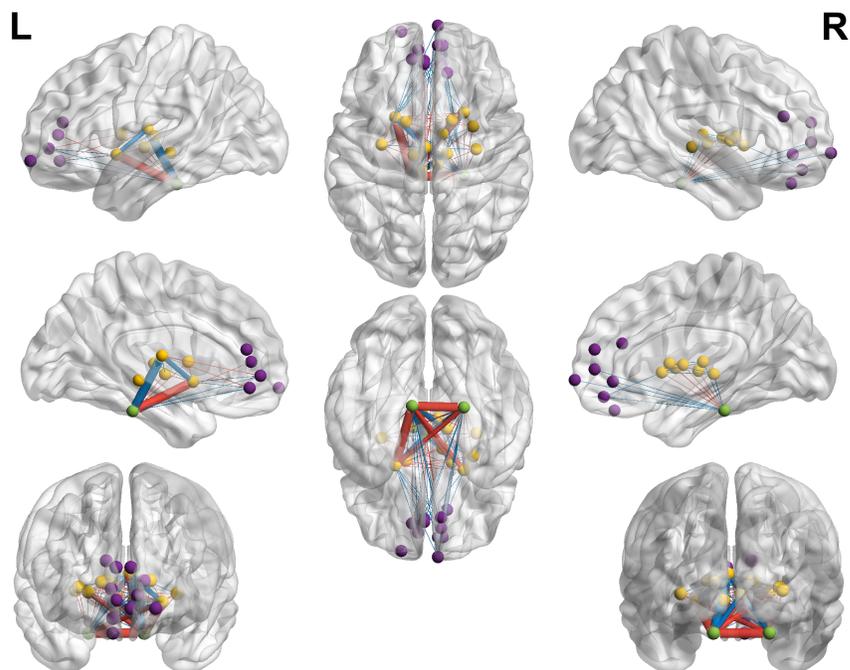
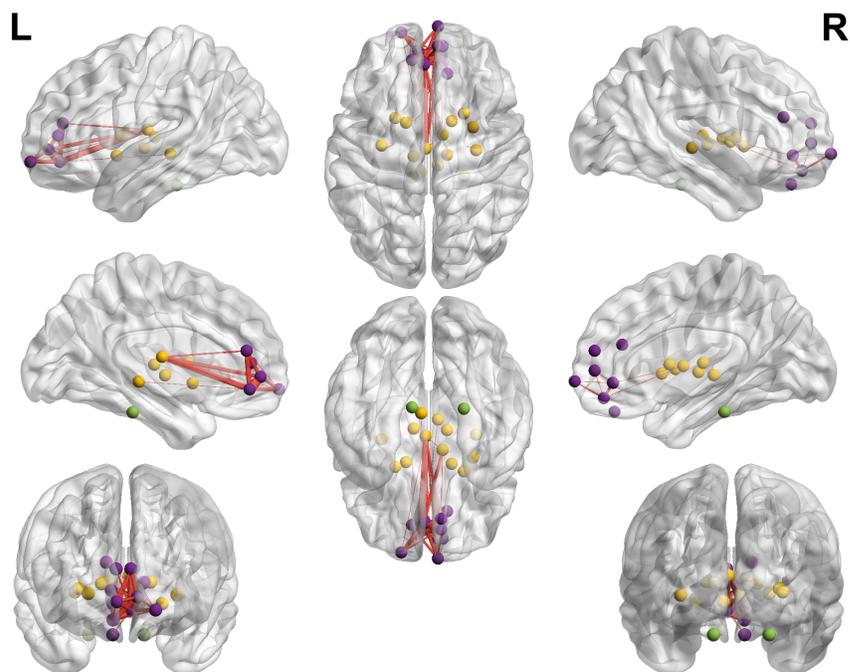


Figure 4.5: The BayDiL covariance estimates between each of the twenty seven brain regions in the first and fourth dictionary atoms. Edges are weighted by the magnitude of the covariance between the regions and the color indicates if it is positive (red) or negative (blue).



(A) First dictionary atom



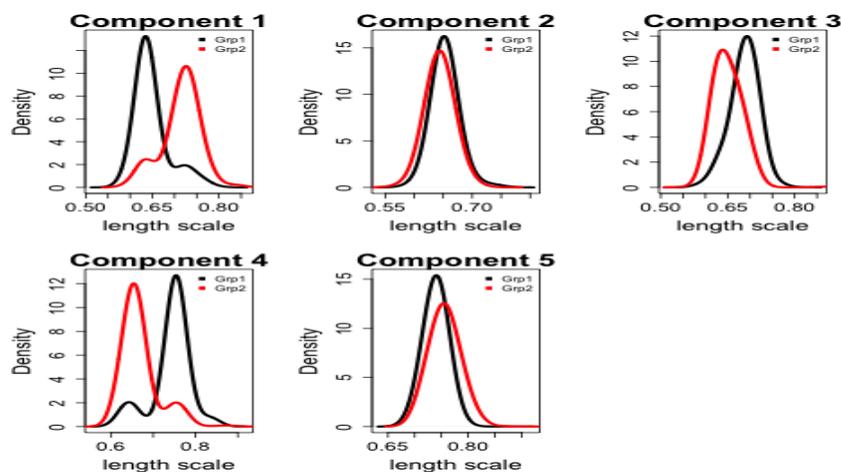
(B) Fourth dictionary atom

The BayDiL clusters the rs-fMRI data into two composed of 46 (group 1) and 10 (group 2) individuals based upon the subjects' utilization of the dictionary atoms. We study several characteristics of the groups and discover clinically relevant differences. Group 1 has lower median childhood trauma questionnaire scores compared to the group 2 (36 and 44.9, respectively). This suggests that subjects assigned to the second population experienced more traumatic childhoods, which was shown by Cisler (2017) to produce irregular functional connectivity between the mPFC and amygdala brain regions. Group based differences are also evident via the Beck Depression Inventory where group 1 on average exhibits minimal depressive symptoms while group 2 has mild to moderate symptoms. Further, 13% of group 1 has received depression treatment and 40% of group 2 ($\chi^2 = 4.10$, p-value=.049). Although our method originally intends to discover PTSD specific modifications in dynamic functional activity, twenty one percent of group 1 self report PTSD symptoms while thirty percent of group 2 reports the symptoms.

We find that the length scale parameter governing the covariance functions is similar in the groups for components two, three, and five (figure 4.6). Components one and four depict mild differences in the length scale parameter, where group one has a larger value (smoother temporal patterns) in the first component and smaller value in component one. In figure 4.5, we see that component one exhibits positive BOLD signals for all subcortical regions and the amygdala while the fourth component exhibits positive signals in the mPFC region. We note that the bimodal densities results from label switching issues inherent in Bayesian mixture modeling.

Finally, we assess three symptom clusters characterizing PTSD (intrusion, avoidance/numbing, and hyperarousal). Table 4.3 shows that group 1 has a larger scores in the intrusion domain while group 2 has higher scores in the avoidance and hyperarousal domains. Scher et al. (2008) find that the clinician administered PTSD scale (CAPS), which is the gold standard for PTSD assessment (Weiss, 2004), "explicitly

Figure 4.6: The posterior distribution of the length scale parameter for each component for the two subgroups in the PTSD data application. Differences in the group’s component parameter is evident in components one and four.



links numbing and hyperarousal symptoms to the experience of a traumatic event.” The authors go on to suggest that high scores in the two domains may indicate symptoms of PTSD, depression, or “simply a tendency to startle easily.” This echoes our earlier findings that subjects assigned to group 2 have stronger depressive symptoms, which is a condition frequently presenting in PTSD populations.

Table 4.3: Average score in the intrusion, avoidance/numbing, and hyperarousal domains. Reported values are the mean and standard deviation.

	Intrusion	Avoid/Numb.	Hyperarousal
Group 1	1.72(1.75)	1.89(1.99)	1.87(1.68)
Group 2	1.00(1.63)	2.30(2.40)	2.60(1.78)

As a check for the consistency and replicability of the results, we consider a larger parcellation scheme encompassing 74 brain regions, where the added regions are in the default mode network (DMN). We add these regions because the DMN has been highly implicated in a wide range of neurological disorders impacting brain function (Mohan et al., 2016). Further, the DMN has been shown to exhibit deficient connectivity in adults with childhood maltreatment-related PTSD (Daniels et al., 2011). We find that two clusters arise—48 subjects in group 1 and 8 subjects in group 2—that are

strongly consistent with the clusters identified in the twenty seven region analysis. Specifically, six subjects in group 2 also clustered together in group 2 for the 27 region analysis.

4.4 Conclusion

In this chapter, we propose a semi-parametric hierarchical Bayesian dictionary learning method that samples the model parameters via an efficient MCMC routine. Our proposed model exhibits several strengths with regards to the detection of dictionary atoms. Additionally, we have a simple framework for clustering individuals into latent groups. In the numerical studies, the BayDiL does an excellent job recovering the dictionary atoms and assigning observations to the correct group. In the data application, the BayDiL detects subgroups exhibiting suggestive, although statistically insignificant group differences. Individuals diagnosed with PTSD often present at least one other psychiatric disorder. While major depressive disorder is the most common co-morbid condition, substance abuse and anxiety disorders frequently exist in PTSD populations (Brady et al., 2000). Our unsupervised learning approach doesn't detect subgroups perfectly aligned with PTSD diagnoses, but does detect a subgroup exhibiting depressive symptoms and higher scores in domains aligned with PTSD as assessed by the CAPS.

We believe our method can be extended to address the more clinically relevant but statistically challenging issue of differential disease classification. The persistence of other psychiatric disorders in PTSD populations is regularly observed various conditions. For example, there is a growing class of methods leveraging functional and structural brain data to differentiate schizophrenia from depression and other mood disorders (Koutsouleris et al., 2015; Yu et al., 2013) as well as depression from bipolar disorder (Grotegerd et al., 2013; Serpa et al., 2014). As currently presented, the

BayDiL will be useful to clinicians for evaluating functional activation patterns as biomarkers to differentiate co-presenting disorders. Furthermore, the shared dictionary permits direct comparison across subjects used in the analysis. We expect the clustering results to help in the selection of treatment plans, for which success is directly connected to a correct initial diagnosis.

Additionally, we think that the selection of regions more closely aligned with the disorder of interest should lead to more clinically relevant findings. As previously discussed, our atlas of twenty seven brain regions spans the mPFC, sub cortical areas, and amygdala. Unfortunately, our atlas does not cover the hippocampus, which is regularly cited as an important brain area in PTSD populations. Crucially, the hippocampus is involved in memory of places and is likely intricately tied to the avoidance/numbing domain characterizing PTSD.

Despite the strengths of our method, there are various computational and statistical challenges to overcome. An obvious limitation to our work that is well documented in the matrix decomposition literature is that of identifiability. Specifically, the dictionary and mixing matrices are approximate up to a finite positive constant. In simulations, this led to numerical instabilities as the mixing matrices became increasingly non-Gaussian. In this work, we fix the global signal variance parameter for all components to constrain the magnitude of the mixing coefficients. In our numerical studies, the dictionary was nearly fully recovered (i.e. positive constant ≈ 1). Although this solved our numerical issues related to model identifiability, fixing the global signal may be too restrictive in practice. Another potential solution to the identifiability issue is to use a more restrictive prior on the columns of the dictionary such as the von-Mises Fisher (vMF) distribution that samples dictionary atoms from the space of vectors on the unit hypersphere. This is probabilistically equivalent to unit norm constraints imposed on the dictionary atoms in optimization routines.

We also found that the clustering performance was severely impacted by periodic

temporal mixing patterns. The BayDiL assumes that correlations overtime follow a Gaussian process with a square exponential covariance kernel function. Although commonly used, the square exponential kernel can miss local patterns that are likely present in neuroimaging data. A potential solution to this is to expand the complexity of the kernel function to incorporate global and local trends, such as the addition of a global square exponential kernel to the product of square exponential and periodic kernels. Williams and Rasmussen (2006) show that this composition produces a valid kernel. This construction will capture local periodic temporal trends likely present in fMRI data. We note that there should be negligible increases in computational time attributed to the expanded number of kernel parameters due to the use of the MALA sampling procedure. Another possible solution is to directly model the observed data as a function of the Gaussian process distributed mixing coefficients. One approach is to impose an orthogonality constraint on the columns of the dictionary such that sampling of \mathbf{A}_i is now based on the regression equation $\mathbf{D}'\mathbf{X}_i = \mathbf{A}_i + \mathbf{D}'\mathbf{B}(C_i \otimes \mathbf{1}'_T) + \mathbf{D}'\mathbf{E}_i$ where \mathbf{D} has been updated.

In future work, we would like to extend our model to task fMRI studies. Specifically, it has been observed that individuals with lower cognitive functioning exhibit irregular functional dynamics between task blocks (Cohen, 2018). While BayDiL has been formulated for rs-fMRI, we can naturally extend it to task-based studies by modifying the length scale parameter to reflect a population value shared by all subjects and random subject specific variation attributed to the subgroup. The BayDiL method requires the user to pre-specify the number of dictionary components prior to analysis. In simulations, we found that the specification of correct number of components is paramount to the classification performance. This was also evident in the data application, where the BayDiL returned uninteresting clustering when three, four, six or seven components was specified. We will incorporate a Dirichlet process prior to have a data-informed selection of the appropriate number of components.

Chapter 5

Summary and Future Directions

In this dissertation, we have investigated the utility of networks for analyzing functional connectivity in the human brain. Networks are an appealing approach for studying the brain because they significantly reduce the complexity of cortical communication. For instance, the brain contains millions of neurons which are densely packed in the cortex. It is computationally infeasible to study the interactions between these small units. However, network based approaches simplify this problem by aggregating the signals from spatially proximal neurons into one region. Thus, the success of the network approaches depend in part on the manner in which the brain is parcellated into non-overlapping regions of interest. Further, the parcellation scheme allows for various scales of analysis (i.e. coarse to dense). Throughout this dissertation, we've used moderately coarse parcellation schemes providing full coverage of the brain.

In Chapter 2, we proposed a difference degree test that detects brain regions incident to a statistically significant number of differentially weighted edges. Current network comparison methods detect global differences in network structure across two populations and typically fall into one of two categories: (1) mass univariate testing (2) differentially weighted edges form large, connected subcomponents spanning multiple brain systems. Our method is a modified mass univariate testing approach, where first threshold meaningful edges with a statistically meaning cutoff and then use the Hirshbaumer-Qi-Steuer (HQS) algorithm to generate meaningful null networks for the detection of functionally compromised brain regions. In application to a study of Major Depressive Disorder, the DDT detects regions supported by literature on the disorder. The DDT is an excellent approach for detecting local differences in network architecture, but can be significantly improved. First, we can extend the HQS algorithm to generate null networks that incorporate edge-specific variability. As currently implemented, the HQS procedure assumes that edges between regions have a common variance. However, this is likely unrealistic due to heterogeneity

of functional connectivity in the brain. Second, the DDT can be tremendously improved by removing of the thresholding step. Network thresholding is a critical step in nearly every human brain network application. Recently, Baggio et al. (2018) proposed threshold-free network based statistics which reduce the impact of thresholding by “averaging” over the tuning parameter. We can extend the DDT by assuming a Markov random field over the nodes, where each region’s height is the last threshold at which it is a statistically meaningful node. The DDT is easily modified to incorporate this step and will not exhibit computational inefficiencies.

