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:

_____ _____

Jin Ming Date

Brain Network-Based Statistical Approaches for Neuroimaging Data

By

Jin Ming
Doctor of Philosophy
Biostatistics

Suprateek Kundu, Ph.D.
Advisor

Robert Lyles, Ph.D.
Committee Member

Joe Nocera, Ph.D.
Committee Member

Lance Waller, Ph.D.
Committee Member

Accepted:

Kimberly Jacob Arriola, Ph.D, MPH
Dean of the James T.Laney School of Graduate Studies

Date

Brain Network-Based Statistical Approaches for Neuroimaging Data

By

Jin Ming
B.A., University of Nottingham Ningbo China, China, 2014
MSPH., Emory University, GA, 2016

Advisor: 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
2022

Abstract

Brain Network-Based Statistical Approaches for Neuroimaging Data
By Jin Ming

Brain Networks derived from neuroimaging data have been widely studied recently to analyze the underlying complex spatial relationships between different regions of the brain. Functional connectivity (FC) derived from brain networks have been used as a potential biomarkers for delineating between healthy and mentally ill groups or subgroups with pre-defined features. Although there has been an intense development of statistical methods for computing brain networks, further advances are needed for developing novel and flexible statistical approaches are needed to address several gaps in the literature.

In chapter one, we propose a novel graph-theoretic approach for estimating a population of individual-specific dynamic functional connectivity that is able to systematically borrow information across multiple heterogeneous samples in a data-adaptive manner and guided by supplementary covariate information. In one of the first such approaches in literature, we develop a Bayesian product mixture model that uses covariates to model the mixture weights, which is able to cluster across heterogeneous samples independently at each time scan. An application to a fMRI block task experiment with behavioral interventions in veterans reveals sub-groups of individuals with homogeneous dynamic connectivity patterns and identifies significant dynamic network changes resulting from the intervention.

In chapter two, we proposed a novel semi-parametric Bayesian Support Vector Machine (SVM) approach that incorporates high-dimensional networks as covariates and is able to assign varying levels of shrinkage to the coefficients in an unsupervised manner via a Dirichlet process mixture of double exponential priors. Although SVM-based methods are heavily used in classifying mental disorders, there are few, if any, semi-parametric Bayesian SVM approaches for classification based on high-dimensional brain networks that naturally provides the ability to conduct inferences. We apply the approach to a connectome fingerprinting problem using the Human Connectome Project (HCP) data as well as a second application involving classification of individuals with attention deficiency hyperactivity disorder (ADHD) and showcase the superior classification accuracy of the proposed approach.

In chapter three, we examine the potential of multimodal dynamic FC, computed by fusing functional magnetic resonance imaging (fMRI) and diffusion tensor imaging data, in terms of predicting continuous clinical measures of disease severity. We develop concrete measures of temporal network variability that are directly linked with disease severity and identify regions whose temporal connectivity fluctuations are significantly related to the disease. Our results illustrate the distinct advantages of prediction of disease severity compared to the usual analysis based on disease phenotype categories, it shows that the multimodal approach is more sensitive to connectivity changes and highlights the predictive prowess of multimodal dynamic FC over existing static and dynamic network models.

Brain Network-Based Statistical Approaches for Neuroimaging Data

By

Jin Ming
B.A., University of Nottingham Ningbo China, China, 2014
MSPH., Emory University, GA, 2016

Advisor: 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
2022

Contents

List of Figures

4.4 Panels (A)-(C) illustrates the edge-wise temporal variation averaged over all individuals, individuals with PTSD and those without PTSD, respectively. Here the temporal variation for an edge was calculated as the ratio of the number of state changes for that edge divided by the number of state changes in the network. Edges in several modules including Visual, SAL, SCOR, VAN and DAN show strong temporal fluctuations resulting from frequent state changes. On the other hand, the Sensory, Cingulomotor, and DMN register the fewest temporal fluctuations over time. Only 213 edges illustrated significant differences in terms of the proportion of edge-specific state changes between the PTSD and non-PTSD groups, which suggests the inadequacy of this measure to distinguish disease severity. Out of these 213 edges, almost all had a higher frequency corresponding to the PTSD group, illustrating higher temporal fluctuations in this cohort. Panel (D) illustrates edges whose fluctuations in terms of the edge strength (measured via edge-specific standard deviations for partial correlations over time) are significantly related to PTSD resilience (multiplicity adjusted). Most of these edges lie between functional modules and are contained between the Visual and other modules, as well as between the DAN and other modules. Blue and red colors imply a negative and positive association with PTSD resilience respectively. It is clear that an increase in temporal edge strength fluctuations in most edges leads to decrease in resilience and vice-versa. However, a small number of connections within DMN and between DMN and other modules lead to increased resilience corresponding to higher fluctuations in edge strength and vice-versa. 110

List of Tables

Chapter 1

Introduction

1.1 Neuroimaging Study

Neuroimaging studies have proved to be a pivotal tool for understanding the neurobiological basis of cognitive and behavioral outcomes in psychiatric studies, as well as for examining the pathopysological mechanisms and atypical brain development underlying mental disorders. Traditional neuroimaging studies focus on identifying brain regions which are activated with respect to a certain task or outcome. However, more studies have focus on analysing the functional connectivity (FC), which measures the temporal correlation between different regions of the brain recently.

1.1.1 Brain Network

Many studies have tried to describe the nervous system as a network of interconnected neurons (Sporns, 2013; Hutchison et al., 2013). In addition, many studies have implicitly assumed that the this interconnected system between different brain regions (Functional Connectivity, FC) is constant throughout the experiments. As a quantity to measure the temporal correlation between different brain regions over time, pair-

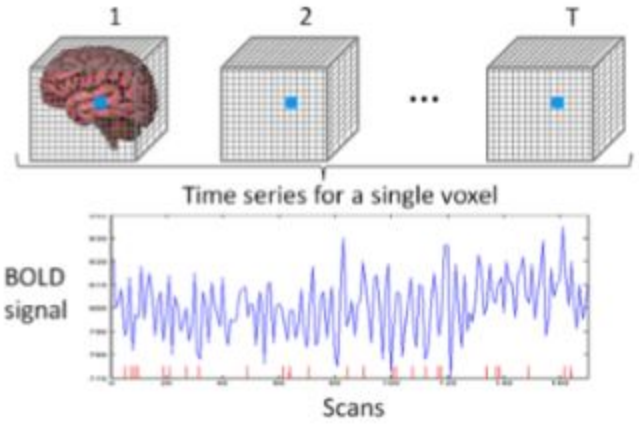

Figure 1.1: An illustration of fMRI study.

wise correlation and covariance have been widely used to measure the FC (Friston, 2011; Bressler and Menon, 2010). Recent works have shown that functional connectivity, especially resting-state functional connectivity, is related with some mental illness and psychiatric disorders. Resting-state functional connectivity could be defined as the brain network when people are at rest, which normally means that the fMRI data is acquired in the absence of a stimulus or a task. For example, many resting state FC studies have shown that cortical networks and cortical-subcortical connectivity is altered in schizophrenia (Baker et al., 2014; Woodward et al., 2012). Researchers are also trying to identify patients with major depression based on the difference in resting-sate FC (Greicius et al., 2007; Zeng et al., 2012).

There has been a steady development of approaches to compute dynamic functional connectivity (FC) that is fueled by an increasing agreement among scientists that the brain network does not remain constant across time and instead undergoes dynamic changes resulting from endogenous and exogenous factors. For example, task-related imaging studies have shown that the brain networks will re-organize when the subjects undergo different modulations of the experimental tasks during the scanning session (Chang and Glover, 2010; Lukemire et al., 2020). Similarly,

dynamic FC has also been observed during resting-state experiments (Bullmore and Sporns, 2009). In general, recent studies have found increasing evidence of underlying neuronal bases for temporal variations in FC which is linked with changes in cognitive and disease states (Hutchinson et al., 2013). Throughout this article, we use a graph-theoretic interpretation of dynamic connectivity involving time-varying correlations that is commonly used in literature.