In Chapter 3, we proposed an approach that incorporates structural connectivity into the estimation of functional connectivity. Logically, structures must be present to permit the observed coordination in functional activation between spatially remote neurons. In fact, a large body of literature suggests that structural connectivity, as measured by white matter fiber tracks, constrains functional connectivity in the brain. While various existing approaches enforce a deterministic relationship between brain function and structure, our siGGM is the first to flexibly incorporate anatomical connectivity into estimation of functional connectivity. Although we note superior performance relative to popular competitors, our approach only considers direct structural connections between brain regions. Skudlarski et al. (2008) show that the structure-function relationship is strongest when indirect structural connections up to order four are accounted for. Our method readily extends to include indirect structural connectivity. However, in simulation studies not presented, we found that naively accounting for indirect structural connections lead to dense precision matrix estimates. There are many options to ensure this does not occur. We could place a Laplace prior on coefficients of the the indirect structural connections or a group lasso prior if we want to impose structure.

Static connectivity dominates the brain network literature, primarily because it is easily calculated and intuitively straightforward. As we discuss in chapter 4, increas-

ingly researchers finding that communication between brain regions rapidly changes during tasks and at rest. Subsequent studies of dynamic functional connectivity indicate that it is a better predictor of various disorders than static connectivity. In that chapter, we propose the BayDiL method, which detects a set of dictionary atoms common to all subjects that linearly mix to reproduce the observed signals. Further, our method classify individuals based upon the utilization of the atoms. There are several innovations that can increase the effectiveness and reliability of results. First, we want to incorporate constraints on the dictionary atoms as implemented in the optimization routines, thereby allowing us to relax the constraint on the global signal parameters governing the covariance kernel functions. In practice, it is likely that that the global signal plays a significant role in how subgroups utilize the dictionary atoms. Furthermore, we want to allow the method to detect the number of components in the data rather than require a prior selection. As discussed in the chapter, classification performance is intricately tied to the selection of the correct number of dictionary atoms. Finally, we want to investigate an appropriate functional connectivity object as the observed signals. In unreported work, we investigated two measures of dynamic functional connectivity—sliding windows and dynamical conditional correlations. The BayDiL performed very poorly when sliding windows was input as the observed signal. This is likely because the sliding windows merge the components into one and average out the temporally varying changes that occur at smaller time scales. Unfortunately, sliding windows breaks down when extremely short window lengths are used due to an insufficient number of timepoints for calculating Pearson correlations. The dynamical conditional correlations provide the temporal resolution desired, but is prohibitively slow for practical use. We believe that advances in these methods should allow us to fully incorporate covariate effects into the classification of subjects based on usage of a common dictionary.

Appendix A

Appendix for Chapter 2

A.1 Proof for HQS procedure

In section 2.2.2.2 we suggest that sampling $f_{ij} \sim N(\mu, \sigma^2)$ appropriately allows for the condition that that $E(c_{ij}) = \bar{e}$ and $Var(c_{ij}) = \bar{v}$. We now provide details for the distribution of $c_{ij} = \sum_k^m f_{ik} \times f_{kj}$. Consider $f_{ik}, f_{kj} \stackrel{i.i.d.}{\sim} N(\mu, \sigma^2)$ for $k = 1, \dots, m$.

$$\sum_{k=1}^m f_{ik} f_{kj} = \frac{1}{4} \sum_{k=1}^m (f_{ik} + f_{kj})^2 - \frac{1}{4} \sum_{k=1}^m (f_{ik} - f_{kj})^2 \quad (\text{A.1})$$

For $X, Y \stackrel{i.i.d.}{\sim} N(\mu, \sigma^2)$,

$$\begin{aligned} \frac{(X+Y)}{\sqrt{2}\sigma} &\sim N\left(\frac{2\mu}{\sqrt{2}\sigma}, 1\right) \implies \frac{(X+Y)^2}{2\sigma^2} \sim \chi_1^2\left(\frac{4\mu^2}{2\sigma^2}\right) \\ \frac{(X-Y)}{\sqrt{2}\sigma} &\sim N(0, 1) \implies \frac{(X-Y)^2}{2\sigma^2} \sim \chi_1^2 \end{aligned}$$

We can introduce constants and rewrite (A.1) as

$$\begin{aligned}\sum_{k=1}^m f_{ik}f_{kj} &= \frac{1}{4} \sum_{k=1}^m (f_{ik} + f_{kj})^2 - \frac{1}{4} \sum_{k=1}^m (f_{ik} - f_{kj})^2 \\ \sum_{k=1}^m f_{ik}f_{kj} &= \frac{2\sigma^2}{4} \sum_{k=1}^m \frac{(f_{ik} + f_{kj})^2}{2\sigma^2} - \frac{2\sigma^2}{4} \sum_{k=1}^m \frac{(f_{ik} - f_{kj})^2}{2\sigma^2} \\ \sum_{k=1}^m f_{ik}f_{kj} &= \frac{2\sigma^2}{4} T - \frac{2\sigma^2}{4} Q\end{aligned}$$

where T is a non-central χ^2 with m df and non-centrality parameter $m \times (\frac{4\mu^2}{2\sigma^2})$ and Q is a central χ^2 with m df. Utilizing the first moment of non-central χ^2 and χ^2 distributions, we see that

$$\begin{aligned}E[\frac{2\sigma^2}{4}T - \frac{2\sigma^2}{4}Q] &= \frac{2\sigma^2}{4}E[T] - \frac{2\sigma^2}{4}E[Q] \\ &= \frac{2\sigma^2}{4}(m + \frac{4m\mu^2}{2\sigma^2}) - \frac{2\sigma^2}{4}m \\ &= m\mu^2 \\ &= \bar{e}\end{aligned}$$

and

$$\begin{aligned}Var[\frac{2\sigma^2}{4}T - \frac{2\sigma^2}{4}Q] &= \frac{4\sigma^4}{16}Var[T] + \frac{4\sigma^4}{16}Var[Q] - (\frac{4\sigma^4}{16})^2Cov(T, Q) \\ &= \frac{4\sigma^4}{16}(2 \times (m + 2(\frac{4m\mu^2}{2\sigma^2}))) + \frac{4\sigma^4}{16}(2m) \\ &= m\sigma^4 + 2\sigma^2m\mu^2 \\ &= m\sigma^4 + 2m\sigma^2\mu^2 + m\mu^4 - m\mu^4 \\ &= m(\sigma^2 + \mu^2)^2 - m\mu^4 \\ &= \bar{v}\end{aligned}$$

To see the $Cov(T, Q)=0$, we note that $(\mathbf{x}, \mathbf{y})^T \sim MVN(\tilde{\mu}, \tilde{\Sigma})$ where $\tilde{\mu} = \mu \times \mathbf{1}_{2m}$ for $\mathbf{1}_{2m}$ a vector of one's in \mathbb{R}^{2m} and $\tilde{\Sigma} = diag(\Sigma, \Sigma)$ is a block matrix with $\Sigma =$

$diag(\sigma^2, \dots, \sigma^2) \in \mathbb{R}^{2m \times 2m}$. Multiplying the multivariate random vector by an appropriate matrix, P , we have $(\mathbf{x}_1 + \mathbf{y}_1, \dots, \mathbf{x}_m + \mathbf{y}_m, \dots, \mathbf{x}_1 - \mathbf{y}_1, \dots, \mathbf{x}_m - \mathbf{y}_m)' \sim MVN((2\mu, \dots, 2\mu, 0, \dots, 0)', P\Sigma P')$. By the partitioning of the full covariance matrix, we see that $(x_1 + y_1, \dots, x_m + y_m)' \perp\!\!\!\perp (x_1 - y_1, \dots, x_m - y_m)'$. Consider $f(\mathbf{x}) = \frac{1}{2\sigma^2}(x_1^2 + \dots + x_m^2)$. Since $f(\cdot)$ is a continuous function, we have $f(x_1 + y_m, \dots, x_m + y_m) \perp\!\!\!\perp f(x_1 - y_m, \dots, x_m - y_m)$. By definition of $f(\cdot)$, we have $T \perp\!\!\!\perp Q$ which implies $Cov(T, Q) = 0$.

A.2 Tables

Table A.1: Within and between functional module DWE in the major depressive disorder study. Bolded values indicate statistically significant number of DWE between the respective functional modules

(A) Pearson, model-free												
	SM	CIO	AUD	DMN	MEM	VIS	FPN	SAL	SUB	VAN	DAN	UNC
SM	55											
CIO	11	1										
AUD	14	4	1									
DMN	36	22	21	42								
MEM	6	0	0	7	0							
VIS	29	14	28	32	1	4						
FPN	10	1	9	26	1	11	7					
SAL	11	2	12	15	1	4	9	2				
SUB	20	3	12	6	5	63	2	1	2			
VAN	25	6	2	15	0	8	2	2	2	0		
DAN	20	0	4	9	0	9	3	2	2	2	0	
UNC	22	10	3	43	5	13	11	9	9	4	2	5

(B) Pearson, model-based												
	SM	CIO	AUD	DMN	MEM	VIS	FPN	SAL	SUB	VAN	DAN	UNC
SM	46											
CIO	9	1										
AUD	12	4	1									
DMN	27	25	22	45								
MEM	3	0	0	5	0							
VIS	29	14	27	26	4	3						
FPN	9	0	7	30	1	12	7					
SAL	10	2	8	16	1	6	12	4				
SUB	16	3	9	10	6	54	1	1	2			
VAN	23	5	2	15	0	6	2	2	2	0		
DAN	17	0	2	9	0	14	3	3	1	2	0	
UNC	20	9	1	48	4	14	10	12	9	4	3	5

(C) Partial, model-free												
	SM	CIO	AUD	DMN	MEM	VIS	FPN	SAL	SUB	VAN	DAN	UNC
SM	6											
CIO	2	0										
AUD	2	3	3									
DMN	25	6	12	21								
MEM	2	1	0	6	0							
VIS	12	5	7	16	1	7						
FPN	13	4	5	19	1	7	3					
SAL	9	6	1	16	0	5	6	1				
SUB	4	0	0	11	1	8	5	5	1			
VAN	2	3	0	11	0	4	0	3	1	0		
DAN	6	3	1	9	0	4	1	5	3	1	1	
UNC	15	6	2	32	2	11	10	8	9	2	4	6

(D) Partial, model-based												
	SM	CIO	AUD	DMN	MEM	VIS	FPN	SAL	SUB	VAN	DAN	UNC
SM	9											
CIO	7	0										
AUD	8	4	3									
DMN	47	15	17	36								
MEM	3	1	1	9	0							
VIS	16	8	11	30	5	11						
FPN	20	8	7	40	3	19	7					
SAL	14	10	2	22	1	9	11	4				
SUB	7	4	2	18	2	12	8	8	1			
VAN	8	3	1	19	1	8	2	6	2	0		
DAN	6	4	4	17	0	6	4	7	4	3	3	
UNC	30	13	7	49	5	22	16	10	13	7	7	10

Appendix B

Appendix for Chapter 3

B.1 Parameter Updates

We optimize (3.6) by iteratively updating model parameters as follows.

Update Ω : Given the data, $\mathbf{Y} = \{\mathbf{y}_1, \dots, \mathbf{y}_T\}$, and current estimates for all other model parameters, we solve

$$\hat{\Omega}^{(m+1)} = \arg \min_{\Omega} -\log \det(\Omega) + \text{tr}(S\Omega) + \frac{\nu}{2} \sum_{j < k} e^{\alpha_{jk}^{(m+1)}} |\omega_{jk}| + \frac{1}{2} \nu \sum_{j=k} |\omega_{kk}| \quad (\text{B.1.0.1})$$

for $\hat{\Omega}$. This resembles the penalized likelihood framework of the traditional Gaussian graphical model. Define $\delta_{ij}^{(k+1)} = \frac{1}{2} \exp(\alpha_{jk}^{(m+1)})$ for $j \neq k$ and $\delta_{jk}^{(m+1)} = \frac{1}{2}$ for $j = k$. We can re-express (B.1.0.1) as

$$\hat{\Omega}^{(m+1)} = \arg \min_{\Omega} -\log \det(\Omega) + \text{tr}(S\Omega) + \nu \sum_{j < k} \delta_{jk}^{(m+1)} |\omega_{jk}|,$$

where we update Ω using a quadratic approximation solver, *QUIC* (Hsieh et al., 2011), available in R.

Update μ_{jk} : Given \mathbf{Y} , $\eta^{(m+1)}$, and $\alpha_{jk}^{(m)}$, we update μ_{jk} via the closed form equation

$$\mu_{jk}^{(m+1)} = \frac{\sigma_\mu^2(\alpha_{jk}^{(m)} + \eta^{(m+1)}p_{jk}) + \sigma_\lambda^2\mu_0}{\sigma_\mu^2 + \sigma_\lambda^2}$$

Update η : Given \mathbf{Y} and $\boldsymbol{\alpha}^{(m)}$, we can update η via closed form equation

$$\hat{\eta}^{(m+1)} = \frac{-\beta^{(m)} + \sqrt{(\beta^{(m)})^2 - 4\gamma\rho}}{2\gamma},$$

where $\beta^{(m)} = b_\eta + \frac{\sum_{j<k} \alpha_{jk}^{(m)} p_{jk}}{\sigma_\lambda^2} - \frac{1}{\sigma_\lambda^2} \sum_{j<k} \mu_{jk}^{(m)} p_{jk}$, $\gamma = \frac{\sum_{j<k} p_{jk}^2}{\sigma_\lambda^2}$, and $\rho = -\frac{1}{\sigma_\lambda^2}(a_\eta - 1)$

Update $\boldsymbol{\alpha}$: Given \mathbf{Y} , $\boldsymbol{\Omega}^{(m)}$, $\boldsymbol{\mu}^{(m+1)}$, and $\eta^{(m+1)}$, we can estimate $\alpha_{jk}^{(m+1)}$ for $1 \leq j < k \leq p$ by solving

$$\hat{\boldsymbol{\alpha}}^{(m+1)} = \arg \min_{\boldsymbol{\alpha}} \nu \sum_{j<k} e^{\alpha_{jk}} |\omega_{jk}^{(m)}| + \sum_{j<k} \frac{(\alpha_{jk} - (\mu_{jk}^{(m+1)} - \eta^{(m+1)}p_{jk}))^2}{2\sigma_\lambda}$$

A closed form solution doesn't exist, so we implement a Newton Raphson solver to find the optimal choice of $\boldsymbol{\alpha}$. Re-expressing this problem, we have

$$\arg \min_{\boldsymbol{\alpha}} \exp(\boldsymbol{\alpha})' |\boldsymbol{\omega}^{(m+1)}| - \frac{1}{2\sigma_\lambda^2} (\boldsymbol{\alpha} - (\boldsymbol{\mu}^{(m+1)} - \eta^{(m+1)}\tilde{P}))' (\boldsymbol{\alpha} - (\boldsymbol{\mu}^{(m+1)} - \eta^{(m+1)}\tilde{P}))$$

where $\boldsymbol{\alpha} = \{\alpha_{12}, \alpha_{13}, \dots, \alpha_{(p-1)p}\}$, \tilde{P} denotes the upper diagonal elements of the structural connectivity matrix P , $e^{\boldsymbol{\alpha}}$ is the element wise exponential for each component of $\boldsymbol{\alpha}$, and $\mathbf{1}$ as a vector of 1's of length $\frac{p(p-1)}{2}$. Since $\boldsymbol{\Omega}$ is symmetric and we do not shrink diagonal elements, we simplify our estimation of $\boldsymbol{\alpha}$ by only focusing upon the upper diagonal elements.