1.1.2 Graphical Model

Graphical models have been used to describe the brain network. Graphical model usually contains two core parameters: nodes and edges. Nodes correspond to random variable, like different regions of interests (ROIs) in neuroimaging study. Edges represent the statistical dependencies between the variables. Graphical network could be either directed or undirected. In a directed network, the value of edges from two nodes might be different in order to show the different dependencies between these two nodes. However, undirected network has edges that do not have a direction. The edge indicate a two-way relationship. Edges in graphical model could also be weighted or unweighted. In a weighted graphical model, the edges might have different value, in terms to represent the strength of different edges. Unweighted models assume that all edges are equally weighted, or the dependence relationship is the same among all edges. In a graphical model, network could be transferred into matrix version. In an unweighted graphical model, edge could be either 1 or 0, whereas 1 indicates there is connection between two corresponding nodes, and 0 represent the absence of edge. Figure 1.2 is a illustration of a graphical model with 5 nodes.

In the study of fMRI, normally the undirected, unweighted graphical models are used to represent the brain network, or functional connectivity. Precision matrix has recently become one of the most widely used and successful metrics to describe brain network, as it could distinguish a true, direct functional connection from an apparent

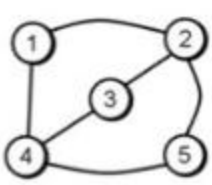

	1	2	3	4	5
1	0	1	0	1	0
2	1	0	1	0	1
3	0	1	0	1	0
4	1	0	1	0	1
5	0	1	0	1	0

Figure 1.2: Illustration of a graphical model with 5 nodes(undirected, unweighted). Left: network. Right: corresponding matrix

connection, caused via confounding by a common third party node. Although it is possible to threshold the pairwise correlation matrix to estimate brain network, the choice of threshold value might be controversial. In addition, the connection based on pairwise correlation is not the direct connection for two nodes.

In addition, the dynamic functional connectivity could be estimated via sparse precision matrix under graphical models. The Gaussian graphical model, like graphical lasso (Friedman et al., 2008) could be Incorporated directly under Bayesian framework. For example, Kundu et al., (2018) estimate the dynamic functional connectivity of each state by using graphical lasso of all time points within that time range.

1.2 Static Functional Connectivity

Many neuroimaging studies have focus on the study of static functional connectivity (sFC), which assume that the brain network keep stationary over the time of fMRI session or the experiment. However, brain regions do not show constant communication with each other but show some fluctuations over time. In other words, sFC is calculated as the average or mean connectivity over the whole fMRI session.

Normally, static functional connectivity could be captured as the pairwise correlation (Pearson‘s correlation) or partial correlation between BOLD signals (time series) of different regions of the brain. Pairwise correlation measures the connections between pairs of nodes without accounting for the third party of nodes, whereas partial corre-

lation reports connectivity that can be interpreted as the association between nodes conditional on the effects of the remaining network nodes. Both pairwise correlation and partial correlation have been widely used as the measure of static functional connectivity (Kim et al., 2015). The advantage of using pairwise correlation is that it is easy to calculate. However, pairwise correlation could not distinguish the association between a pair of nodes is direct or not. In other words, it might be the case that the association of two nodes is caused by a third node or third party of nodes.

In order to capture the direct association between two nodes, many studies have used partial correlation to measure the functional connectivity. Partial correlation quantifies the association between two regions, conditioning on the other regions. In other words, partial correlations measure the direct connection between two regions. A zero value of partial correlation indicates there is no direct connection between two regions, which could also be interpreted as conditional independence. Despite the advantage of quantifying direct association, there are some limitations of using partial correlation as sFC as well. First of all, the computation of partial correlation is much more challenging for neuroimaging data. Precision matrix is normally used for the calculation of partial correlation. Friedman et al., (2008) proposed the using for graphical lasso penalty when estimating precision matrix. However, as neuroimaing data normally is high dimensional, with over thousands of different nodes, the penalization framework took much longer time than pairwise correlation directly. Secondly, the estimation of precision matrix through penalization need the use of penazlied parameter to control the overall density of the precision matrix. But the performance of different brain network with different density level might be different as well. So researchers might need to test the performance based on all different level of densities. Finally, it is unclear that the statistical power of using partial correlation is higher than that of using pairwise correlation when testing group different.

1.3 Methods for Dynamic Functional Connectivity

Dynamic connectivity may be computed via (i) change point methods (Cribben et al., 2012; Kundu et al., 2018) that assume stable phases of connectivity interspersed with connectivity jumps at unknown locations, which results in piece-wise constant connectivity; (ii) Hidden Markov Models (HMMs) that model fast transient networks suitable for electrophysiological data (Vidaurre et al., 2016; Quinn et al., 2018) and which have also been recently applied to fMRI data (Warnick et al., 2018); and (iii) sliding window approaches that enforce temporal homogeneity in FC (Chang and Glover, 2010; Monti et al., 2014) that reflects the biologically plausible assumption of slowly varying temporal correlations resulting in gradual changes in connectivity.

1.3.1 Sliding Window Method

Arguably, the sliding window methods are most widely used to compute dynamic functional connectivity. Allen et al.(2012) combined sliding window with Independent Component Analysis (ICA) and K-means clustering to assess whole-brain dynamic FC. Sakoglu et al.(2010) applied sliding window method to a schizophrenia data-set. They have also combined time-frequency analysis.

While the sliding window methods is a valuable tool for investigating temporal dynamics of functional brain networks, however, there are some known limitations associated with this approach such as the choice of the window length (Lindquist et. al, 2014). A narrower window length may not be strong enough to capture the real underlying network structure because limited number of time scans are used, which results in poor network estimation. However, the rapid change of the functional connectivity could not be detected when using a large window size. The sliding window method with adaptive window size was proposed to solve this problem, however, their performance under real data is still not clear. Another problem relating to sliding

window methods is that computationally expensive model-based approaches relying on sliding window correlations (Monti et. al, 2014) may not be suitable for high-dimensional networks. In addition, the sliding window methods require secondary criteria to determine if variations in the edge structure are significant, which might not be straightforward to implement.

1.3.2 Change Point Method

Dynamic functional connectivity (dFC) is often conceptualized as a collection of state phases corresponding to various modulations in the brain (Hutchison et al., 2013). Change point models assume stable state phases where the transitions between consecutive state phases are defined by connectivity jumps. Gribben et al. (2012) has proposed dynamic connectivity regression (DCR), which could detect temporal change points in functional connectivity and estimates a graph, or set of relationships between ROIs, for data in the temporal partition that falls between pairs of change points. They combined graphical Lasso (Glasso) (Friedman et al., 2008) and regression tree to search for the potential change points by using the information of all subjects. In 2013, Gribben et al. (2013) has also extend the previous DCR algorithm to a single-subject version, which increases accuracy for individual subject data with a small number of observations and reduces the number of false positives in the estimated undirected graphs. In addition, Shakil et al. (2015) used the similarity measures of two images, statistical sign change (SSC), to adaptive detect change points for a single subject. Kundu et al. (2018) proposed a fully automated two-stage approach which pools information across multiple subjects to estimate change points in functional connectivity, and subsequently estimates the brain network within each state phase lying between consecutive change points.

Although existing change point methods for dynamic functional connectivity have

been somewhat successful in describing the temporal changes in the brain network, there are some existing challenges. First of all, most existing change point methods could not detect rapid changes in brain organizations, that can evolve within as little as 30-60s. Such rapid fluctuations would divide the scanning period into a collection of narrow state phases. The number of scans might be smaller than the minimal scans required to detect the change under assumption. Secondly, it is hard to pool information across heterogeneous individuals to inference the subject-level functional connectivity. DCR has both multi-subject version and single-subject version, but neither of them could solve this problem.

1.3.3 Hidden Markov Model

Hidden Markov Models (HMM) also assume a collection of state phases exists. However, HMM estimate transient states that can change instantaneously but are reinforced or revisited over time. In addition, there is no provision to ensure temporal smoothness of connectivity estimates under HMMs, which is a practical feature encountered in fMRI data as evident from the vast literature on sliding window based approaches. Hidden Markov Models have been widely applied in the studies of neuroimaging. Vidaurre et al. (2016) presents a method combined HMM with multivariate autoregressive (MAR) for EEG data. This approach models brain activity using a discrete set of sequential states, with each state distinguished by its own multi-region spectral properties. Warnick et al. (2018) applied HMM in fMRI, and they also defined a super-graph to better estimate the brain graph structure. In addition, Fan et al. (2020) propose a continuous HMM by considering the mixture of von Misea Fisher (VMF) distributions as its emission densities.