The Newton Raphson updating equation based on step size Δ is $\boldsymbol{\alpha}^{m+1} = \boldsymbol{\alpha}^m - \Delta g(\boldsymbol{\alpha}^m) H(\boldsymbol{\alpha}^m)^{-1}$, where $g(\boldsymbol{\alpha}) = \nu \sigma_\lambda^2 D_{|\boldsymbol{\omega}^{(m)}|} e^{\boldsymbol{\alpha}} + [\boldsymbol{\alpha} - (\boldsymbol{\mu}^{(m+1)} - \eta^{(m+1)}\tilde{P})]$ and $H(\boldsymbol{\alpha}) = \nu \sigma_\lambda^2 D_{|\boldsymbol{\omega}^{(m)}|} D_{|e^{\boldsymbol{\alpha}}|} + I$, $D_{|\boldsymbol{\omega}^{(m)}|}$ is a $\frac{p(p-1)}{2} \times \frac{p(p-1)}{2}$ diagonal matrix with elements as the upper triangular elements of $\boldsymbol{\Omega}$, and similarly for $D_{|e^{\boldsymbol{\alpha}}|}$, and I is an identity matrix. Since H is a diagonal matrix, it is easily inverted and serves as an appropriate Hessian

matrix. We search for the step size (Δ) using a back tracking line search for each update of α as in Chang et. al (2017).

B.2 Hyperparameter Choice and Initial Values

The proposed siGGM approach iteratively solves for the MAP estimator and works best when reasonable starting values are provided. We first find an initial estimate for the graph structure and the sparse inverse precision matrix (Ω_0), using the graphical lasso. We initialize all edge specific penalty parameters as λ_0 , which is the global tuning parameter corresponding to Ω_0 . We set $\sigma_\mu^2 = 5$, corresponding to an uninformative prior which reflects our lack of knowledge regarding the baseline effects and choose $\mu_0 = 0$ as a default setting. We randomly generate the edge specific baseline effects μ_{jk} from the prior distribution $N(\mu_0, \sigma_\mu^2)$ and use these as initial values. The initial value of η is chosen by averaging $-\frac{\sum_{l < k} (\exp(\lambda_0) - \mu_{jk}) / p_{jk}}{p(p-1)/2}$, which is the average of all possible η values under the relationship $\exp(\lambda_0) = \mu_{jk} - \eta p_{jk}, j < k$ corresponding to $\sigma_\lambda = 0$. We choose $\sigma_\lambda^2 = \frac{1}{p(p-1)/2} \sum_{j < k} \sum_{j,k=1}^p (\exp(\lambda_0) - \mu_{jk} - \eta p_{jk})^2$.

Finally, we found that choosing a_η and b_η to attain $E[\eta] \approx 6$ and $\text{Var}[\eta] \approx 1$ incorporates structural information in a flexible manner. However, larger first moments for the prior on η may lead to increased false positives as our method places more weight on smaller structural connections, and similarly, smaller first moment may decrease the overall impact of structural information. For example, when $a_\eta > 1$ and $b_\eta \rightarrow \infty$, we have $E[\eta] \rightarrow 0$, which makes the siGGM indistinguishable from SC naive methods. In simulations, we found values selected for the set of hyperparameters— $(a_\eta, b_\eta, \sigma_\mu, \sigma_\lambda)$ —led to fast estimation of model parameters and accurate results across a wide array of settings.

B.3 Measure for computing between module differences

We define the goodness of fit measure

$$X_{g_1, g_2}^2 = \frac{(Q_{(g_1, g_2)} - E_{(g_1, g_2)})^2}{E_{(g_1, g_2)}}, \quad (\text{B.3.0.1})$$

where $g_1, g_2 \in \{1, \dots, G\}$ are the indices corresponding to one of the G functional modules, Q_{g_1, g_2} represents the observed number of differentially weighted edges in the (g_1, g_2) block, E_{g_1, g_2} represents the expected number of differentially weighted edges in the (g_1, g_2) block when edges distribute randomly across the module blocks. X_{g_1, g_2}^2 measures the goodness of fit for each within-module block ($g_1 = g_2$) or between-module block ($g_1 \neq g_2$). In equation (B.3.0.1), the expected value can be derived in a straightforward manner as $E_{g_1, g_2} = 0.5p^*\{|g_1|(|g_2| - 1)\}$ for within module blocks ($g_1 = g_2$) and $E_{g_1, g_2} = p^*|g_1||g_2|$ for between-module blocks ($g_1 \neq g_2$), where $|g|$ represents the total number of nodes within the g th module, and p^* represents the proportion of differentially weighted edges among all the edges across the network. Using 5000 permutations of group labels at each edge, the differentially weighted edges are identified as those connections with significant FDR-adjusted p-values.

B.4 Calculation of ICC

The intraclass correlation coefficient is a widely used reliability metric for assessing test-retest reliability of brain network topology in neuroimaging applications. Using ICC(3,1), (two-way mixed single measures testing for consistency) we investigate the reliability of graph metrics across two scanning session (Guo et al., 2012; Telesford et al., 2010). The quantity is calculated as

$$\text{ICC}(3, 1) = \frac{\text{BMS} - \text{EMS}}{\text{BMS} + (k - 1)\text{EMS}}, \quad (\text{B.4.0.1})$$

where k is the number of scanning sessions per participant, BMS is the between mean square and EMS is the mean residual sum of squares. BMS captures the variability between subjects while EMS measures unexplained within-subject variation in functional connectivity across scanning sessions (see Shrout and Fleiss (1979)). This metric is commonly used to measure test-retest network stability in brain networks (Braun et al., 2012) with agreement scale $0 < ICC \leq .2$ (slight), $.2 < ICC \leq .4$ (fair), $.4 < ICC \leq .6$ (moderate), $.6 < ICC \leq .8$ (strong), and $.8 < ICC \leq 1$ (near perfect) as suggested by Telesford et al. (2010).

B.5 Tables

Appendix B presents the results for siGGM and SC naive approaches on simulated data.

Table B.1: Performance of siGGM and SC naive approaches on simulated network data with $p=100$ and 200 nodes. Eglob is the bias in global efficiency.

	p=100				p=200			
	Eglob	MCC	AUC	L1	Eglob	MCC	AUC	L1
Small World								
Glasso	0.177	0.327	0.827	0.575	0.128	0.333	0.757	0.668
Space	-0.206	0.585	0.839	0.407	-0.374	0.597	0.763	0.430
siGGM($\eta = 0$)	0.061	0.538	0.847	0.509	-0.019	0.506	0.843	0.587
siGGM	0.078	0.590	0.884	0.478	0.121	0.526	0.906	0.532
Scale Free								
Glasso	0.117	0.365	0.798	0.560	0.038	0.324	0.657	0.605
Space	-0.219	0.495	0.772	0.491	-0.403	0.358	0.664	0.555
siGGM($\eta = 0$)	0.005	0.509	0.808	0.528	-0.100	0.411	0.769	0.573
siGGM	0.054	0.562	0.853	0.428	-0.075	0.469	0.868	0.442
Erdos-Renyi								
Glasso	0.245	0.247	0.789	0.859	0.065	0.182	0.659	0.837
Space	-0.125	0.580	0.824	0.465	-0.415	0.253	0.638	0.577
siGGM($\eta = 0$)	0.020	0.363	0.792	0.679	-0.204	0.208	0.661	0.700
siGGM	0.124	0.442	0.861	0.624	0.049	0.514	0.862	0.689

B.6 Network Metrics

Network metrics quantify properties of graphs, such as the local connectedness of regions and global connectivity across the network. These features distill complex organizing principles into scalar values that have led to critical insights into the brain's functionality in health and disease (Lord et al., 2012). In the following section, we provide interpretations and mathematical formulas for the seven network metrics investigated in this work. All functions assume the brain network is a $p \times p$ binary adjacency matrix, A , where a_{ij} is equal to one if there is an edge between regions i and j or zero otherwise. We first describe two metrics that are building blocks for our network features of interest. The first is the number of triangles around node i , $t_i = \frac{1}{2} \sum_{j,m=1}^p a_{ij}a_{im}a_{jm}$, which measures connectivity among regions adjacent to region i . The second is the shortest path length between regions i and j which is defined as $d_{ij} = \sum_{a_{uv} \in g_{i \leftrightarrow j}} a_{uv}$, where $g_{i \leftrightarrow j}$ is the set of edges comprising the shortest geodesic distance between the regions. This metric is a basis for quantifying global integration across the brain (Rubinov and Sporns, 2010a). For additional details, please see <https://sites.google.com/site/bctnet/measures>.

1. Degree

The degree of region i is the number of edges incident to it. Mathematically, it is defined as $k_i = \sum_{j \neq i}^p a_{ij}$. Regions i and j are considered neighbors (or adjacent) if $a_{ij} = 1$.

2. Characteristic Path Length

The characteristic path length is the average shortest path length in the network and is defined as $L = \frac{1}{p} \sum_i L_i = \frac{1}{p} \sum_i \frac{\sum_{j \neq i} d_{ij}}{p-1}$, where d_{ij} is the shortest path length and L_i is the characteristic path length of region i .

3. Clustering coefficient

The clustering coefficient measures the number of closed triangles contain-

ing a ROI. It is equivalent to the number of neighbors of region i that are also neighbors of each other (Watts and Strogatz, 1998). It is calculated as $C = \frac{1}{p} \sum_i C_i = \frac{1}{p} \sum_i \frac{2t_i}{k_i(k_i-1)}$, where C_i is the clustering coefficient of region i .

4. Local Efficiency

Braun et al. (2012) define local efficiency as “the efficiency of the local subgraph of a node that only contains direct neighbors of the region” and is a measure of local connectedness. It is defined as $E_{loc} = \frac{1}{p} \sum_i E_{loc,i} = \frac{1}{p} \sum_i \frac{\sum_{j,m=1,j \neq i} a_{ij} a_{im} [d_{jm}(N_i)]^{-1}}{k_i(k_i-1)}$, where $E_{loc,i}$ is the local efficiency of region i and $d_{jm}(N_i)$ is the length of the shortest path between regions j and m that contains only neighbors of i .

5. Global Efficiency

The global efficiency is the average inverse shortest path length in the network and quantifies integration across spatially distant regions in the brain. It is calculated as $E = \frac{1}{p} \sum_i E_i = \frac{1}{p} \sum_i \frac{\sum_{j=1,j \neq i}^p d_{ij}^{-1}}{p-1}$, where E_i is the global efficiency of region i .

6. Modularity

This metric measures how well the network decomposes into non-overlapping clusters of connected regions and is calculated as $Q = \sum_{u \in M} \left[e_{uu} - \left(\sum_{v \in M} e_{uv} \right)^2 \right]$, where M constitutes the sets of non-overlapping modules, e_{uv} is the proportion of all links that connect regions in module u with regions in module v .

7. Hierarchy

Network hierarchy is an ordering of regions such that high degree regions tend to have lower local clustering coefficient than low degree regions. Mathematically, this ratio is quantified as $C \sim k^{-\beta}$, where C is the clustering coefficient, k is the region’s degree, and β is the hierarchy coefficient that is large for strongly hierarchical structured networks. β is estimated as the slope coefficient from a linear regression of $\log(C)$ on $\log(k)$ (Braun et al., 2012).

Appendix C

Appendix for Chapter 4

C.1 Hyper-parameter selection and initial values

Our model has a few hyper-parameters that must be a priori selected. The set of parameters are $(\alpha_{\sigma_\epsilon}, \alpha_{\sigma_\beta}, \alpha_d, b_{\sigma_\epsilon}, b_{\sigma_\beta}, b_d)$, which govern the variance parameters in our model. Following Yang et al. (2016), we set $\alpha_{\sigma_\epsilon} = \alpha_{\sigma_\beta} = \alpha_d = .5$ and $b_{\sigma_\epsilon} = b_{\sigma_\beta} = b_d = 1e-6$, where these choices have been shown to induce sparsity promoting marginal distributions on the respective higher level parameters.

We initialize the BayDiL method with random starting values for all parameters. Since the mixing coefficient information is subject-specific and dependent upon group membership, we randomly assign all subjects to a group and sample the \mathbf{A}^i from a zero-mean Gaussian distribution with common covariance $\mathbf{R}_\ell \forall i$. For a non-random initialization, we utilize a two step procedure. First, using all subjects data, we use k-means clustering to determine the K centroids which serve as the columns of the shared dictionary. Next, we use the alternating direction method of multipliers (ADMM) basis pursuit (Boyd et al., 2011) in the ADMM R package to solve

$$\arg \min_{\mathbf{a}_t^i} \|\mathbf{a}_t^i\|_1 \text{ s.t. } \mathbf{D}\mathbf{a}_t^i = \mathbf{x}_t^i \text{ for } t = 1, \dots, T; i = 1, \dots, n$$

where \mathbf{D} is fixed from step one, \mathbf{x}_t^i is the observed vector of signals at the t th time step, and \mathbf{a}_t^i linearly mixes the basis dictionary to reproduce the observed signal. For all implementations of the BayDiL, we acquire 25,000 MCMC samples with 10,000 burnin.

C.2 Figures

Figure C.2.1 : Spatial maps of the true basis signal juxtaposed with the estimates from the BayDiL, DL $\lambda \in \{.2, .3\}$, and k-means centroids. The mixing coefficients are generated from a zero-mean gaussian distribution with square exponential covariance function. The BayDiL does an adequate job recovering the true atoms whereas the dictionary learners recover them with the addition of noise. K-means clustering produces estimates that merge the three true signals.

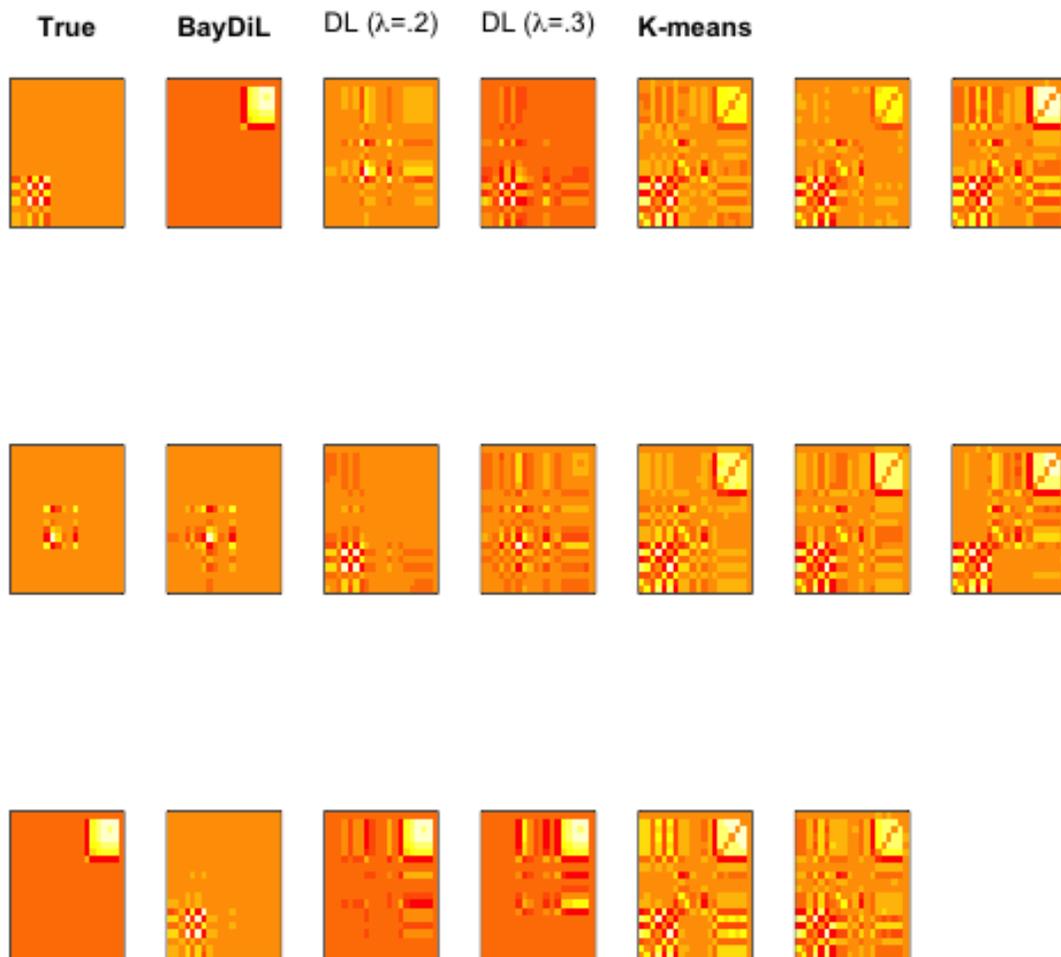


Figure C.2.2 : Spatial maps of the true basis signal juxtaposed with the estimates from the BayDiL, DL $\lambda \in \{.2, .3\}$, and k-means centroids. The mixing coefficients are generated from a zero-mean gaussian distribution with sinusoidal function. The BayDiL does an adequate job recovering the true atoms whereas the dictionary learners recover them with the addition of noise. K-means clustering produces estimates that merge the three true signals.