1.4 Prediction and Classification using Brain Networks

It has been a long history of using brain network do make some prediction and classification for mental illness or disorders. Many disease are considered to be related with human brain. Analysis of the brain network provides powerful abilities to aid brain disease detection. For example, existing literature try to use brain network to classify subjects with Alzheimer's disease (AD) with normal controls (Bi et al., 2020). There is great interest in understanding the neural underpinnings of individual differences in intelligence, because it is one of the most important predictors of long-term life success. Intelligence may be measured via cognitive measures that may include fluid intelligence, defined as the ability to use inductive and deductive reasoning (independent of previously acquired knowledge) to solve new problems (Kyllonen and Kell, 2017), or crystallized intelligence that involves knowledge that comes from prior learning and past experiences, among others. Such intelligence measures are assumed to be tied to brain structure; however, it is a major challenge to relate structural and functional properties of brain to complex behavioural expression or function (Bullmore and Sporns, 2009; Le Bihan et al., 2001; Raichle et al., 2001). Existing literature as used neuroimaging-derived features such as whole brain volume, regional gray and white matter volumes or regional cortical volume/thickness and diffusion indices, which may smooth over discriminative features at a finer resolution and are often inadequate predictors of intelligence (Chen, Chen, Hsu, and Tseng, 2020; Paul et al., 2016; Ritchie et al., 2018; Yuan, Voelkle, and Raz, 2018). A more recently emerging line of work has started to investigate prediction strategies for intelligence based on fMRI and derived features such as functional connectivity that has shown greater promise (Shen et al., 2017; Ferguson, Anderson, and Spreng, 2017; He et al., 2020; Kashyap et al., 2019; Liu, Liao, Xia, and He, 2018; Fan, Su, Qin, Hu, and Shen,

2020; Sen and Parhi, 2021; Liu et al., 2018; Dubois, Galdi, Paul, and Adolphs, 2018). There is increasing evidence linking temporal fluctuations in brain functional connectivity (FC) with mental disorders, which is often characterized by distinct disease phenotypes. However, the potential of dynamic FC as a neuroimaging biomarker for more heterogeneous disorders such as posttraumatic stress disorder (PTSD) that often may not have clearly demarcated phenotypes has yet to be explored. Recently, dynamic network differences have been discovered between PTSD and non-PTSD groups (Jin, et al., 2017; Fu, et al., 2019). However, disease phenotype classification, while a worthy objective, may be of secondary interest in heterogeneous mental disorders where there is no gold standard for classification. For such disorders, experts may prefer modeling continuous measures of disease progression as a clinical outcome of interest. Unfortunately, there have been limited efforts for using dynamic functional connectivity to model continuous clinical measures (Haslam, 2003; Widiger, et al., 2005).

Chapter 2

Integrative Learning for Population of Dynamic Networks with Covariates

2.1 Introduction

Essentially, almost the entirety of the existing dynamic connectivity literature has focused on data from single individuals, due to the fact that connectivity changes are expected to be subject-specific and may not be replicated across individuals. However, recent evidence suggests that combining information across individuals in a group provides more accurate estimates for connectivity (Hindriks et al., 2016), which adheres to the commonly used statistical principle of pooling information from multiple samples to obtain more robust estimates. Along these lines, Kundu et. al (2018) proposed a sub-sampling approach to compute time varying dynamic connectivity measures using multi-subject fMRI data that were then used to estimate dynamic network changes at the level of individuals for fMRI task experiments. Compared to analyses that use data separately for each individual, their approach was shown to

result in considerable gains in dynamic network estimation when the heterogeneity in the sample was limited. However in many practical settings (involving both resting state and task-based fMRI experiments), the heterogeneity across samples may not be restricted. In such cases, one needs a carefully calibrated approach for pooling information across individuals for accurately estimating a population of single-subject dynamic networks, while adequately tackling heterogeneity between subjects. To our knowledge, such approaches have received little or no attention in literature, which is likely due to considerable methodological and computational challenges involved. For example, in order to estimate dynamic networks with V nodes for N individuals each having T time scans, one needs to compute NT distinct $V \times V$ connectivity matrices, which may not be straightforward for high-dimensional fMRI data. Moreover, it is not immediately clear how to effectively borrow information across individuals in a data-adaptive manner that also respects the inherent connectivity differences between heterogeneous samples. Hence, there is a critical need to overcome such challenges via significant methodological advances.

2.2 Methods

In this section, we propose a novel approach for estimating a population of dynamic networks using heterogeneous multi-subject fMRI data with the same number of brain volumes across all individuals. For modeling purposes, we will assume that the demeaned fMRI measurements are normally distributed with zero mean (Kundu et al., 2018) at each time scan, and that pre-whitening steps have been performed to minimize temporal correlations (see Supplementary Materials for details). We will fix some notations here. Suppose fMRI data are collected for T scans and V nodes (voxels or regions of interest) for N individuals. Denote the fMRI measurements across all the nodes at time point t as $\mathbf{y}_t^{(i)} = (y_{1,t}^{(i)}, \ldots, y_{V,t}^{(i)})'$, and denote the $V \times T$