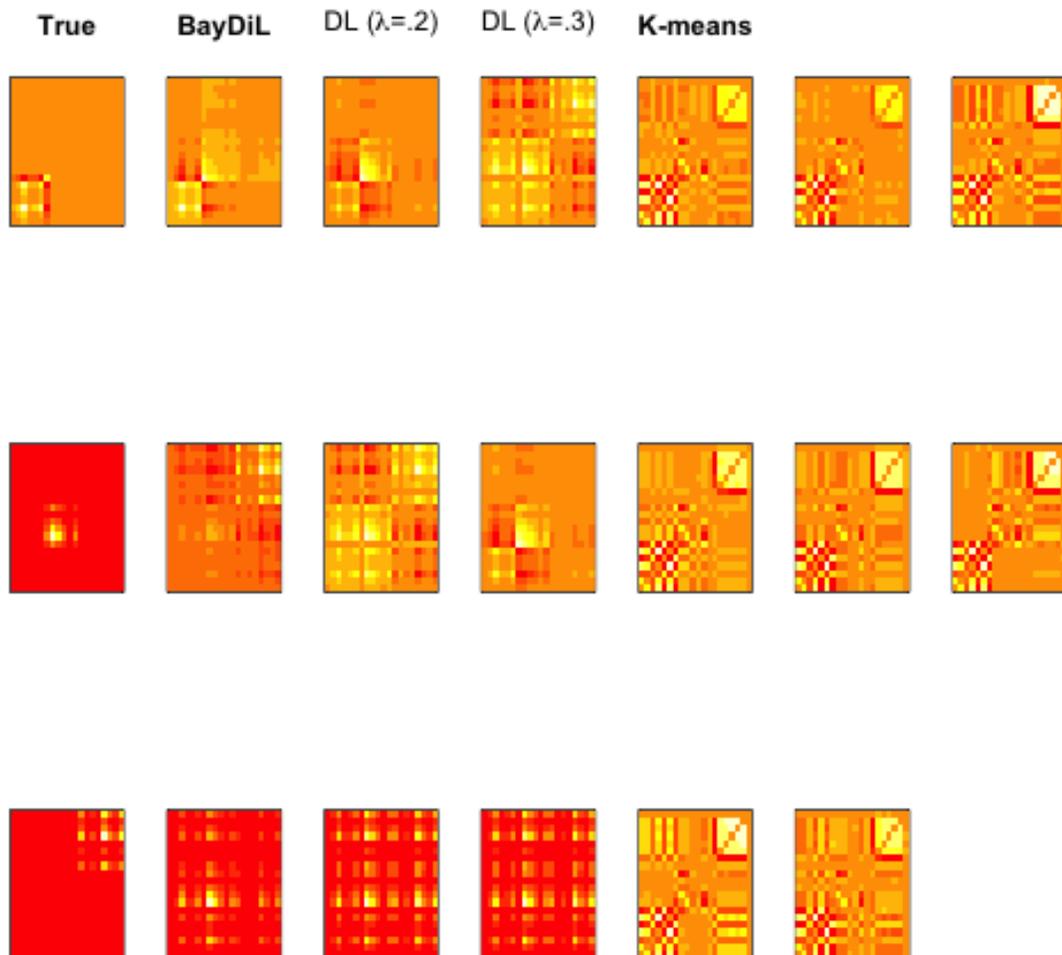
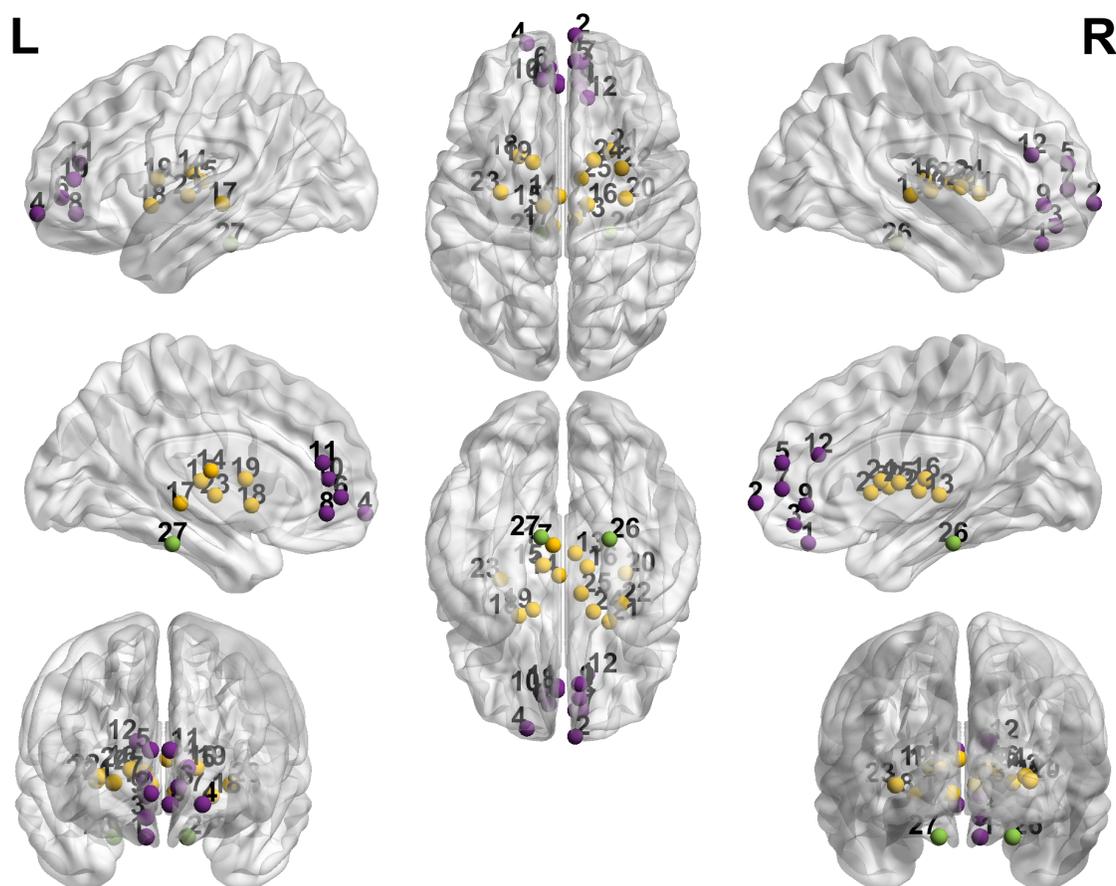


Figure C.2.3 : The twenty seven brain regions analyzed in the PTSD data application organized by the mPFC (*purple*), subcortical (*yellow*), and amygdala (*green*) functional areas. Regions are superimposed on a brain template using the BrainNetViewer matlab toolbox (Xia et al., 2013).



Bibliography

- Achard, S., Delon-Martin, C., Vértes, P. E., Renard, F., Schenck, M., Schneider, F., Heinrich, C., Kremer, S. and Bullmore, E. T. (2012), ‘Hubs of brain functional networks are radically reorganized in comatose patients’, Proceedings of the National Academy of Sciences **109**(50), 20608–20613.
- Adelstein, J. S., Shehzad, Z., Mennes, M., DeYoung, C. G., Zuo, X.-N., Kelly, C., Margulies, D. S., Bloomfield, A., Gray, J. R., Castellanos, F. X. et al. (2011), ‘Personality is reflected in the brain’s intrinsic functional architecture’, PloS one **6**(11), e27633.
- Agrawal, N. and Rickards, H. (2011), ‘Detection and treatment of depression in neurological disorders’, Journal of Neurology, Neurosurgery & Psychiatry **82**(8), 828–829.
- Aharon, M., Elad, M., Bruckstein, A. et al. (2006), ‘K-svd: An algorithm for designing overcomplete dictionaries for sparse representation’, IEEE Transactions on signal processing **54**(11), 4311.
- Allen, E. A., Damaraju, E., Plis, S. M., Erhardt, E. B., Eichele, T. and Calhoun, V. D. (2014a), ‘Tracking whole-brain connectivity dynamics in the resting state’, Cerebral cortex **24**(3), 663–676.
- Allen, E. A., Damaraju, E., Plis, S. M., Erhardt, E. B., Eichele, T. and Calhoun, V. D. (2014b), ‘Tracking whole-brain connectivity dynamics in the resting state’, Cerebral cortex **24**(3), 663–676.

- Allen, E. A., Erhardt, E. B., Wei, Y., Eichele, T. and Calhoun, V. D. (2011), ‘A simulation toolbox for fmri data: Simtb’, online user manual .
- Allen, E. A., Erhardt, E. B., Wei, Y., Eichele, T. and Calhoun, V. D. (2012), ‘Capturing inter-subject variability with group independent component analysis of fmri data: a simulation study’, Neuroimage **59**(4), 4141–4159.
- Andersen, M., Winther, O., Hansen, L. K., Poldrack, R. and Koyejo, O. (2018), Bayesian structure learning for dynamic brain connectivity, in ‘International Conference on Artificial Intelligence and Statistics’, pp. 1436–1446.
- Athanasίου, A., Terzopoulos, N., Pandria, N., Xygonakis, I., Foroglou, N., Polyzoidis, K. and Bamidis, P. D. (2018), ‘Functional brain connectivity during multiple motor imagery tasks in spinal cord injury’, Neural plasticity **2018**.
- Baggio, H. C., Abos, A., Segura, B., Campabadal, A., Garcia-Diaz, A., Uribe, C., Compta, Y., Marti, M. J., Valldeoriola, F. and Junque, C. (2018), ‘Statistical inference in brain graphs using threshold-free network-based statistics’, Human brain mapping **39**(6), 2289–2302.
- Banerjee, O., Ghaoui, L. E. and d’Aspremont, A. (2008), ‘Model selection through sparse maximum likelihood estimation for multivariate gaussian or binary data’, Journal of Machine learning research **9**(Mar), 485–516.
- Bansal, S., Khandelwal, S. and Meyers, L. A. (2009), ‘Exploring biological network structure with clustered random networks’, BMC bioinformatics **10**(1), 405.
- Barabási, A.-L. and Albert, R. (1999), ‘Emergence of scaling in random networks’, science **286**(5439), 509–512.
- Bassett, D. S. and Bullmore, E. (2006), ‘Small-world brain networks’, The neuroscientist **12**(6), 512–523.

- Bassett, D. S., Bullmore, E. T., Meyer-Lindenberg, A., Apud, J. A., Weinberger, D. R. and Coppola, R. (2009), ‘Cognitive fitness of cost-efficient brain functional networks’, Proceedings of the National Academy of Sciences **106**(28), 11747–11752.
- Bassett, D. S., Wymbs, N. F., Porter, M. A., Mucha, P. J., Carlson, J. M. and Grafton, S. T. (2011), ‘Dynamic reconfiguration of human brain networks during learning’, Proceedings of the National Academy of Sciences **108**(18), 7641–7646.
- Bernal-Casas, D., Lee, H. J., Weitz, A. J. and Lee, J. H. (2017), ‘Studying brain circuit function with dynamic causal modeling for optogenetic fmri’, Neuron **93**(3), 522–532.
- Besag, J., Green, P., Higdon, D., Mengersen, K. et al. (1995), ‘Bayesian computation and stochastic systems’, Statistical science **10**(1), 3–41.
- Betzler, R. F., Avena-Koenigsberger, A., Goñi, J., He, Y., De Reus, M. A., Griffa, A., Vértés, P. E., Mišić, B., Thiran, J.-P., Hagmann, P. et al. (2016), ‘Generative models of the human connectome’, Neuroimage **124**, 1054–1064.
- Biswal, B. B., Kylen, J. V. and Hyde, J. S. (1997), ‘Simultaneous assessment of flow and bold signals in resting-state functional connectivity maps’, NMR in Biomedicine **10**(4-5), 165–170.
- Biswal, B. B., Mennes, M., Zuo, X.-N., Gohel, S., Kelly, C., Smith, S. M., Beckmann, C. F., Adelstein, J. S., Buckner, R. L., Colcombe, S. et al. (2010), ‘Toward discovery science of human brain function’, Proceedings of the National Academy of Sciences **107**(10), 4734–4739.
- Biswal, B., Zerrin Yetkin, F., Haughton, V. M. and Hyde, J. S. (1995), ‘Functional connectivity in the motor cortex of resting human brain using echo-planar mri’, Magnetic resonance in medicine **34**(4), 537–541.

- Bowman, F. D., Zhang, L., Derado, G. and Chen, S. (2012), ‘Determining functional connectivity using fmri data with diffusion-based anatomical weighting’, NeuroImage **62**(3), 1769–79.
- Boyd, S., Parikh, N., Chu, E., Peleato, B., Eckstein, J. et al. (2011), ‘Distributed optimization and statistical learning via the alternating direction method of multipliers’, Foundations and Trends® in Machine learning **3**(1), 1–122.
- Brady, K. T., Killeen, T. K., Brewerton, T. and Lucerini, S. (2000), ‘Comorbidity of psychiatric disorders and posttraumatic stress disorder.’, The Journal of clinical psychiatry .
- Braun, U., Plichta, M. M., Esslinger, C., Sauer, C., Haddad, L., Grimm, O., Mier, D., Mohnke, S., Heinz, A., Erk, S. et al. (2012), ‘Test–retest reliability of resting-state connectivity network characteristics using fmri and graph theoretical measures’, Neuroimage **59**(2), 1404–1412.
- Bressler, S. L. and Tognoli, E. (2006), ‘Operational principles of neurocognitive networks’, International Journal of Psychophysiology **60**(2), 139–148.
- Brown, E. S., Murray, M., Carmody, T. J., Kennard, B. D., Hughes, C. W., Khan, D. A. and Rush, A. J. (2008), ‘The quick inventory of depressive symptomatology-self-report: a psychometric evaluation in patients with asthma and major depressive disorder’, Annals of Allergy, Asthma & Immunology **100**(5), 433–438.
- Brzezicka, A. (2013), ‘Integrative deficits in depression and in negative mood states as a result of fronto-parietal network dysfunctions’, Acta Neurobiol Exp **73**(3), 313–325.
- Bullmore, E. T. and Bassett, D. S. (2011), ‘Brain graphs: graphical models of the human brain connectome’, Annual review of clinical psychology **7**, 113–140.

- Bullmore, E. T., Suckling, J., Overmeyer, S., Rabe-Hesketh, S., Taylor, E. and Brammer, M. J. (1999), ‘Global, voxel, and cluster tests, by theory and permutation, for a difference between two groups of structural mr images of the brain’, IEEE transactions on medical imaging **18**(1), 32–42.
- Calhoun, V. D., Miller, R., Pearlson, G. and Adalı, T. (2014), ‘The chronnectome: time-varying connectivity networks as the next frontier in fmri data discovery’, Neuron **84**(2), 262–274.
- Catani, M. and Ffytche, D. H. (2005), ‘The rises and falls of disconnection syndromes’, Brain **128**(10), 2224–2239.
- Chang, C., Kundu, S. and Long, Q. (2016), ‘Scalable bayesian variable selection for structured high-dimensional data’, arXiv preprint arXiv:1604.07264 .
- Chen, S., Kang, J., Xing, Y. and Wang, G. (2015), ‘A parsimonious statistical method to detect groupwise differentially expressed functional connectivity networks’, Human brain mapping **36**(12), 5196–5206.
- Chen, Y. and Welling, M. (2012), ‘Bayesian structure learning for markov random fields with a spike and slab prior’, arXiv preprint arXiv:1206.1088 .
- Cheng, W., Rolls, E. T., Qiu, J., Xie, X., Lyu, W., Li, Y., Huang, C.-C., Yang, A. C., Tsai, S.-J., Lyu, F. et al. (2018), ‘Functional connectivity of the human amygdala in health and in depression’, Social cognitive and affective neuroscience **13**(6), 557–568.
- Choe, A. S., Nebel, M. B., Barber, A. D., Cohen, J. R., Xu, Y., Pekar, J. J., Caffo, B. and Lindquist, M. A. (2017), ‘Comparing test-retest reliability of dynamic functional connectivity methods’, Neuroimage **158**, 155–175.

- Cisler, J. M. (2017), 'Childhood trauma and functional connectivity between amygdala and medial prefrontal cortex: a dynamic functional connectivity and large-scale network perspective', Frontiers in Systems Neuroscience **11**, 29.
- Cohen, A. L., Fair, D. A., Dosenbach, N. U., Miezin, F. M., Dierker, D., Van Essen, D. C., Schlaggar, B. L. and Petersen, S. E. (2008), 'Defining functional areas in individual human brains using resting functional connectivity mri', Neuroimage **41**(1), 45–57.
- Cohen, J. R. (2018), 'The behavioral and cognitive relevance of time-varying, dynamic changes in functional connectivity', NeuroImage **180**, 515–525.
- Cole, M. W., Repovš, G. and Anticevic, A. (2014), 'The frontoparietal control system: a central role in mental health', The Neuroscientist **20**(6), 652–664.
- Cordes, D., Haughton, V. M., Arfanakis, K., Wendt, G. J., Turski, P. A., Moritz, C. H., Quigley, M. A. and Meyerand, M. E. (2000), 'Mapping functionally related regions of brain with functional connectivity mr imaging', American Journal of Neuroradiology **21**(9), 1636–1644.
- Craddock, R. C., Holtzheimer, P. E., Hu, X. P. and Mayberg, H. S. (2009), 'Disease state prediction from resting state functional connectivity', Magnetic resonance in Medicine **62**(6), 1619–1628.
- Craddock, R. C., James, G. A., Holtzheimer, P. E., Hu, X. P. and Mayberg, H. S. (2012), 'A whole brain fmri atlas generated via spatially constrained spectral clustering', Human brain mapping **33**(8), 1914–1928.
- Cuthbert, B. N. and Insel, T. R. (2013), 'Toward the future of psychiatric diagnosis: the seven pillars of rdoc', BMC medicine **11**(1), 126.

- Dahl, J., Vandenberghe, L. and Roychowdhury, V. (2008), ‘Covariance selection for nonchordal graphs via chordal embedding’, Optimization Methods & Software **23**(4), 501–520.
- Damaraju, E., Allen, E. A., Belger, A., Ford, J., McEwen, S., Mathalon, D., Mueller, B., Pearlson, G., Potkin, S., Preda, A. et al. (2014a), ‘Dynamic functional connectivity analysis reveals transient states of dysconnectivity in schizophrenia’, NeuroImage: Clinical **5**, 298–308.
- Damaraju, E., Allen, E. A., Belger, A., Ford, J., McEwen, S., Mathalon, D., Mueller, B., Pearlson, G., Potkin, S., Preda, A. et al. (2014b), ‘Dynamic functional connectivity analysis reveals transient states of dysconnectivity in schizophrenia’, NeuroImage: Clinical **5**, 298–308.
- Damoiseaux, J., Rombouts, S., Barkhof, F., Scheltens, P., Stam, C., Smith, S. M. and Beckmann, C. (2006), ‘Consistent resting-state networks across healthy subjects’, Proceedings of the national academy of sciences **103**(37), 13848–13853.
- Damoiseaux, J. S. and Greicius, M. D. (2009), ‘Greater than the sum of its parts: a review of studies combining structural connectivity and resting-state functional connectivity’, Brain Structure and Function **213**(6), 525–533.
- Daniels, J. K., Frewen, P., McKinnon, M. C. and Lanius, R. A. (2011), ‘Default mode alterations in posttraumatic stress disorder related to early-life trauma: a developmental perspective’, Journal of psychiatry & neuroscience: JPN **36**(1), 56.
- Deco, G. and Corbetta, M. (2011), ‘The dynamical balance of the brain at rest’, The Neuroscientist **17**(1), 107–123.
- Deen, B., Pitskel, N. B. and Pelphrey, K. A. (2010), ‘Three systems of insular functional connectivity identified with cluster analysis’, Cerebral cortex **21**(7), 1498–1506.

- Desikan, R. S., Ségonne, F., Fischl, B., Quinn, B. T., Dickerson, B. C., Blacker, D., Buckner, R. L., Dale, A. M., Maguire, R. P., Hyman, B. T. et al. (2006), ‘An automated labeling system for subdividing the human cerebral cortex on mri scans into gyral based regions of interest’, Neuroimage **31**(3), 968–980.
- Deuker, L., Bullmore, E. T., Smith, M., Christensen, S., Nathan, P. J., Rockstroh, B. and Bassett, D. S. (2009), ‘Reproducibility of graph metrics of human brain functional networks’, Neuroimage **47**(4), 1460–1468.
- Di, X. and Biswal, B. B. (2015), ‘Dynamic brain functional connectivity modulated by resting-state networks’, Brain Structure and Function **220**(1), 37–46.
- Dohmatob, E., Mensch, A., Varoquaux, G. and Thirion, B. (2016), Learning brain regions via large-scale online structured sparse dictionary learning, in ‘Advances in Neural Information Processing Systems’, pp. 4610–4618.
- Dong, Q., Welsh, R. C., Chenevert, T. L., Carlos, R. C., Maly-Sundgren, P., Gomez-Hassan, D. M. and Mukherji, S. K. (2004), ‘Clinical applications of diffusion tensor imaging’, Journal of Magnetic Resonance Imaging: An Official Journal of the International Society for Magnetic Resonance in Medicine **19**(1), 6–18.
- Drysdale, A. T., Grosenick, L., Downar, J., Dunlop, K., Mansouri, F., Meng, Y., Fetcho, R. N., Zebley, B., Oathes, D. J., Etkin, A. et al. (2017), ‘Resting-state connectivity biomarkers define neurophysiological subtypes of depression’, Nature medicine **23**(1), 28.
- Duda, J. T., Cook, P. A. and Gee, J. C. (2014), ‘Reproducibility of graph metrics of human brain structural networks’, Frontiers in neuroinformatics **8**.
- Dunlop, B. W., Kelley, M. E., Aponte-Rivera, V., Mletzko-Crowe, T., Kinkead, B., Ritchie, J. C., Nemeroff, C. B., Craighead, W. E., Mayberg, H. S. and Team, P.

- (2017), ‘Effects of patient preferences on outcomes in the predictors of remission in depression to individual and combined treatments (predict) study’, American Journal of Psychiatry **174**(6), 546–556.
- Eavani, H., Filipovych, R., Davatzikos, C., Satterthwaite, T. D., Gur, R. E. and Gur, R. C. (2012), Sparse dictionary learning of resting state fmri networks, in ‘Pattern Recognition in NeuroImaging (PRNI), 2012 International Workshop on’, IEEE, pp. 73–76.
- Eavani, H., Satterthwaite, T. D., Filipovych, R., Gur, R. E., Gur, R. C. and Davatzikos, C. (2015), ‘Identifying sparse connectivity patterns in the brain using resting-state fmri’, Neuroimage **105**, 286–299.
- Engle, R. (2002), ‘Dynamic conditional correlation: A simple class of multivariate generalized autoregressive conditional heteroskedasticity models’, Journal of Business & Economic Statistics **20**(3), 339–350.
- Estrada, E. and Rodriguez-Velazquez, J. A. (2005), ‘Subgraph centrality in complex networks’, Physical Review E **71**(5), 056103.
- Eyre, H. A., Yang, H., Leaver, A. M., Van Dyk, K., Siddarth, P., Cyr, N. S., Narr, K., Ercoli, L., Baune, B. T. and Lavretsky, H. (2016), ‘Altered resting-state functional connectivity in late-life depression: a cross-sectional study’, Journal of affective disorders **189**, 126–133.
- Fallani, F. D. V., Richiardi, J., Chavez, M. and Achard, S. (2014), ‘Graph analysis of functional brain networks: practical issues in translational neuroscience’, Phil. Trans. R. Soc. B **369**(1653), 20130521.
- Fischer, A. S., Keller, C. J. and Etkin, A. (2016), ‘The clinical applicability of functional connectivity in depression: Pathways toward more targeted intervention’, Biological Psychiatry: Cognitive Neuroscience and Neuroimaging **1**(3), 262–270.

- Fischl, B., Salat, D. H., Busa, E., Albert, M., Dieterich, M., Haselgrove, C., Van Der Kouwe, A., Killiany, R., Kennedy, D., Klaveness, S. et al. (2002), ‘Whole brain segmentation: automated labeling of neuroanatomical structures in the human brain’, Neuron **33**(3), 341–355.
- Fischl, B., Salat, D. H., Van Der Kouwe, A. J., Makris, N., Ségonne, F., Quinn, B. T. and Dale, A. M. (2004), ‘Sequence-independent segmentation of magnetic resonance images’, Neuroimage **23**, S69–S84.
- Flory, J. D. and Yehuda, R. (2015), ‘Comorbidity between post-traumatic stress disorder and major depressive disorder: alternative explanations and treatment considerations’, Dialogues in clinical neuroscience **17**(2), 141.
- Fornito, A., Zalesky, A. and Breakspear, M. (2013), ‘Graph analysis of the human connectome: promise, progress, and pitfalls’, Neuroimage **80**, 426–444.
- Fornito, A., Zalesky, A. and Bullmore, E. T. (2010), ‘Network scaling effects in graph analytic studies of human resting-state fmri data’, Frontiers in systems neuroscience **4**, 22.
- Frazier, J. A., Chiu, S., Breeze, J. L., Makris, N., Lange, N., Kennedy, D. N., Herbert, M. R., Bent, E. K., Koneru, V. K., Dieterich, M. E. et al. (2005), ‘Structural brain magnetic resonance imaging of limbic and thalamic volumes in pediatric bipolar disorder’, American Journal of Psychiatry **162**(7), 1256–1265.
- Frederikse, M. E., Lu, A., Aylward, E., Barta, P. and Pearlson, G. (1999), ‘Sex differences in the inferior parietal lobule’, Cerebral Cortex **9**(8), 896–901.
- Friedman, J., Hastie, T. and Tibshirani, R. (2008), ‘Sparse inverse covariance estimation with the graphical lasso’, Biostatistics **9**(3), 432–441.

- Friston, K. (2009), 'Causal modelling and brain connectivity in functional magnetic resonance imaging', PLoS biology **7**(2), e1000033.
- Friston, K., Buechel, C., Fink, G., Morris, J., Rolls, E. and Dolan, R. J. (1997), 'Psychophysiological and modulatory interactions in neuroimaging', Neuroimage **6**(3), 218–229.
- Friston, K. J. (1994), 'Functional and effective connectivity in neuroimaging: a synthesis', Human brain mapping **2**(1-2), 56–78.
- Friston, K. J., Harrison, L. and Penny, W. (2003), 'Dynamic causal modelling', Neuroimage **19**(4), 1273–1302.
- Fu, C. H., Williams, S. C., Cleare, A. J., Brammer, M. J., Walsh, N. D., Kim, J., Andrew, C. M., Pich, E. M., Williams, P. M., Reed, L. J. et al. (2004), 'Attenuation of the neural response to sad faces in major depression by antidepressant treatment: a prospective, event-related functional magnetic resonance imaging study', Archives of general psychiatry **61**(9), 877–889.
- Ghosh, A., Rho, Y., McIntosh, A. R., Kötter, R. and Jirsa, V. K. (2008), 'Noise during rest enables the exploration of the brain's dynamic repertoire', PLoS computational biology **4**(10), e1000196.
- Ginestet, C. E., Fournel, A. P. and Simmons, A. (2014), 'Statistical network analysis for functional mri: summary networks and group comparisons', Frontiers in computational neuroscience **8**, 51.
- Glasser, M. F., Coalson, T. S., Robinson, E. C., Hacker, C. D., Harwell, J., Yacoub, E., Ugurbil, K., Andersson, J., Beckmann, C. F., Jenkinson, M. et al. (2016), 'A multi-modal parcellation of human cerebral cortex', Nature **536**(7615), 171–178.