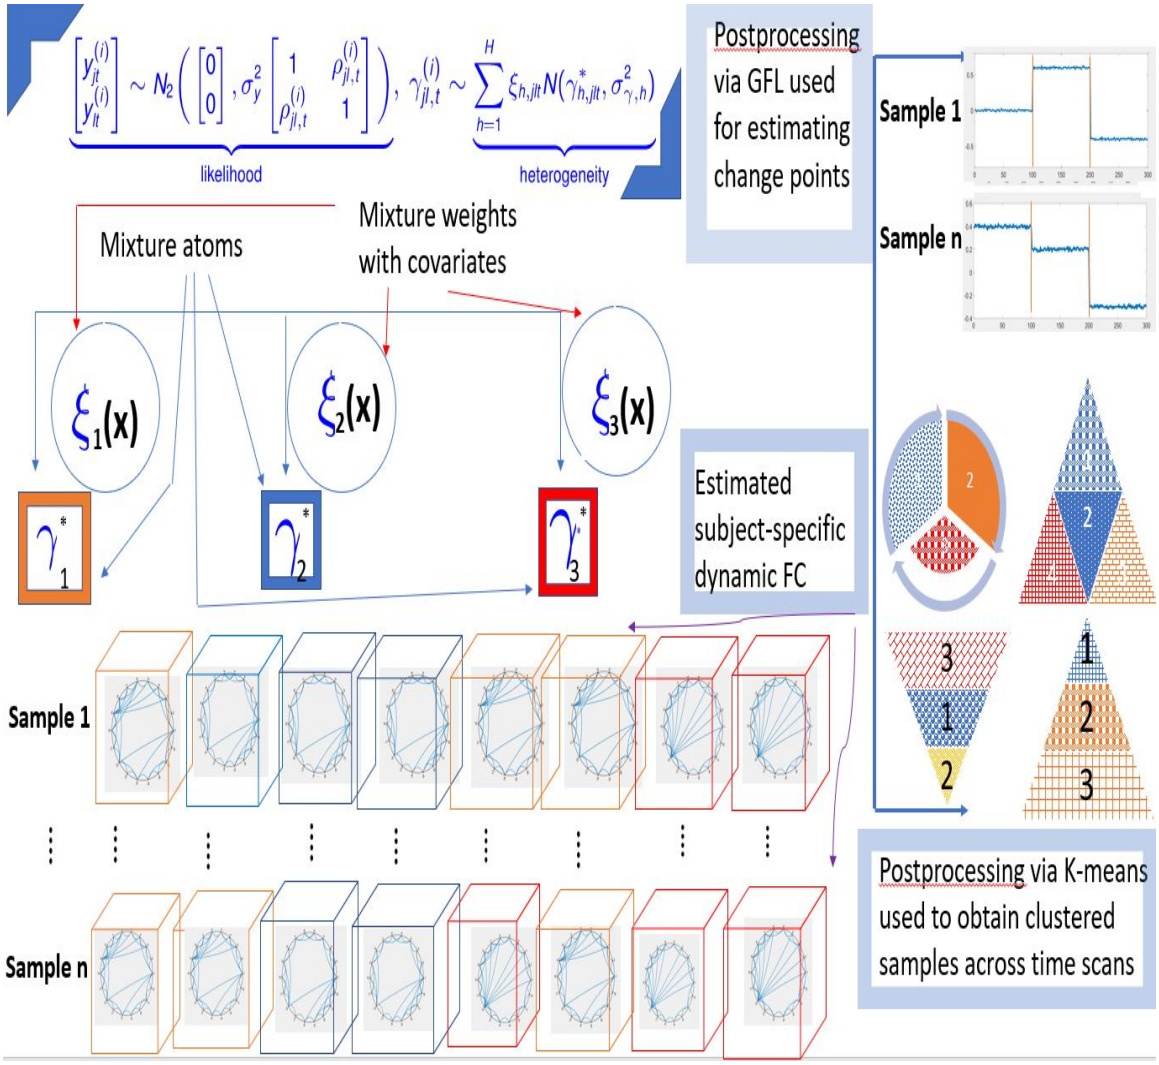

Figure 2.1: A schematic diagram illustrating the proposed dynamic pairwise correlation method. A mixture prior with $H = 3$ components is used to model dynamic correlations, where the mixture weights are modeled using covariates. The resulting networks at each time scan for each sample are allocated to one of the H clusters representing distinct network states that are represented by red, orange and blue cubes. Although the proposed method does not cluster transient states across time, the simplified representation in the Figure illustrates the similarity of brain states contained in identical colored cubes across the experimental session. Such temporal smoothness of the network is imposed via hierarchical fused lasso priors on the mixture atoms. Once, the dynamic FC is estimated, a post-processing step using K-means (Section 2.2) is applied to compute sub-groups of samples that exhibit similar dynamic connectivity patterns summarized across all time scans. The subgroups are represented by the circle, pyramid, triangle and inverted triangle shapes in the Figure and correspond to different modes of dynamic connectivity with different number of brain states represented by different patterns within each shape. The connectivity change points for each individual, as well as at a cluster level, are computed via another post-processing step that employs a group fused lasso penalty (Section 2.3). The method reports both individual and cluster-level network features.

matrix of fMRI measurements for the i-th individual as $Y^{(i)}$ that has the t-th column as $\mathbf{y}_t^{(i)}, i = 1, \ldots, N$. Further, denote the vector of $q \times 1$ covariates as $\mathbf{x}_i$ for the i-th sample, and represent the collection of fMRI data matrices across all individuals as Y.

In what follows, the idPAC method for pairwise correlations (Section 2.1) and idPMAC method for partial correlations (Section 3), will involve a combination of likelihood terms and priors on the model parameters that are combined into a posterior distribution, which is used to estimate model parameters. The posterior distribution for parameter θ given data Y is defined as $P(\theta|Y) = \frac{L(Y|\theta)\times\pi(\theta)}{P(Y)}$ using Bayes theorem, where $L(Y|\theta)$ denotes the data likelihood given the parameter value θ, $\pi(\theta)$ represents the prior on θ under the Bayesian model, and $P(Y) = \int L(Y|\theta)\pi(\theta)d\theta$ is the marginal likelihood after integrating out all possible values of θ. Full details of the posterior distributions are provided in the Appendix.

2.2.1 Dynamic Connectivity via Pair-wise Correlations

Let the unknown dynamic pairwise correlation of individual i be denoted as $\rho^{(i)} := \{\rho_{jl,t}^{(i)}, j < l, j, l = 1 \ldots V, t = 1 \ldots T\}$, and the corresponding Fisher-transformed pairwise correlations be denoted as $\gamma_{jl,t}^{(i)} = arctanh(\rho_{jl,t}^{(i)})$. We propose a Bayesian hierarchical approach that models the edge-wise dynamic correlations, using data from multiple individuals. We propose the following model for edge (j, l), and jointly for $t = 1, \ldots, T$, and $i = 1, \ldots, n$,

$$\begin{bmatrix} y_{jt}^{(i)} \\ y_{lt}^{(i)} \end{bmatrix} \sim N_2\left(\begin{bmatrix} 0 \\ 0 \end{bmatrix}, \sigma_y^2 \begin{bmatrix} 1 & \rho_{jl,t}^{(i)} \\ \rho_{jl,t}^{(i)} & 1 \end{bmatrix} \right), \gamma_{jl,t}^{(i)} \sim \sum_{h=1}^{H} \xi_{h,jlt}(\mathbf{x}_i) N(\gamma_{h,jlt}^*, \sigma_{\gamma,h}^2),\ \sigma_{\gamma,h}^{-2} \sim Ga(a_\sigma, b_\sigma),$$

$$\pi(\gamma_{h,jl1}^*, \ldots \gamma_{h,jlT}^*) \propto \exp(-\lambda \sum_{t=1}^{T-1} |\gamma_{h,jlt}^* - \gamma_{h,jl,t-1}^*|),\ \xi_{h,jlt}^{(i)}(\mathbf{x}_i) = \frac{exp(x_i{}_{h,jlt}^T)}{1 + \sum_{h=1}^{H-1} exp(x_i{}_{h,jlt}^T)}, \quad (2.1)$$

where $|\cdot|$ denotes the L_1 norm, σ_y^2 denotes the residual variance in the likelihood term, the Fisher-transformed correlations $\gamma_{jl,t}^{(i)}$ are modeled under a mixture of Gaussians prior having H components denoted as $\gamma_{h,jlt}^*, h = 1, \ldots, H$, with the prior probability for the h-th mixture component denoted as $\xi_{h,jlt}(\mathbf{x}_i)$ that depends on covariates, such that $\sum_{h=1}^{H} \xi_{h,jlt}(\mathbf{x}_i) = 1$ for all $t = 1, \ldots, T$, $\sigma_{\gamma,h}^2$ captures the (unknown) variability of the pairwise correlations under the mixture prior specification, and $N_v(\mu, \Sigma)$ denotes a multivariate Gaussian distribution with mean μ and $V \times V$ covariance matrix Σ. Under a hierarchical Bayesian specification, $\sigma_{\gamma,h}^{-2}$ is estimated under the conjugate Gamma prior with shape and scale parameters a_σ, b_σ, respectively. The mixture prior specifies that for any given time scan t, the functional connectivity for each individual can take values revolving around any one of the H mixture atoms denoted by $(\gamma_{1,jlt}^*, \ldots, \gamma_{H,jlt}^*)$, that are themselves unknown and modeled under a fused lasso prior as in (4.1). These values are realized with respective prior probabilities $(\xi_{1,jlt}(\mathbf{x}_i), \ldots, \xi_{H,jlt}(\mathbf{x}_i))$ that are modulated via covariates with effect sizes $(_{1,jlt}, \ldots, _{H,jlt})$ respectively, where $_{h,jlt} \sim N(0, \Sigma)$ with $_{H,jlt} = 0$ fixed as the reference group.

Modeling mixture atoms via fused lasso: The mixture atoms are modeled under a fused lasso prior in (4.1) that encourages temporal smoothness of pairwise correlations by assigning small prior probabilities for large changes in the values between consecutive time scans. Although temporal smoothness in correlations is encouraged, the Bayesian approach is still equipped to accommodate sharp jumps in connectivity that may arise due to changes in experimental design or other factors. Such connectivity jumps are detected using a post-processing step (see Section 2.3) applied to the estimated dynamic connectivity under the proposed model.

Modeling mixture weights via covariates: In order to effectively tackle heterogeneity, we incorporate supplementary covariate information when modeling the mixture weights under our mixture modeling framework in (4.1). By incorporating covariate

information, the model is designed to achieve more accurate identification of clusters, which then naturally translates to improved estimates for dynamic FC at the level of each individual. In particular, we model $(\xi^{(i)}_{1,jl}, \ldots, \xi^{(i)}_{H,jl})$ via a Multinomial Logistic regression (Engel, 1988), where $_{h,jlt} \sim N(0, \Sigma)$ (with $_{H,jlt} = 0$) represents the vector of unknown regression coefficients that control the contribution of the covariates to the mixture probabilities for the h-th component ($h = 1, \ldots, H-1$), in contrast to the H-th component. A large value of these regression coefficients implies increased importance of the corresponding covariate with respect to modeling a particular edge under consideration, whereas $_{1,jlt} \approx \ldots \approx_{H-1,jlt} \approx 0$ for all $t = 1, \ldots, T$, indicates spurious covariates unrelated to the dynamic pairwise correlations.