- Goldsmith, J., Wand, M. P. and Crainiceanu, C. (2011), ‘Functional regression via variational bayes’, Electronic journal of statistics **5**, 572.
- Gordon, E. M., Laumann, T. O., Adeyemo, B., Huckins, J. F., Kelley, W. M. and Petersen, S. E. (2014), ‘Generation and evaluation of a cortical area parcellation from resting-state correlations’, Cerebral cortex **26**(1), 288–303.
- Granger, C. W. (1969), ‘Investigating causal relations by econometric models and cross-spectral methods’, Econometrica: Journal of the Econometric Society pp. 424–438.
- Greicius, M. D., Flores, B. H., Menon, V., Glover, G. H., Solvason, H. B., Kenna, H., Reiss, A. L. and Schlaggar, B. L. (2007a), ‘Resting-state functional connectivity in major depression: abnormally increased contributions from subgenual cingulate cortex and thalamus’, Biological psychiatry **62**(5), 429–437.
- Greicius, M. D., Flores, B. H., Menon, V., Glover, G. H., Solvason, H. B., Kenna, H., Reiss, A. L. and Schlaggar, B. L. (2007b), ‘Resting-state functional connectivity in major depression: abnormally increased contributions from subgenual cingulate cortex and thalamus’, Biological psychiatry **62**(5), 429–437.
- Greicius, M. D., Krasnow, B., Reiss, A. L. and Menon, V. (2003), ‘Functional connectivity in the resting brain: a network analysis of the default mode hypothesis’, Proceedings of the National Academy of Sciences **100**(1), 253–258.
- Greicius, M. D., Supekar, K., Menon, V. and Dougherty, R. F. (2009), ‘Resting-state functional connectivity reflects structural connectivity in the default mode network’, Cerebral cortex **19**(1), 72–78.
- Griffanti, L., Rolinski, M., Szewczyk-Krolikowski, K., Menke, R. A., Filippini, N., Zamboni, G., Jenkinson, M., Hu, M. T. and Mackay, C. E. (2016), ‘Challenges in

- the reproducibility of clinical studies with resting state fmri: An example in early parkinson's disease', Neuroimage **124**, 704–713.
- Grotegerd, D., Suslow, T., Bauer, J., Ohrmann, P., Arolt, V., Stuhrmann, A., Heindel, W., Kugel, H. and Dannlowski, U. (2013), 'Discriminating unipolar and bipolar depression by means of fmri and pattern classification: a pilot study', European archives of psychiatry and clinical neuroscience **263**(2), 119–131.
- Guo, C. C., Kurth, F., Zhou, J., Mayer, E. A., Eickhoff, S. B., Kramer, J. H. and Seeley, W. W. (2012), 'One-year test–retest reliability of intrinsic connectivity network fmri in older adults', Neuroimage **61**(4), 1471–1483.
- Gur, R. C., Richard, J., Calkins, M. E., Chiavacci, R., Hansen, J. A., Bilker, W. B., Loughead, J., Connolly, J. J., Qiu, H., Mentch, F. D. et al. (2012), 'Age group and sex differences in performance on a computerized neurocognitive battery in children age 8- 21.', Neuropsychology **26**(2), 251.
- Gur, R. C., Turetsky, B. I., Matsui, M., Yan, M., Bilker, W., Hughett, P. and Gur, R. E. (1999), 'Sex differences in brain gray and white matter in healthy young adults: correlations with cognitive performance', Journal of Neuroscience **19**(10), 4065–4072.
- Haimovici, A., Tagliazucchi, E., Balenzuela, P. and Laufs, H. (2017), 'On wakefulness fluctuations as a source of bold functional connectivity dynamics', Scientific reports **7**(1), 5908.
- Halpern, D. F. (2000), Sex Differences in Cognitive Abilities, Psychology Press.
- Halpern, D. F., Benbow, C. P., Geary, D. C., Gur, R. C., Hyde, J. S. and Gernsbacher, M. A. (2007), 'The science of sex differences in science and mathematics', Psychological science in the public interest **8**(1), 1–51.

- Handwerker, D. A., Roopchansingh, V., Gonzalez-Castillo, J. and Bandettini, P. A. (2012), 'Periodic changes in fmri connectivity', Neuroimage **63**(3), 1712–1719.
- Hasenkamp, W. and Barsalou, L. W. (2012), 'Effects of meditation experience on functional connectivity of distributed brain networks', Frontiers in human neuroscience **6**, 38.
- Hayasaka, S. and Laurienti, P. J. (2010), 'Comparison of characteristics between region-and voxel-based network analyses in resting-state fmri data', Neuroimage **50**(2), 499–508.
- Heberlein, K. A. and Hu, X. (2004), 'Simultaneous acquisition of gradient-echo and asymmetric spin-echo for single-shot z-shim: Z-saga', Magnetic resonance in medicine **51**(1), 212–216.
- Hermundstad, A. M., Bassett, D. S., Brown, K. S., Aminoff, E. M., Clewett, D., Freeman, S., Frithsen, A., Johnson, A., Tipper, C. M., Miller, M. B. et al. (2013), 'Structural foundations of resting-state and task-based functional connectivity in the human brain', Proceedings of the National Academy of Sciences **110**(15), 6169–6174.
- Hilgetag, C. C. and Goulas, A. (2016), 'Is the brain really a small-world network?', Brain Structure and Function **221**(4), 2361–2366.
- Hines, M. (2010), 'Sex-related variation in human behavior and the brain', Trends in cognitive sciences **14**(10), 448–456.
- Hinne, M., Ambrogioni, L., Janssen, R. J., Heskes, T. and van Gerven, M. A. (2014), 'Structurally-informed bayesian functional connectivity analysis', NeuroImage **86**, 294–305.

- Hinne, M., Janssen, R. J., Heskes, T. and van Gerven, M. A. (2015), ‘Bayesian estimation of conditional independence graphs improves functional connectivity estimates’, PLoS computational biology **11**(11), e1004534.
- Hipp, J. F., Engel, A. K. and Siegel, M. (2011), ‘Oscillatory synchronization in large-scale cortical networks predicts perception’, Neuron **69**(2), 387–396.
- Hirschberger, M., Qi, Y. and Steuer, R. E. (2007), ‘Randomly generating portfolio-selection covariance matrices with specified distributional characteristics’, European Journal of Operational Research **177**(3), 1610–1625.
- Honey, C. J., Thivierge, J.-P. and Sporns, O. (2010), ‘Can structure predict function in the human brain?’, Neuroimage **52**(3), 766–776.
- Honey, C., Sporns, O., Cammoun, L., Gigandet, X., Thiran, J.-P., Meuli, R. and Hagmann, P. (2009), ‘Predicting human resting-state functional connectivity from structural connectivity’, Proceedings of the National Academy of Sciences **106**(6), 2035–2040.
- Hsieh, C.-J., Dhillon, I. S., Ravikumar, P. K. and Sustik, M. A. (2011), Sparse inverse covariance matrix estimation using quadratic approximation, in ‘Advances in neural information processing systems’, pp. 2330–2338.
- Huang, S., Li, J., Sun, L., Ye, J., Fleisher, A., Wu, T., Chen, K., Reiman, E., Initiative, A. D. N. et al. (2010), ‘Learning brain connectivity of alzheimer’s disease by sparse inverse covariance estimation’, NeuroImage **50**(3), 935–949.
- Hutchison, R. M. and Morton, J. B. (2015), ‘Tracking the brain’s functional coupling dynamics over development’, Journal of Neuroscience **35**(17), 6849–6859.
- Hutchison, R. M., Womelsdorf, T., Allen, E. A., Bandettini, P. A., Calhoun, V. D., Corbetta, M., Della Penna, S., Duyn, J. H., Glover, G. H., Gonzalez-Castillo, J.

- et al. (2013), ‘Dynamic functional connectivity: promise, issues, and interpretations’, Neuroimage **80**, 360–378.
- Hwang, J., Egorova, N., Yang, X., Zhang, W., Chen, J., Yang, X., Hu, L., Sun, S., Tu, Y. and Kong, J. (2015), ‘Subthreshold depression is associated with impaired resting-state functional connectivity of the cognitive control network’, Translational psychiatry **5**(11), e683.
- Hyde, J. S. (1981), ‘How large are cognitive gender differences? a meta-analysis using w^2 and $d..$ ’, American Psychologist **36**(8), 892.
- Ingalhalikar, M., Smith, A., Parker, D., Satterthwaite, T. D., Elliott, M. A., Ruparel, K., Hakonarson, H., Gur, R. E., Gur, R. C. and Verma, R. (2014), ‘Sex differences in the structural connectome of the human brain’, Proceedings of the National Academy of Sciences **111**(2), 823–828.
- James, G. A., Kelley, M. E., Craddock, R. C., Holtzheimer, P. E., Dunlop, B. W., Nemeroff, C. B., Mayberg, H. S. and Hu, X. P. (2009), ‘Exploratory structural equation modeling of resting-state fmri: applicability of group models to individual subjects’, Neuroimage **45**(3), 778–787.
- Jin, C., Jia, H., Lanka, P., Rangaprakash, D., Li, L., Liu, T., Hu, X. and Deshpande, G. (2017), ‘Dynamic brain connectivity is a better predictor of ptsd than static connectivity’, Human brain mapping **38**(9), 4479–4496.
- Johnstone, D., Milward, E. A., Berretta, R., Moscato, P., Initiative, A. D. N. et al. (2012), ‘Multivariate protein signatures of pre-clinical alzheimer’s disease in the alzheimer’s disease neuroimaging initiative (adni) plasma proteome dataset’, PLoS one **7**(4), e34341.
- Jurafsky, D. and Martin, J. H. (2014), Speech and language processing, Vol. 3, Pearson London.

- Kaiser, R. H., Andrews-Hanna, J. R., Wager, T. D. and Pizzagalli, D. A. (2015), ‘Large-scale network dysfunction in major depressive disorder: a meta-analysis of resting-state functional connectivity’, JAMA psychiatry **72**(6), 603–611.
- Kemmer, P. B., Bowman, F. D., Mayberg, H. and Guo, Y. (2017), ‘Quantifying the strength of structural connectivity underlying functional brain networks’, arXiv preprint arXiv:1703.04056 .
- Kemmer, P. B., Wang, Y., Bowman, F. D., Mayberg, H. and Guo, Y. (2018), ‘Evaluating the strength of structural connectivity underlying brain functional networks’, Brain Connectivity **8**(10), 579–594.
- Kilaru, V., Iyer, S., Almli, L., Stevens, J., Lori, A., Jovanovic, T., Ely, T., Bradley, B., Binder, E., Koen, N. et al. (2016), ‘Genome-wide gene-based analysis suggests an association between neuroligin 1 (*nlg1*) and post-traumatic stress disorder’, Translational psychiatry **6**(5), e820.
- Kim, J., Wozniak, J. R., Mueller, B. A. and Pan, W. (2015), ‘Testing group differences in brain functional connectivity: using correlations or partial correlations?’, Brain connectivity **5**(4), 214–231.
- Kim, J., Wozniak, J. R., Mueller, B. A., Shen, X. and Pan, W. (2014a), ‘Comparison of statistical tests for group differences in brain functional networks’, NeuroImage **101**, 681–694.
- Kim, J., Wozniak, J. R., Mueller, B. A., Shen, X. and Pan, W. (2014b), ‘Comparison of statistical tests for group differences in brain functional networks’, NeuroImage **101**, 681–694.
- Koutsouleris, N., Meisenzahl, E. M., Borgwardt, S., Riecher-Rössler, A., Frodl, T., Kambeitz, J., Köhler, Y., Falkai, P., Möller, H.-J., Reiser, M. et al. (2015), ‘In-

- dividualized differential diagnosis of schizophrenia and mood disorders using neuroanatomical biomarkers', Brain **138**(7), 2059–2073.
- Kumari, V., Mitterschiffthaler, M. T., Teasdale, J. D., Malhi, G. S., Brown, R. G., Giampietro, V., Brammer, M. J., Poon, L., Simmons, A., Williams, S. C. et al. (2003), 'Neural abnormalities during cognitive generation of affect in treatment-resistant depression', Biological psychiatry **54**(8), 777–791.
- Kundu, S., Mallick, B. K., Baladandayuthapani, V. et al. (2018), 'Efficient bayesian regularization for graphical model selection', Bayesian Analysis .
- Landis, J. R. and Koch, G. G. (1977), 'The measurement of observer agreement for categorical data', biometrics pp. 159–174.
- Lang, S. and Brezger, A. (2004), 'Bayesian p-splines', Journal of computational and graphical statistics **13**(1), 183–212.
- Langer, N., Pedroni, A. and Jäncke, L. (2013a), 'The problem of thresholding in small-world network analysis', PloS one **8**(1), e53199.
- Langer, N., Pedroni, A. and Jäncke, L. (2013b), 'The problem of thresholding in small-world network analysis', PloS one **8**(1), e53199.
- Latora, V. and Marchiori, M. (2001), 'Efficient behavior of small-world networks', Physical review letters **87**(19), 198701.
- Lazarov, A., Zhu, X., Suarez-Jimenez, B., Rutherford, B. R. and Neria, Y. (2017), 'Resting-state functional connectivity of anterior and posterior hippocampus in posttraumatic stress disorder', Journal of psychiatric research **94**, 15–22.
- Le Bihan, D., Mangin, J.-F., Poupon, C., Clark, C. A., Pappata, S., Molko, N. and Chabriat, H. (2001), 'Diffusion tensor imaging: concepts and applications', Journal

- of Magnetic Resonance Imaging: An Official Journal of the International Society for Magnetic Resonance in Medicine **13**(4), 534–546.
- Lebo, M. J. and Box-Steffensmeier, J. M. (2008), ‘Dynamic conditional correlations in political science’, American Journal of Political Science **52**(3), 688–704.
- Leistedt, S. J., Coumans, N., Dumont, M., Lanquart, J.-P., Stam, C. J. and Linkowski, P. (2009), ‘Altered sleep brain functional connectivity in acutely depressed patients’, Human brain mapping **30**(7), 2207–2219.
- Lenroot, R. K. and Giedd, J. N. (2010), ‘Sex differences in the adolescent brain’, Brain and cognition **72**(1), 46–55.
- Lenroot, R. K., Gogtay, N., Greenstein, D. K., Wells, E. M., Wallace, G. L., Clasen, L. S., Blumenthal, J. D., Lerch, J., Zijdenbos, A. P., Evans, A. C. et al. (2007), ‘Sexual dimorphism of brain developmental trajectories during childhood and adolescence’, Neuroimage **36**(4), 1065–1073.
- Leonardi, N., Richiardi, J., Gschwind, M., Simioni, S., Annoni, J.-M., Schlupe, M., Vuilleumier, P. and Van De Ville, D. (2013), ‘Principal components of functional connectivity: a new approach to study dynamic brain connectivity during rest’, NeuroImage **83**, 937–950.
- Leonardi, N., Shirer, W. R., Greicius, M. D. and Van De Ville, D. (2014), ‘Disentangling dynamic networks: separated and joint expressions of functional connectivity patterns in time’, Human brain mapping **35**(12), 5984–5995.
- Leonardi, N. and Van De Ville, D. (2015), ‘On spurious and real fluctuations of dynamic functional connectivity during rest’, Neuroimage **104**, 430–436.
- Li, X., Zhu, D., Jiang, X., Jin, C., Zhang, X., Guo, L., Zhang, J., Hu, X., Li, L. and

- Liu, T. (2014), ‘Dynamic functional connectomics signatures for characterization and differentiation of ptsd patients’, Human brain mapping **35**(4), 1761–1778.
- Liang, X., Wang, J., Yan, C., Shu, N., Xu, K., Gong, G. and He, Y. (2012), ‘Effects of different correlation metrics and preprocessing factors on small-world brain functional networks: a resting-state functional mri study’, PloS one **7**(3), e32766.
- Liao, W., Zhang, Z., Pan, Z., Mantini, D., Ding, J., Duan, X., Luo, C., Lu, G. and Chen, H. (2010), ‘Altered functional connectivity and small-world in mesial temporal lobe epilepsy’, PloS one **5**(1), e8525.
- Lindquist, M. A., Xu, Y., Nebel, M. B. and Caffo, B. S. (2014), ‘Evaluating dynamic bivariate correlations in resting-state fmri: a comparison study and a new approach’, NeuroImage **101**, 531–546.
- Liu, X. and Duyn, J. H. (2013), ‘Time-varying functional network information extracted from brief instances of spontaneous brain activity’, Proceedings of the National Academy of Sciences p. 201216856.
- Liu, X., Kong, X. and Ragin, A. B. (2017), Unified and contrasting graphical lasso for brain network discovery, in ‘Proceedings of the 2017 SIAM International Conference on Data Mining’, SIAM, pp. 180–188.
- Liu, X., Ward, B. D., Binder, J. R., Li, S.-J. and Hudetz, A. G. (2014), ‘Scale-free functional connectivity of the brain is maintained in anesthetized healthy participants but not in patients with unresponsive wakefulness syndrome’, PloS one **9**(3), e92182.
- Liu, Y., Liang, M., Zhou, Y., He, Y., Hao, Y., Song, M., Yu, C., Liu, H., Liu, Z. and Jiang, T. (2008), ‘Disrupted small-world networks in schizophrenia’, Brain **131**(4), 945–961.

- Lord, A., Horn, D., Breakspear, M. and Walter, M. (2012), ‘Changes in community structure of resting state functional connectivity in unipolar depression’, PloS one **7**(8), e41282.
- Lorenzetti, V., Allen, N. B., Fornito, A. and Yücel, M. (2009), ‘Structural brain abnormalities in major depressive disorder: a selective review of recent mri studies’, Journal of affective disorders **117**(1), 1–17.
- Luo, X. (2014), ‘A hierarchical graphical model for big inverse covariance estimation with an application to fmri’, arXiv preprint arXiv:1403.4698 .
- Lynall, M.-E., Bassett, D. S., Kerwin, R., McKenna, P. J., Kitzbichler, M., Muller, U. and Bullmore, E. (2010), ‘Functional connectivity and brain networks in schizophrenia’, Journal of Neuroscience **30**(28), 9477–9487.
- Mairal, J., Bach, F., Ponce, J. and Sapiro, G. (2009), Online dictionary learning for sparse coding, in ‘Proceedings of the 26th annual international conference on machine learning’, ACM, pp. 689–696.
- Mansouri, F. A., Fehring, D. J., Gaillard, A., Jaberzadeh, S. and Parkinson, H. (2016), ‘Sex dependency of inhibitory control functions’, Biology of sex differences **7**(1), 11.
- Maslov, S. and Sneppen, K. (2002), ‘Specificity and stability in topology of protein networks’, Science **296**(5569), 910–913.
- Matthews, B. W. (1975), ‘Comparison of the predicted and observed secondary structure of t4 phage lysozyme’, Biochimica et Biophysica Acta (BBA)-Protein Structure **405**(2), 442–451.
- Mayberg, H. S., Liotti, M., Brannan, S. K., McGinnis, S., Mahurin, R. K., Jerabek, P. A., Silva, J. A., Tekell, J. L., Martin, C. C., Lancaster, J. L. et al. (1999),

- ‘Reciprocal limbic-cortical function and negative mood: converging pet findings in depression and normal sadness’, American Journal of Psychiatry **156**(5), 675–682.
- Mazumder, R. and Hastie, T. (2012), ‘The graphical lasso: New insights and alternatives’, Electronic journal of statistics **6**, 2125.
- Meinshausen, N. and Bühlmann, P. (2006), ‘High-dimensional graphs and variable selection with the lasso’, The annals of statistics pp. 1436–1462.
- Messé, A., Rudrauf, D., Benali, H. and Marrelec, G. (2014), ‘Relating structure and function in the human brain: relative contributions of anatomy, stationary dynamics, and non-stationarities’, PLoS computational biology **10**(3), e1003530.
- Mohan, A., Roberto, A. J., Mohan, A., Lorenzo, A., Jones, K., Carney, M. J., Liogier-Weyback, L., Hwang, S. and Lapidus, K. A. (2016), ‘Focus: the aging brain: the significance of the default mode network (dmn) in neurological and neuropsychiatric disorders: a review’, The Yale journal of biology and medicine **89**(1), 49.
- Monti, R. P., Hellyer, P., Sharp, D., Leech, R., Anagnostopoulos, C. and Montana, G. (2014), ‘Estimating time-varying brain connectivity networks from functional mri time series’, NeuroImage **103**, 427–443.
- Morgan, V. L., Mishra, A., Newton, A. T., Gore, J. C. and Ding, Z. (2009), ‘Integrating functional and diffusion magnetic resonance imaging for analysis of structure-function relationship in the human language network’, PLoS One **4**(8).
- Mumford, J. A. and Ramsey, J. D. (2014), ‘Bayesian networks for fmri: a primer’, Neuroimage **86**, 573–582.
- Naumova, V. and Schnass, K. (2018), ‘Fast dictionary learning from incomplete data’, EURASIP journal on advances in signal processing **2018**(1), 12.

- Newton, M. A., Noueir, A., Sarkar, D. and Ahlquist, P. (2004), ‘Detecting differential gene expression with a semiparametric hierarchical mixture method’, Biostatistics **5**(2), 155–176.
- Ng, B., Varoquaux, G., Poline, J.-B. and Thirion, B. (2012), ‘A novel sparse graphical approach for multimodal brain connectivity inference’, Medical Image Computing and Computer-Assisted Intervention–MICCAI 2012 pp. 707–714.
- Nichols, T. E. and Holmes, A. P. (2002), ‘Nonparametric permutation tests for functional neuroimaging: a primer with examples’, Human brain mapping **15**(1), 1–25.
- Niu, H., Li, Z., Liao, X., Wang, J., Zhao, T., Shu, N., Zhao, X. and He, Y. (2013), ‘Test-retest reliability of graph metrics in functional brain networks: a resting-state fmri study’, PLoS One **8**(9), e72425.
- Olshausen, B. A. and Field, D. J. (1997), ‘Sparse coding with an overcomplete basis set: A strategy employed by v1?’, Vision research **37**(23), 3311–3325.
- Onias, H., Viol, A., Palhano-Fontes, F., Andrade, K. C., Sturzbecher, M., Viswanathan, G. and de Araujo, D. B. (2014), ‘Brain complex network analysis by means of resting state fmri and graph analysis: Will it be helpful in clinical epilepsy?’, Epilepsy & Behavior **38**, 71–80.
- Padmanabhan, N., White, M., Zhou, H. H. and O’Connell, R. (2016), ‘Estimating sparse precision matrices’, Monthly Notices of the Royal Astronomical Society **460**(2), 1567–1576.
- Pan, W., Kim, J., Zhang, Y., Shen, X. and Wei, P. (2014), ‘A powerful and adaptive association test for rare variants’, Genetics **197**(4), 1081–1095.
- Pandya, M., Altinay, M., Malone, D. A. and Anand, A. (2012), ‘Where in the brain is depression?’, Current psychiatry reports **14**(6), 634–642.

- Papan Thaipisuttikul, P. I., Waleeprakhon, P., Wisajun, P. and Jullagate, S. (2014), ‘Psychiatric comorbidities in patients with major depressive disorder’, Neuropsychiatric disease and treatment **10**, 2097.
- Peng, J., Wang, P., Zhou, N. and Zhu, J. (2009), ‘Partial correlation estimation by joint sparse regression models’, Journal of the American Statistical Association **104**(486), 735–746.
- Pineda-Pardo, J. A., Bruña, R., Woolrich, M., Marcos, A., Nobre, A. C., Maestú, F. and Vidaurre, D. (2014), ‘Guiding functional connectivity estimation by structural connectivity in meg: an application to discrimination of conditions of mild cognitive impairment’, Neuroimage **101**, 765–777.
- Power, J. D., Cohen, A. L., Nelson, S. M., Wig, G. S., Barnes, K. A., Church, J. A., Vogel, A. C., Laumann, T. O., Miezin, F. M., Schlaggar, B. L. et al. (2011a), ‘Functional network organization of the human brain’, Neuron **72**(4), 665–678.
- Power, J. D., Cohen, A. L., Nelson, S. M., Wig, G. S., Barnes, K. A., Church, J. A., Vogel, A. C., Laumann, T. O., Miezin, F. M., Schlaggar, B. L. et al. (2011b), ‘Functional network organization of the human brain’, Neuron **72**(4), 665–678.
- Power, J. D., Schlaggar, B. L., Lessov-Schlaggar, C. N. and Petersen, S. E. (2013), ‘Evidence for hubs in human functional brain networks’, Neuron **79**(4), 798–813.
- Qiu, A., Lee, A., Tan, M. and Chung, M. K. (2015), ‘Manifold learning on brain functional networks in aging’, Medical image analysis **20**(1), 52–60.
- Raznahan, A., Lerch, J. P., Lee, N., Greenstein, D., Wallace, G. L., Stockman, M., Clasen, L., Shaw, P. W. and Giedd, J. N. (2011), ‘Patterns of coordinated anatomical change in human cortical development: a longitudinal neuroimaging study of maturational coupling’, Neuron **72**(5), 873–884.

- Robins, G., Pattison, P., Kalish, Y. and Lusher, D. (2007), 'An introduction to exponential random graph (p^*) models for social networks', Social networks **29**(2), 173–191.
- Roehrig, C. (2016), 'Mental disorders top the list of the most costly conditions in the united states: \$201 billion', Health Affairs **35**(6), 1130–1135.
- Rogers, B. P., Morgan, V. L., Newton, A. T. and Gore, J. C. (2007), 'Assessing functional connectivity in the human brain by fmri', Magnetic resonance imaging **25**(10), 1347–1357.
- Rosa, M. J., Portugal, L., Hahn, T., Fallgatter, A. J., Garrido, M. I., Shawe-Taylor, J. and Mourao-Miranda, J. (2015), 'Sparse network-based models for patient classification using fmri', Neuroimage **105**, 493–506.
- Rubinov, M., Knock, S. A., Stam, C. J., Micheloyannis, S., Harris, A. W., Williams, L. M. and Breakspear, M. (2009), 'Small-world properties of nonlinear brain activity in schizophrenia', Human brain mapping **30**(2), 403–416.
- Rubinov, M. and Sporns, O. (2010a), 'Complex network measures of brain connectivity: uses and interpretations', Neuroimage **52**(3), 1059–1069.
- Rubinov, M. and Sporns, O. (2010b), 'Complex network measures of brain connectivity: uses and interpretations', Neuroimage **52**(3), 1059–1069.
- Rudie, J. D., Brown, J., Beck-Pancer, D., Hernandez, L., Dennis, E., Thompson, P., Bookheimer, S. and Dapretto, M. (2013), 'Altered functional and structural brain network organization in autism', NeuroImage: clinical **2**, 79–94.
- Ruigrok, A. N., Salimi-Khorshidi, G., Lai, M.-C., Baron-Cohen, S., Lombardo, M. V., Tait, R. J. and Suckling, J. (2014), 'A meta-analysis of sex differences in human brain structure', Neuroscience & Biobehavioral Reviews **39**, 34–50.

- Rydén, T. et al. (2008), ‘Em versus markov chain monte carlo for estimation of hidden markov models: A computational perspective’, Bayesian Analysis **3**(4), 659–688.
- Sacher, J., Neumann, J., Okon-Singer, H., Gotowiec, S. and Villringer, A. (2013), ‘Sexual dimorphism in the human brain: evidence from neuroimaging’, Magnetic resonance imaging **31**(3), 366–375.
- Sadeghi, M., Babaie-Zadeh, M. and Jutten, C. (2014), ‘Learning overcomplete dictionaries based on atom-by-atom updating’, IEEE Transactions on Signal Processing **62**(4), 883–891.
- Salvador, R., Suckling, J., Coleman, M. R., Pickard, J. D., Menon, D. and Bullmore, E. (2005), ‘Neurophysiological architecture of functional magnetic resonance images of human brain’, Cerebral cortex **15**(9), 1332–1342.
- Satterthwaite, T. D., Connolly, J. J., Ruparel, K., Calkins, M. E., Jackson, C., Elliott, M. A., Roalf, D. R., Hopson, R., Prabhakaran, K., Behr, M. et al. (2016), ‘The philadelphia neurodevelopmental cohort: a publicly available resource for the study of normal and abnormal brain development in youth’, Neuroimage **124**, 1115–1119.
- Satterthwaite, T. D., Elliott, M. A., Ruparel, K., Loughead, J., Prabhakaran, K., Calkins, M. E., Hopson, R., Jackson, C., Keefe, J., Riley, M. et al. (2014), ‘Neuroimaging of the philadelphia neurodevelopmental cohort’, Neuroimage **86**, 544–553.
- Satterthwaite, T. D., Wolf, D. H., Roalf, D. R., Ruparel, K., Erus, G., Vandekar, S., Gennatas, E. D., Elliott, M. A., Smith, A., Hakonarson, H. et al. (2014a), ‘Linked sex differences in cognition and functional connectivity in youth’, Cerebral cortex **25**(9), 2383–2394.
- Satterthwaite, T. D., Wolf, D. H., Roalf, D. R., Ruparel, K., Erus, G., Vandekar, S., Gennatas, E. D., Elliott, M. A., Smith, A., Hakonarson, H. et al. (2014b), ‘Linked

- sex differences in cognition and functional connectivity in youth', Cerebral cortex **25**(9), 2383–2394.
- Scher, C. D., McCreary, D. R., Asmundson, G. J. and Resick, P. A. (2008), 'The structure of post-traumatic stress disorder symptoms in three female trauma samples: A comparison of interview and self-report measures', Journal of Anxiety Disorders **22**(7), 1137–1145.
- Schweitzer, I., Tuckwell, V., Ames, D. and O'brien, J. (2001), 'Structural neuroimaging studies in late-life depression: a review', The World Journal of Biological Psychiatry **2**(2), 83–88.
- Scrucca, L., Fop, M., Murphy, T. B. and Raftery, A. E. (2016), 'mclust 5: clustering, classification and density estimation using Gaussian finite mixture models', The R Journal **8**(1), 205–233.
URL: <https://journal.r-project.org/archive/2016-1/scrucca-fop-murphy-et-al.pdf>
- Serpa, M. H., Ou, Y., Schaufelberger, M. S., Doshi, J., Ferreira, L. K., Machado-Vieira, R., Menezes, P. R., Scazufca, M., Davatzikos, C., Busatto, G. F. et al. (2014), 'Neuroanatomical classification in a population-based sample of psychotic major depression and bipolar i disorder with 1 year of diagnostic stability', BioMed research international **2014**.
- Shaw, P., Kabani, N. J., Lerch, J. P., Eckstrand, K., Lenroot, R., Gogtay, N., Greenstein, D., Clasen, L., Evans, A., Rapoport, J. L. et al. (2008), 'Neurodevelopmental trajectories of the human cerebral cortex', Journal of Neuroscience **28**(14), 3586–3594.
- Sheline, Y. I., Gado, M. H. and Price, J. L. (1998), 'Amygdala core nuclei volumes are decreased in recurrent major depression', Neuroreport **9**(9), 2023–2028.

- Shen, K., Mišić, B., Cipollini, B. N., Bezgin, G., Buschkuehl, M., Hutchison, R. M., Jaeggi, S. M., Kross, E., Peltier, S. J., Everling, S. et al. (2015), ‘Stable long-range interhemispheric coordination is supported by direct anatomical projections’, Proceedings of the National Academy of Sciences **112**(20), 6473–6478.
- Shen, X., Tokoglu, F., Papademetris, X. and Constable, R. T. (2013), ‘Groupwise whole-brain parcellation from resting-state fmri data for network node identification’, Neuroimage **82**, 403–415.
- Shine, J. M., Bissett, P. G., Bell, P. T., Koyejo, O., Balsters, J. H., Gorgolewski, K. J., Moodie, C. A. and Poldrack, R. A. (2016), ‘The dynamics of functional brain networks: integrated network states during cognitive task performance’, Neuron **92**(2), 544–554.
- Shine, J. M., Koyejo, O., Bell, P. T., Gorgolewski, K. J., Gilat, M. and Poldrack, R. A. (2015), ‘Estimation of dynamic functional connectivity using multiplication of temporal derivatives’, NeuroImage **122**, 399–407.
- Shrout, P. E. and Fleiss, J. L. (1979), ‘Intraclass correlations: uses in assessing rater reliability.’, Psychological bulletin **86**(2), 420.
- Simpson, S. L., Hayasaka, S. and Laurienti, P. J. (2011), ‘Exponential random graph modeling for complex brain networks’, PloS one **6**(5), e20039.
- Simpson, S. L., Lyday, R. G., Hayasaka, S., Marsh, A. P. and Laurienti, P. J. (2013a), ‘A permutation testing framework to compare groups of brain networks’, Frontiers in computational neuroscience **7**.
- Simpson, S. L., Lyday, R. G., Hayasaka, S., Marsh, A. P. and Laurienti, P. J. (2013b), ‘A permutation testing framework to compare groups of brain networks’, Frontiers in computational neuroscience **7**, 171.