The multinomial logistic regression model incorporating covariates suggests that the log-odds for each component $\xi^{(i)}_{h,jlt}(\mathbf{x}_i)/\xi^{(i)}_{H,jlt}(\mathbf{x}_i), h = 1, H-1$, can be expressed as a linear combination of covariates. When two or more samples have similar covariate information, the prior specification in (4.1) will encourage similar mixture components to characterise the dynamic connectivity for all these samples that will result in analogous connectivity patterns. However the posterior distribution (that is used to derive parameter estimates) should be flexible enough to accurately estimate varying connectivity patterns between individuals even when they share similar covariate values, by leveraging information present in the data (as evident from extensive numerical studies in Section 5).

Role of clustering in tackling heterogeneity and pooling information: Under model (4.1), each sample will be assigned to one of the H clusters at each time scan in an unsupervised manner and guided by their covariate profiles in order to model the edge-level dynamic connectivity. Due to independent clustering at each time scan, these cluster configurations change over the experimental session in a data-adaptive manner to characterize connectivity fluctuations across individuals. Such time scan specific clusters represent subgroups of individuals with similar connectivity profiles

over a subset of time scans, which are learnt by pooling information across all samples within a cluster. Here, it is important to note that model (4.1) does not impose identical dynamic connectivity across all time scans between multiple individuals (that is biologically unrealistic), but instead encourages common connectivity patterns within subgroups of samples for a subset of time points that are learnt in a data-adaptive manner. Hence, the proposed method is designed to result in more accurate estimation compared to a single subject analysis that is not equipped to pool information across samples or a group level analysis that does not account for within sample heterogeneity. We note that although the estimation is performed separately for each edge, the connectivity estimates across all edges are consolidated to obtain connectivity change point estimates (Section 2.3) or identify subgroups with common dynamic connectivity profiles (Section 2.2).

2.2.2 Post-processing steps for sub-group detection

In practical neuroimaging applications, it is often of interest to detect dissimilar modes of dynamic connectivity patterns that are embodied by distinct subgroups of individuals who also differ in terms of demographic or clinical characteristics, or other factors. For example in our fMRI task study, one of the objectives is to assess variations in dynamic connectivity with respect to subgroups of samples that were assigned different interventions, and who also had varying demographic characteristics. Instead of comparing network differences between pre-specified subgroups that are likely to contain individuals with heterogeneous connectivity patterns, it is more appealing to develop a data-adaptive approach to identify subgroups that comprise individuals with homologous dynamic connectivity, and then examine connectivity variations across such subgroups and how these variations are related to intervention and other factors of interest. When estimating these subgroups, we do not require identical dynamic connectivity patterns for all individuals within subgroups, but rather expect

them to have limited network differences in terms of edge strengths and connectivity change points. An inherently appealing feature of subgroup detection is that is allows one to compute cluster level change points and other aggregate network features (see Section 2.3) which are more reproducible in the presence of noise and heterogeneity, compared to a single-subject analysis. Subgroup level network summaries may be particularly beneficial in certain scenarios such as fMRI block task experiments where it may be challenging for single-subject analyses to detect rapidly evolving network features induced via quick transitions between rest and task blocks within the experimental design.

We propose an approach that consolidates the time-varying clusters of samples under the BPMM approach to detect subgroups which comprises samples with similar network-level dynamic connectivity patterns. In order to identify these subgroups, we first create a $N \times N$ similarity matrix that measures the propensity of each pair of samples to belong to the same cluster over the experimental session. This matrix is created by examining the proportion of time scans during which a pair of samples belonged to the same cluster across the experimental session, averaged across all edges. Once this similarity matrix has been computed, a K-means algorithm is applied to identify clusters of samples that exhibit similar dynamic connectivity patterns across the experimental session. The number of clusters K is determined using some goodness of fit score such as the elbow method (Thorndike, 1953), or it is fixed as the maximum number of mixture components (H) under the BPMM approach. Finally, we note that the subgroup identification step is not strictly needed under the proposed BPMM framework for dynamic network estimation, but it is an optional analysis that can be used to identify cluster-level network features in certain scenarios of interest.

2.2.3 Post-processing steps for connectivity change point estimation

The estimated dynamic correlations can be used to detect connectivity change points in scenarios involving sharp changes in the network during the session, such as in fMRI task experiments. Our strategy involves computing change points for each individual network (a) at the edge level that captures localized changes; and (b) at the global level that captures major disruptions in connectivity over the entire network. We compute the change points using the total variation penalty (Bleakley and Vert, 2010) that was also used in CCPD approach by Kundu et. al (2018). However the proposed methods are distinct from the two-stage CCPD approach; the latter estimates connectivity change points based on empirical time-varying connectivity measures in the first stage, and then in the second stage, computes piecewise constant networks conditional on the estimated change points that represent connectivity jumps. In contrast, proposed method pools information across samples in order to first estimate dynamic correlations that does not depend on change points and can vary continuously over time, and subsequently uses a post-processing step to compute connectivity change points without requiring piecewise constant connectivity assumptions. An appealing feature of the proposed mixture modeling framework guided by covariates is that it is more suitable for tackling divergent dynamic connectivity across samples, in contrast to empirical correlations under the CCPD approach.

Denote the vector of estimated correlations over all edges for the i-th individual and at time scan t as $\widehat{\mathbf{r}}_t^{(i)} \in \Re^{V(V-1)/2}, t = 1, \ldots, T, i = 1, \ldots, N$. Then the functional connectivity change points for the i-th individual may be estimated using connections across all edges via a total variation norm penalty that is defined as $||\mathbf{u}_{t+1}^{(i)} - \mathbf{u}_t^{(i)}|| = \frac{1}{V(V-1)/2}\sqrt{\sum_{m=1}^{V(V-1)/2}(u_{t+1,m}^{(i)} - u_{t,m}^{(i)})^2}$. In particular, the following penalized criteria is

used as in Kundu et al. (2018) for detecting network level connectivity change points:

$$min_{\mathbf{u}\in\Re^{V(V-1)/2}} \sum_{t=1}^{T} ||\widehat{\mathbf{r}}_t^{(i)} - \mathbf{u}_t^{(i)}||^2 + \lambda_u \sum_{t=1}^{T-1} ||\mathbf{u}_{t+1}^{(i)} - \mathbf{u}_t^{(i)}||, \tag{2.2}$$

where λ_u represents the penalty parameter and $\mathbf{u}_t^{(i)} \in \Re^{p(p-1)/2}$ represents the piece-wise constant approximation to the time series of correlations at time point t for the i-th individual that also assumes the presence of an unknown number of connectivity jumps. The first term in (2.2) measures the error between the observed correlations and the piece-wise constant connectivity, while the second term controls the temporal smoothness of correlations for $V(V-1)/2$ edges. The increment $||\mathbf{u}_{t+1}^{(i)} - \mathbf{u}_t^{(i)}||$ in the second term becomes negligible when the multivariate time series does not change significantly between times t and $t+1$, but it takes large values corresponding to significant connectivity changes. The network change points computed via (2.2) represent global changes functional connectivity resulting from a subset of edges that exhibit large connectivity changes. It is important to note that not all edges are expected to exhibit changes at these estimated change points. When it is of interest to compute edge-level connectivity change points, one can simply use criteria (2.2) separately for each edge, so that the total variation term translates to the L_1 penalty. However, it is important to note that edge-level connectivity changes represent granular fluctuations that are typically more challenging to detect in the presence of noise in fMRI.

The number of change points is determined by the penalty parameter λ_u, with a smaller value yielding a greater number of change points and vice-versa. Tibshirani and Wang (2007) proposed an estimate of λ_u based on a pre-smoothed fit of a univariate time series using a lowess estimator (Becker et al., 1988). We adapt this approach for a multivariate time series to obtain an initial estimate for λ_u, and then propose post-processing steps to tune this estimate in order to obtain change points,

as in the CCPD approach in Kundu et al. (2018). Full details for these steps are provided in Section 3 of Supplementary Materials.