- Simpson, S. L., Moussa, M. N. and Laurienti, P. J. (2012), 'An exponential random graph modeling approach to creating group-based representative whole-brain connectivity networks', Neuroimage **60**(2), 1117–1126.
- Skudlarski, P., Jagannathan, K., Calhoun, V. D., Hampson, M., Skudlarska, B. A. and Pearlson, G. (2008), 'Measuring brain connectivity: diffusion tensor imaging validates resting state temporal correlations', Neuroimage **43**(3), 554–561.
- Smith, S. M., Fox, P. T., Miller, K. L., Glahn, D. C., Fox, P. M., Mackay, C. E., Filippini, N., Watkins, K. E., Toro, R., Laird, A. R. et al. (2009a), 'Correspondence of the brain's functional architecture during activation and rest', Proceedings of the National Academy of Sciences **106**(31), 13040–13045.
- Smith, S. M., Fox, P. T., Miller, K. L., Glahn, D. C., Fox, P. M., Mackay, C. E., Filippini, N., Watkins, K. E., Toro, R., Laird, A. R. et al. (2009b), 'Correspondence of the brain's functional architecture during activation and rest', Proceedings of the National Academy of Sciences **106**(31), 13040–13045.
- Smith, S. M., Miller, K. L., Moeller, S., Xu, J., Auerbach, E. J., Woolrich, M. W., Beckmann, C. F., Jenkinson, M., Andersson, J., Glasser, M. F. et al. (2012a), 'Temporally-independent functional modes of spontaneous brain activity', Proceedings of the National Academy of Sciences **109**(8), 3131–3136.
- Smith, S. M., Miller, K. L., Moeller, S., Xu, J., Auerbach, E. J., Woolrich, M. W., Beckmann, C. F., Jenkinson, M., Andersson, J., Glasser, M. F. et al. (2012b), 'Temporally-independent functional modes of spontaneous brain activity', Proceedings of the National Academy of Sciences **109**(8), 3131–3136.
- Smith, S. M., Miller, K. L., Salimi-Khorshidi, G., Webster, M., Beckmann, C. F., Nichols, T. E., Ramsey, J. D. and Woolrich, M. W. (2011), 'Network modelling methods for fmri', Neuroimage **54**(2), 875–891.

- Soares, J., Marques, P., Alves, V. and Sousa, N. (2013), ‘A hitchhiker’s guide to diffusion tensor imaging’, Frontiers in neuroscience **7**, 31.
- Song, M., Zhou, Y., Li, J., Liu, Y., Tian, L., Yu, C. and Jiang, T. (2008), ‘Brain spontaneous functional connectivity and intelligence’, Neuroimage **41**(3), 1168–1176.
- Sporns, O. (2013), ‘Structure and function of complex brain networks’, Dialogues in clinical neuroscience **15**(3), 247.
- Stam, C., Jones, B., Nolte, G., Breakspear, M. and Scheltens, P. (2006), ‘Small-world networks and functional connectivity in alzheimer’s disease’, Cerebral cortex **17**(1), 92–99.
- Sui, J., Pearlson, G., Caprihan, A., Adali, T., Kiehl, K. A., Liu, J., Yamamoto, J. and Calhoun, V. D. (2011), ‘Discriminating schizophrenia and bipolar disorder by fusing fmri and dti in a multimodal cca+ joint ica model’, Neuroimage **57**(3), 839–855.
- Sun, Y., Collinson, S. L., Suckling, J. and Sim, K. (n.d.), ‘Dynamic reorganization of functional connectivity reveals abnormal temporal efficiency in schizophrenia’, Schizophrenia Bulletin .
- Telesford, Q. K., Joyce, K. E., Hayasaka, S., Burdette, J. H. and Laurienti, P. J. (2011), ‘The ubiquity of small-world networks’, Brain connectivity **1**(5), 367–375.
- Telesford, Q. K., Morgan, A. R., Hayasaka, S., Simpson, S. L., Barret, W., Kraft, R. A., Mozolic, J. L. and Laurienti, P. J. (2010), ‘Reproducibility of graph metrics in fmri networks’, Frontiers in neuroinformatics **4**.
- Tian, L., Wang, J., Yan, C. and He, Y. (2011), ‘Hemisphere-and gender-related differences in small-world brain networks: a resting-state functional mri study’, Neuroimage **54**(1), 191–202.

- Tibshirani, R. (1996), ‘Regression shrinkage and selection via the lasso’, Journal of the Royal Statistical Society. Series B (Methodological) pp. 267–288.
- Tomasi, D. and Volkow, N. D. (2012), ‘Gender differences in brain functional connectivity density’, Human brain mapping **33**(4), 849–860.
- Tyszka, J. M., Kennedy, D. P., Paul, L. K. and Adolphs, R. (2013), ‘Largely typical patterns of resting-state functional connectivity in high-functioning adults with autism’, Cerebral cortex **24**(7), 1894–1905.
- Tyszka, J. M. and Pauli, W. M. (2016), ‘In vivo delineation of subdivisions of the human amygdaloid complex in a high-resolution group template’, Human brain mapping **37**(11), 3979–3998.
- Tzourio-Mazoyer, N., Landeau, B., Papathanassiou, D., Crivello, F., Etard, O., Delcroix, N., Mazoyer, B. and Joliot, M. (2002), ‘Automated anatomical labeling of activations in spm using a macroscopic anatomical parcellation of the mni mri single-subject brain’, Neuroimage **15**(1), 273–289.
- Van Den Heuvel, M. P., Mandl, R. C., Kahn, R. S., Pol, H. and Hilleke, E. (2009), ‘Functionally linked resting-state networks reflect the underlying structural connectivity architecture of the human brain’, Human brain mapping **30**(10), 3127–3141.
- Van Den Heuvel, M. P. and Pol, H. E. H. (2010), ‘Exploring the brain network: a review on resting-state fmri functional connectivity’, European neuropsychopharmacology **20**(8), 519–534.
- Varoquaux, G., Gramfort, A., Poline, J.-B. and Thirion, B. (2010), Brain covariance selection: better individual functional connectivity models using population prior, in ‘Advances in neural information processing systems’, pp. 2334–2342.

- Venkataraman, A., Rathi, Y., Kubicki, M., Westin, C.-F. and Golland, P. (2012), ‘Joint modeling of anatomical and functional connectivity for population studies’, IEEE transactions on medical imaging **31**(2), 164–182.
- Vértés, P. E., Alexander-Bloch, A. F., Gogtay, N., Giedd, J. N., Rapoport, J. L. and Bullmore, E. T. (2012), ‘Simple models of human brain functional networks’, Proceedings of the National Academy of Sciences **109**(15), 5868–5873.
- Vidaurre, D., Smith, S. M. and Woolrich, M. W. (2017), ‘Brain network dynamics are hierarchically organized in time’, Proceedings of the National Academy of Sciences **114**(48), 12827–12832.
- Volz, E. (2004), ‘Random networks with tunable degree distribution and clustering’, Physical Review E **70**(5), 056115.
- Wang, H. (2012), ‘Bayesian graphical lasso models and efficient posterior computation’, Bayesian Analysis **7**(4), 867–886.
- Wang, J.-H., Zuo, X.-N., Gohel, S., Milham, M. P., Biswal, B. B. and He, Y. (2011), ‘Graph theoretical analysis of functional brain networks: test-retest evaluation on short-and long-term resting-state functional mri data’, PloS one **6**(7), e21976.
- Wang, J., Wang, X., Xia, M., Liao, X., Evans, A. and He, Y. (2015a), ‘Gretna: a graph theoretical network analysis toolbox for imaging connectomics’, Frontiers in human neuroscience **9**, 386.
- Wang, J., Wang, X., Xia, M., Liao, X., Evans, A. and He, Y. (2015b), ‘Gretna: a graph theoretical network analysis toolbox for imaging connectomics’, Frontiers in human neuroscience **9**.
- Wang, Y., Kang, J., Kemmer, P. B. and Guo, Y. (2016), ‘An efficient and reliable sta-

- tistical method for estimating functional connectivity in large scale brain networks using partial correlation', Frontiers in neuroscience **10**.
- Watts, D. J. and Strogatz, S. H. (1998), 'Collective dynamics of "small-world" networks', nature **393**(6684), 440–442.
- Weiss, D. S. (2004), 'Structured clinical interview techniques for ptsd', Assessing psychological trauma and PTSD **2**, 103–21.
- Welton, T., Kent, D. A., Auer, D. P. and Dineen, R. A. (2015), 'Reproducibility of graph-theoretic brain network metrics: a systematic review', Brain connectivity **5**(4), 193–202.
- Welvaert, M., Durnez, J., Moerkerke, B., Verdoolaege, G. and Rosseel, Y. (2011), 'neurosim: An r package for generating fmri data', Journal of Statistical Software **44**(10), 1–18.
- Whittaker, J. (1990), Graphical models in applied multivariate statistics, Wiley Publishing.
- Williams, C. K. and Rasmussen, C. E. (2006), Gaussian processes for machine learning, Vol. 2, MIT Press Cambridge, MA.
- Woolrich, M. W., Baker, A., Luckhoo, H., Mohseni, H., Barnes, G., Brookes, M. and Rezek, I. (2013), 'Dynamic state allocation for meg source reconstruction', Neuroimage **77**, 77–92.
- Wozniak, J. R., Mueller, B. A., Bell, C. J., Muetzel, R. L., Hoecker, H. L., Boys, C. J. and Lim, K. O. (2013), 'Global functional connectivity abnormalities in children with fetal alcohol spectrum disorders', Alcoholism: Clinical and experimental research **37**(5), 748–756.

- Wu, K., Taki, Y., Sato, K., Hashizume, H., Sassa, Y., Takeuchi, H., Thyreau, B., He, Y., Evans, A. C., Li, X. et al. (2013), ‘Topological organization of functional brain networks in healthy children: differences in relation to age, sex, and intelligence’, PloS one **8**(2), e55347.
- Xia, M., Wang, J. and He, Y. (2013), ‘Brainnet viewer: a network visualization tool for human brain connectomics’, PloS one **8**(7), e68910.
- Xue, W., Bowman, F. D., Pileggi, A. V. and Mayer, A. R. (2015), ‘A multimodal approach for determining brain networks by jointly modeling functional and structural connectivity’, Frontiers in computational neuroscience **9**.
- Yaesoubi, M., Adalı, T. and Calhoun, V. D. (2018), ‘A window-less approach for capturing time-varying connectivity in f mri data reveals the presence of states with variable rates of change’, Human brain mapping **39**(4), 1626–1636.
- Yang, L., Fang, J. and Li, H. (2016), Sparse bayesian dictionary learning with a gaussian hierarchical model, in ‘2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)’, IEEE, pp. 2564–2568.
- Yeo, B. T., Krienen, F. M., Sepulcre, J., Sabuncu, M. R., Lashkari, D., Hollinshead, M., Roffman, J. L., Smoller, J. W., Zöllei, L., Polimeni, J. R. et al. (2011), ‘The organization of the human cerebral cortex estimated by intrinsic functional connectivity’, Journal of neurophysiology **106**(3), 1125–1165.
- Yu, Y., Shen, H., Zeng, L.-L., Ma, Q. and Hu, D. (2013), ‘Convergent and divergent functional connectivity patterns in schizophrenia and depression’, PLoS One **8**(7), e68250.
- Yuan, J., Li, X., Zhang, J., Luo, L., Dong, Q., Lv, J., Zhao, Y., Jiang, X., Zhang, S., Zhang, W. et al. (2017), ‘Spatio-temporal modeling of connectome-scale brain network interactions via time-evolving graphs’, NeuroImage .

- Yuan, M. and Lin, Y. (2006), ‘Model selection and estimation in regression with grouped variables’, Journal of the Royal Statistical Society: Series B (Statistical Methodology) **68**(1), 49–67.
- Yuan, M. and Lin, Y. (2007), ‘Model selection and estimation in the gaussian graphical model’, Biometrika **94**(1), 19–35.
- Zalesky, A. and Breakspear, M. (2015), ‘Towards a statistical test for functional connectivity dynamics’, Neuroimage **114**, 466–470.
- Zalesky, A., Cocchi, L., Fornito, A., Murray, M. M. and Bullmore, E. (2012), ‘Connectivity differences in brain networks’, Neuroimage **60**(2), 1055–1062.
- Zalesky, A., Fornito, A. and Bullmore, E. (2012), ‘On the use of correlation as a measure of network connectivity’, Neuroimage **60**(4), 2096–2106.
- Zalesky, A., Fornito, A. and Bullmore, E. T. (2010), ‘Network-based statistic: identifying differences in brain networks’, Neuroimage **53**(4), 1197–1207.
- Zeng, L.-L., Shen, H., Liu, L., Wang, L., Li, B., Fang, P., Zhou, Z., Li, Y. and Hu, D. (2012), ‘Identifying major depression using whole-brain functional connectivity: a multivariate pattern analysis’, Brain **135**(5), 1498–1507.
- Zhang, J., Wang, J., Wu, Q., Kuang, W., Huang, X., He, Y. and Gong, Q. (2011), ‘Disrupted brain connectivity networks in drug-naive, first-episode major depressive disorder’, Biological psychiatry **70**(4), 334–342.
- Zhang, Y., Wang, Z., Cai, Z., Lin, Q. and Hu, Z. (2016), ‘Nonlinear estimation of bold signals with the aid of cerebral blood volume imaging’, Biomedical engineering online **15**(1), 22.
- Zhou, M., Chen, H., Paisley, J., Ren, L., Li, L., Xing, Z., Dunson, D., Sapiro, G. and

- Carin, L. (2012), ‘Nonparametric bayesian dictionary learning for analysis of noisy and incomplete images’, IEEE Transactions on Image Processing **21**(1), 130–144.
- Zhu, D., Li, K., Faraco, C. C., Deng, F., Zhang, D., Guo, L., Miller, L. S. and Liu, T. (2012), ‘Optimization of functional brain rois via maximization of consistency of structural connectivity profiles’, NeuroImage **59**(2), 1382–1393.
- Zhu, D., Li, K., Guo, L., Jiang, X., Zhang, T., Zhang, D., Chen, H., Deng, F., Faraco, C., Jin, C. et al. (2012), ‘Dicccol: dense individualized and common connectivity-based cortical landmarks’, Cerebral cortex **23**(4), 786–800.
- Zou, H. and Hastie, T. (2005), ‘Regularization and variable selection via the elastic net’, Journal of the Royal Statistical Society: Series B (Statistical Methodology) **67**(2), 301–320.
- Zuo, X.-N., Kelly, C., Di Martino, A., Mennes, M., Margulies, D. S., Bangaru, S., Grzadzinski, R., Evans, A. C., Zang, Y.-F., Castellanos, F. X. et al. (2010), ‘Growing together and growing apart: regional and sex differences in the lifespan developmental trajectories of functional homotopy’, Journal of Neuroscience **30**(45), 15034–15043.