Cluster-level connectivity change point estimation: For fMRI task experiments involving multiple subjects, subgroups of individuals are expected to share analogous dynamic connectivity patterns with limited variations across samples, as discussed in Section 2.2. The proposed total variation penalty norm in (2.2) is equipped to leverage information across samples within a cluster for identifying cluster level change points, which reflect aggregated dynamic connectivity changes across all samples within a cluster at the global network level. These cluster level connectivity changes are obtained by aggregating the change points obtained via (2.2) applied separately to each sample within the cluster, and then choosing those change points that show up repeatedly within the cluster. One can define a threshold such that all change points that appear with a high frequency (above the chosen threshold) across samples within the cluster are determined to represent cluster level change points (Kundu et al., 2018). We note that under the proposed method, it is entirely possible for individuals within a cluster to have unique connectivity changes in addition to the common cluster level change points, which reflect within sample heterogeneity. In our experience, this method typically works well in accurately recovering aggregated cluster-level connectivity changes, in certain scenarios such as block task experiments, or more generally in the presence of subgroups of individuals with similar dynamic connectivity patterns.

2.3 Extension to Dynamic Precision Matrix Estimation

We now propose a mixture model for dynamic precision matrix estimation (idPMAC) that looks at the totality of all nodes in the network, in contrast to the edge-wise

analysis in Section 2.1. While the idPMAC also uses a mixture modeling framework, it is fundamentally distinct compared to the idPAC method in Section 2.1, with respect to the manner in which the mixture prior is specified and in terms of how the network edges are constructed and interpreted. The proposed approach estimates the network by computing the $V \times V$ precision matrix involving $V(V-1)/2$ distinct partial correlations that are learnt by borrowing information across V nodes at each time scan. The partial correlations measure interactions between pairs of regions after removing the influence of third party nodes, which is successful in filtering out spurious correlations. Hence a zero partial correlation between two nodes implies conditional independence. The proposed idPMAC approach enables one to report graph-theoretic network summary measures that capture important patterns of network information transmission (Lukemire et al., 2020), which may not be straightforward to report using pairwise correlations (Smith et al., 2012).

Denote the $V \times V$ precision matrix over all nodes for the i-th individual at the t-th time point as $\Omega_t^{(i)} = \begin{bmatrix} \omega_{t,11}^{(i)} & \omega_{1,t}^{(i)} \\ \omega_{1,t}^{(i)\prime} & \Omega_{11,t}^{(i)} \end{bmatrix}$, and note that the partial correlation between nodes k and l is given directly as $-\omega_{kl}/\sqrt{\omega_{kk}\omega_{ll}}$ (ignoring the subject-specific and time-scan specific notations). We propose a Gaussian graphical model involving product mixture priors as:

$$\mathbf{y}_t^{(i)} \sim N\Big[\mathbf{0}, \Omega_t^{(i)}\Big],\ \omega_{v,t}^{(i)} \sim \sum_{h=1}^{H} \xi_{h,t}(\mathbf{x}_i) N_{V-1}(\omega_{h,t}^*, \sigma_{\omega,h}^2 I_{V-1}),\ \xi_{h,t}^{(i)}(\mathbf{x}_i) = \frac{e^{x_i {}_{h,t}^{T}}}{1 + \sum_{h=1}^{H-1} e^{x_i {}_{h,t}^{T}}},$$
$$\omega_{t,vv}^{(i)} \sim E(\frac{\alpha}{2}),\ \pi(\omega_{h,1}^*, \ldots, \omega_{h,T}^*) \propto \exp(-\lambda \sum_{t=1}^{T-1} |\omega_{h,t}^* - \omega_{h,t-1}^*|),\ \sigma_{\omega,h}^{-2} \sim Ga(a_\sigma, b_\sigma),\ {}_{h,t} \sim N(0, \Sigma), \tag{2.3}$$

for $i = 1, \ldots, N, t = 1, \ldots, T$, where $\Omega_t^{(i)} \in M_V^+$, the space of symmetric positive definite matrices, $E(\alpha)$ denotes the Exponential distribution with scale parameter α, and $\omega_{v,t}^{(i)}$ denotes the vector of $(V-1)$ off-diagonal elements corresponding to the

v-th row of $\Omega_t^{(i)}$ that are modeled using a mixture of multivariate Gaussians prior. Specifically, the dynamic connectivity at time scan t is likely to be characterised via the hth mixture component with prior probability $\xi_{h,t}(\mathbf{x}_i)$ depending on covariates, where the prior mean and precision for this unknown mixture component is given by $\omega^*_{h,t}$ and $\sigma^{-2}_{\omega,h}$ respectively. The idPMAC approach in (2.3) specifies independent mixture priors on the set of all edges related to each node and at each time scan, while ensuring that the precision matrices are symmetric and positive definite. Full details for the computational steps are presented in Section 4.

Modeling mixture atoms: Under a hierarchical Bayesian specification, the mixture atoms or component-specific means $\omega^*_{h,t}$ are themselves unknown and modeled via a fused lasso prior, which encourages temporal homogeneity of partial correlations by assigning small prior probabilities for large changes in the values. In addition, systematic changes in connectivity reflected by sharp jumps may be still identified via a post-processing step in Section 2.3. The unknown prior variance on mixture components $(\sigma^{-2}_{\omega,h})$ is assigned a Gamma prior and is estimated adaptively via the posterior distribution.

Modeling mixture weights via covariates: The node level mixture weights incorporating covariates are modeled via a Multinomial Logistic regression with unknown covariate effects $(_{1,t}, \ldots, _{H,t})$ corresponding to time scan t that are assigned Gaussian priors, and we fix $_{H,t} = 0, t = 1 \ldots T$, as the reference group. The prior in (2.3) encourages similar clustering configurations resulting in analogous time-varying partial correlations for individuals with similar covariate profiles. However in the presence of heterogeneity, the posterior distribution under the idPMAC method is still able to identify divergent dynamic connectivity patterns even among individuals with same covariate profiles (see numerical studies in Section 5).

Role of clustering in tackling heterogeneity and pooling information: Under model (2.3), each column of the precision matrix is assigned to one of the H clusters at each

time scan in an unsupervised manner. Hence, the mixture modeling framework allows subsets of rows/columns of $\Omega_t^{(i)}$ to have the same values depending on their clustering allocation at each given time scan, which is an unique feature under the idPMAC approach that is not shared by the idPAC method. This feature results in robust estimates by pooling information across nodes and samples to estimate common partial correlations, and is a necessary dimension reduction step for scenarios involving large networks. For example, all weak or absent edges can be subsumed into one cluster which yields model parsimony. In addition, divergent connectivity patterns are captured via distinct time-varying clustering configurations across individuals as derived from the posterior distribution, which accommodates heterogeneity. Hence, the clustering mechanism under the idPMAC method not only enables dimension reduction, but also provides a desirable balance between leveraging common connectivity patterns within and across networks and addressing inherent network differences across individuals.

Post-processing steps: The post-processing steps for sub-group detection and connectivity change point estimation under the idPMAC approach can be applied in a similar manner as outlined in Sections 2.2 and 2.3. They proceed by replacing the estimated pairwise correlations with estimated dynamic partial correlations derived under the idPMAC approach in the K-means algorithm and the fused lasso criteria (2.2) in Sections 2.2 and 2.3 respectively.

2.4 Computational Details for Parameter Estimation

Although one can use Markov chain Monte Carlo (MCMC) to sample the parameters from the posterior distribution, we use a *maximum-a-posteriori* or MAP estimators for our purposes in this article that bypasses the computational burden under a MCMC

implementation. The MAP estimators are obtained by maximizing the posterior distribution for the model parameters and are derived via the Expectation-Maximization or EM algorithm. The EM algorithm is scalable to high-dimensional fMRI applications of interest that requires one to compute $N \times T$ distinct dynamic networks each involving $V \times V$ connectivity matrices. Table 2.1 provides a list of model parameters to be estimated via the EM steps for both the dynamic pairwise correlations and the dynamic precision matrix estimation methods.

Notation	Description	Update
	DATA	
$\mathbf{Y}$	fMRI scanning data for all individuals	observed
$y_{lt}^{(i)}$	observed fMRI data for individual i, node l, time point t	observed
$\mathbf{x_i}$	covariate information for individual i	observed
σ_y^2	Variance of $y_{lt}^{(i)}$ that is empirically estimated	Pre-Fixed
Σ_β	prior covariance for covariate effects β	Pre-Fixed
H	Number of components in mixture	Pre-Fixed
	Dynamic Pairwise Correlation	
$\rho_{jl,t}^i$	pairwise corr for edge (j,l) at time t ($\gamma_{jl,t}^{(i)}$ = Fisher-transformed $\rho_{jl,t}^i$)	M-step
$\gamma_{h,jlt}^*$	mean of $h-$th component for edge (j,l) at time point t	M-step
$\sigma_{\gamma,h}^2$	Variance for $h-$th mixture component	M-step
$\beta_{h,jlt}$	unknown regression coefficient used for modelling $\xi_{h,jlt}(\mathbf{x_i})$	M-step
$\psi_{h,jlt}^{(i)}$	posterior probability of $\gamma_{jl,t}^{(i)}$ taking values from $h-$th component	E-step
	Dynamic Precision Matrix	
$\omega_{h,t}^*$	h-th mixture component	M-step
$\sigma_{\omega,h}^2$	mixture variance of h-th component	M-step
$\beta_{h,t}$	unknown regression coefficient used in Multinomial Logistic regression	M-step
$\psi_{h,vt}^{(i)}$	posterior prob of $\omega_{\mathbf{v,t}}^{(\mathbf{i})}$ taking values from component h	E-step
$\Omega_t^{(i)}$	precision matrix for individual i at time point t	MC E-step

Table 2.1: Summary for all model parameters under the dynamic pairwise correlations and the dynamic precision matrix methods. MC E-step refers to Monte Carlo E-step.

2.4.1 EM Algorithm for Pair-wise dynamic connectivity

EM Algorithm: Denote the matrix containing the fMRI time series for the lth node as $Y_l = (\mathbf{y}_{1,l}, \ldots, \mathbf{y}_{T,l})$ where $\mathbf{y}_{t,l} = (y_{l,t}^{(1)}, \ldots, y_{l,t}^{(N)})'$ represents the fMRI observations across all samples for node l and time scan t. Further, denote Δ_h as a latent indicator

variable for the hth mixture component (that is not observed and is imputed in the proposed EM algorithm) and finally, denote by Θ^{jl} the collection of all model parameters under the specification (4.1) corresponding to edge (j, l). Note that under the proposed multinomial logistic regression model for incorporating covariates in (4.1), one has an equivalent specification under the binary latent variables distributed as $(\Delta^{(i)}_{1,jlt}, \ldots, \Delta^{(i)}_{H,jlt}) \sim MN\big(1, (\xi_{1,jlt}(\mathbf{x}_i; _{h,jlt}), \ldots, \xi_{H,jlt}(\mathbf{x}_i; _{h,jlt}))\big)$, where $MN(1; \mathbf{p}_0)$ denotes a multinomial distribution with probability vector $\mathbf{p}_0$, $B_{jlt} = (_{1,jlt}, \ldots, _{H-1,jlt})$ and one can marginalize out $(\Delta^{(i)}_{1,jlt}, \ldots, \Delta^{(i)}_{H,jlt})$ to recover the original specification in (4.1). The EM algorithm uses the augmented log-posterior derived in the Appendix involving the above latent mixture indicators, to compute MAP estimates for the model parameters by iteratively applying the Expectation (E) and Maximization (M) steps. The latent indicators $\{\Delta^{(i)}_{h,jlt}, h = 2, \ldots, H, t = 1, \ldots, T, i = 1, \ldots, N\}$ are imputed via the E-Step by using the posterior probability of $\gamma^{(i)}_{jl,t}$ taking values from the h-th mixture component, which is denoted by $\psi^{(i)}_{h,jlt} = Pr(\Delta^{(i)}_{h,jlt} = 1 \mid -)$ and updated as:

E-step: Compute the posterior expectation for the latent cluster membership indicators as $\hat{\psi}^{(i)}_{h,jlt} = \big\{ \sum_{r=1}^{H} \xi_{r,jlt}(\mathbf{x}_i; _{h,jlt}) \phi(\gamma^{(i)}_{jl,t} \mid \gamma^{*}_{r,jlt}, \sigma^2_{\gamma,h}) \big\}^{-1} \big\{ \xi_{h,jlt}(\mathbf{x}_i; _{h,jlt}) \times \phi(\gamma^{(i)}_{jl,t} \mid \gamma^{*}_{h,jlt}, \sigma^2_{\gamma,h}) \big\}$, where $\phi(\gamma^{(i)}_{jl,t} \mid \gamma^{*}, \sigma^2_{\gamma})$ denotes the normal density with mean γ^{*} and variance σ^2_{γ}.

The remaining parameters are updated via M-steps using closed form solutions except $\gamma^{(i)}_{jl,t}$ that is updated using Newton-Raphson steps. These M-steps comprise several mathematically involved derivations and are detailed in the Appendix. The E and M steps are repeated till convergence, which occurs when the absolute change in the log-posterior between successive iterations falls below a certain threshold (we use 10^{-4} in our implementation).

2.4.2 EM Algorithm for Dynamic Precision Matrix Estimation

Let us denote the collection of all the precision matrices as Θ, and $\mathbf{y}_{t,-v}^{(i)'}$ as the $(V-1)$-dimensional vector of fMRI measurements at time scan t over all nodes except node v. The prior on the precision matrix can be expressed as $\pi(\Omega_t^{(i)}) = \prod_{v=1}^{V} \pi(\omega_{t,vv}^{(i)})\pi(\omega_{vt}^{(i)})$, with the corresponding prior distributions $\pi(\cdot)$ being defined in (2.3). Denote by $|\cdot|_1$, the element-wise L_1 norm, denote $\kappa_{1,t}^{(i)} = \omega_{t,11}^{(i)} - \omega_{1,t}^{(i)'}\Omega_{11,t}^{(i)-1}\omega_{1,t}^{(i)}$ to represent the conditional variance corresponding to the fMRI measurements for the vth node given all other nodes, and let $\omega_{t,vv}^{(i)}$ and $\omega_{v,t}^{(i)}$ respectively denote the diagonal and the vector of off-diagonal elements of the vth row in $\Omega_t^{(i)}$. Moreover use $det(A)$ to denote the determinant of the matrix A, and write $S_t^{(i)} = \mathbf{y}_t^{(i)}\mathbf{y}_t^{(i)'} = \begin{bmatrix} s_{t,11}^{(i)} & \mathbf{s}_{1,t}^{(i)} \\ \mathbf{s}_{1,t}^{(i)'} & S_{11,t}^{(i)} \end{bmatrix}$ as the matrix of cross-products of the response variable, where $s_{vv,t}^{(i)}$ and $\mathbf{s}_{v,t}^{(i)}$ denote the v-th diagonal element and the off-diagonal elements for the v-th row respectively. Introduce latent indicator variables $(\Delta_{1,vt}^{(i)}, \ldots, \Delta_{H,vt}^{(i)})$ that follow a multinomial distribution with probability vector $(\xi_{1,t}(\mathbf{x}_i), \ldots, \xi_{H,t}(\mathbf{x}_i))$ such that $\sum_{h=1}^{H} \xi_{h,t}(\mathbf{x}_i) = 1$.

Denote by $\Omega_{vv,t}^{(i)}$, the $(V-1)\times(V-1)$ obtained by deleting the v-th row and column from $\Omega_t^{(i)}$. The EM algorithm uses an E step for the latent mixture indicators, as well as a Monte Carlo E step that samples from the posterior distribution in order to obtain estimates for the precision matrix. These steps are described below:

E-step for mixture component indicator: For $v = 1, \ldots, V$, use the expression: $\hat{\psi}_{h,vt}^{(i)} = \left\{ \sum_{r=1}^{H} \xi_{r,t}\big(\mathbf{x}_i;_{h,t}\big)\phi_{V-1}\big(\omega_{v,t}^{(i)} \mid \omega_{r,t}^{*}, \sigma_{\omega,r}^{2} I_{V-1}\big) \right\}^{-1} \times \left\{ \xi_{h,t}\big(\mathbf{x}_i;_{h,t}\big)\phi_{V-1}\big(\omega_{v,t}^{(i)} \mid \omega_{h,t}^{*}, \sigma_{\omega,h}^{2} I_{V-1}\big) \right\}$, where $\phi_{V-1}(\cdot \mid \omega^*, \Sigma)$ denotes the probability density function for the $(V-1)$-dimensional normal density with mean and variance as (ω^*, Σ) respectively.

Monte Carlo E-step for precision matrix: We use an E-step to update the precision matrix that computes the posterior mean by averaging MCMC samples drawn

from the posterior distribution, which is equivalent to a Monte Carlo EM method (Wei and Tanner, 1990). We use this Monte Carlo approximation for the conditional expectation since it provides a computationally efficient approach to sample symmetric positive definite precision matrices via closed form posteriors. The posterior distribution for the precision off-diagonal elements are given as $\pi(\hat{\omega}_{vt}^{(i)} \mid -) \sim N\left[V_{\omega_{vt}}\left(\sum_{h=1}^{H}\frac{\Delta_{h,vt}^{(i)}\omega_{h,t}^{*}}{\sigma_{\omega,h}^{2}}+2(\mathbf{s}_{v,t}^{(i)})\right), V_{\omega_{v,t}}\right]$, where $V_{\omega_{vt}} = \left(\sigma_{\omega,h}^{2}I_{V-1}+(s_{vv,t}^{(i)}+\alpha)(\Omega_{vv,t}^{(i)-1})+\sum_{h=1}^{H}\frac{\Delta_{h,vt}^{(i)}}{\sigma_{\omega,h}^{2}}\right)^{-1}$ is the posterior covariance. Moreover, writing $\omega_{t,vv}^{(i)} = \kappa_{v,t}^{(i)}+\omega_{v,t}^{(i)'}\Omega_{vv,t}^{(i)-1}\omega_{v,t}^{(i)}$, the diagonal precision matrix elements are updated via the posterior $\kappa_{vt}^{(i)} \sim GA(\frac{1}{2}+1, \frac{s_{vv,t}^{(i)}+\alpha}{2})$ where α is pre-specified. The above steps can be alternated to sample positive definite precision matrices as in Wang (2012), and we draw several MCMC samples and average over them to approximate the conditional expectation.

The remaining parameters are updated via closed form expressions under the M step, which involve mathematically involved derivations and are detailed in the Appendix. The algorithm iterates through the E and M steps until convergence.

2.4.3 Tuning Parameter Selection

Certain tuning parameters in the BPMM need to be selected properly or pre-specified, in order to ensure optimal performance. For both dynamic pair-wise correlations and precision matrix estimation, λ is the tuning parameter used in fused lasso penalty for the mixture atoms that controls the temporal smoothness of the dynamic connectivity. We choose an optimal value for λ over a pre-specified grid of values, as the value of the tuning parameter that minimizes the BIC score. In model (4.1) for the dynamic pairwise correlation, the σ_y^2 is also pre-specified as the initial mean variance over all edges and across all samples. Moreover when updating covariate effects, Σ_β is pre-fixed as a diagonal matrix with the diagonal terms as 100, although it is also possible

to impose a hierarchical prior on Σ_β and update it using the posterior distribution. Extensive simulation studies revealed that the proposed approach is not sensitive to the choices of Σ_β as long as the variances are not chosen to be exceedingly small. Other hyper-parameters in the hierarchical Bayesian specification include α in the prior on the precision matrices (chosen as in Wang (2012)), and $a_\sigma = 0.1, b_\sigma = 1$, that results in an uninformative prior on the mixture variance.

The number of mixture components H also needs to be chosen appropriately. On the one hand, a large value of H may be used to address inherent heterogeneity, but it will also increase the running time and may generate redundant clusters that overcompensates for the variations across samples. On the other hand, a small value of H may restrict the approach to overlook connectivity variations across individuals, resulting in inaccurate estimates. One may use a data adaptive approach to select H in certain scenarios where it is reasonable to assume that the dynamic connectivity can be approximated by piecewise constant connectivity. In such cases that potentially involve block task experiments (Kundu et al., 2018), one can evaluate criteria (2.2) separately for each individual under different values of H, and fix the optimal choice as that which minimizes the average value of the criteria (2.2) across all individuals. Based on extensive empirical studies, we noticed the need for larger values for H when fitting the model for cases involving large number of nodes and samples.

2.5 Numerical Experiments

2.5.1 Simulation set-up

Data generation: We generate observations from Gaussian distributions with sparse and piecewise constant precision matrices that change at a finite set of change points. Moreover, the network change points are generated based on covariate information where individuals with identical covariates have partially overlapping connectivity

change points. Broadly, we use the following few steps to generate the data, each of which is described in greater detail in the sequel: (i) generate a given number of change points for each subject using corresponding covariate information; (ii) conditional on the generated change points, piecewise constant networks are simulated such that the connectivity changes occur only at the given change points; (iii) conditional on the network for a given state phase, a corresponding positive definite precision matrix is generated for each time scan where non-zero off-diagonal elements represent edge strengths and zero off-diagonals represent absent edges; and (iv) the response variable for a given time point is generated from a Gaussian distribution having zero mean and the precision matrix in step (iii). Four clusters are created with 10 samples each, where the samples with each cluster have the same number of connectivity change points, common state phase specific networks and identical covariate values. However within each cluster, there are differences in locations of connectivity change points and the network edge strengths are free to vary across individuals even when they share the same network structure. All samples in the first two clusters have 3 connectivity change points each, whereas the samples in the other two clusters have 4 change points, out of a total of $T = 300$ time scans.

Conditional on the change points in step (i), several types of networks are constructed for each state phase in step (ii) that include: (a) Erdos Renyi network where each edge can randomly appear with a fixed probability; (b) small-world network, where the mean geodesic distance between nodes are relatively small compared with the number of nodes and which mimics several practical brain network configurations; and (c) scale-free network that resembles a hub network where the degree of network follows a power distribution. Given these networks, the corresponding precision matrix was generated in step (iii) by assigning zeros to off-diagonals for absent edges, and randomly generating edge weights from uniform [-1,1] for all important edges. To ensure the positive definiteness, the diagonal values of the precision matrix

were rescaled by adding the sum of the absolute values of all elements in each row with one. Finally, the response variables were generated either (a) independently at each time point via a Gaussian graphical model, or (b) via a vector autoregressive (VAR) model where the response variables are autocorrelated across time. In both cases, sparse time-varying precision matrices having dimensions $V = 40, 100$, were used. The 'VARM' function in Matlab was used to generate temporally correlated observations under a lag-1 VAR model, where the elements in autocorrelation matrix were generated from a uniform random variable with range $(-0.2, 0.2)$.

We generated two binary features that resulted in four distinct covariate configurations, i.e. (0,0), (0,1), (1,0), (1,1), and all samples with identical covariates were allocated to the same cluster. In addition, we also evaluated the performance of proposed method in the presence of spurious covariates that are not related to dynamic connectivity patterns. Specifically, we introduced anywhere between 1 to 8 spurious covariates for each sample (in addition to the two true covariates described earlier), which were randomly generated using uniform as well as from random normal distributions. We then investigated the performance of the proposed approach over varying number of spurious covariates. While the proposed approach is expected to work best in practical experiments involving a carefully selected set of covariates that influence dynamic connectivity patterns, our goal was also to investigate the change in performance as the number of spurious covariates increase.

Competing methods: We perform extensive simulation studies to evaluate the performance of the proposed approach, and compare the performance with (a) change point estimation approaches such as the CCPD (Kundu et al., 2018) that can estimate single subject connectivity using multi-subject data in the presence of limited heterogeneity, and the dynamic connectivity regression (DCR) approach for single subjects proposed in Cribben et al. (2013); (b) an empirical sliding window based approach (SD) and the model-based SINGLE (Monti et al., 2014) method that uses sliding

window correlations; and (c) a covariate-naive version of the proposed approach using the methods in Section 2.1 and Section 3 (denoted as BPMM-PC and BPMM-PR respectively) that employs a multinomial distribution to model the mixture weights without covariates. While methods in (a) and (c) are designed to report connectivity change points, we were also able to compute change points under the sliding window approaches in (b) by applying a post-processing step in (2.2) on the estimated sliding window correlations. Moreover, the data under the VAR case were prewhitened via an autoregressive integrated moving average (ARIMA) before fitting the proposed models. In particular, the '*auto.arima*' in *R* was used to prewhiten the raw data, which yielded residuals that were subsequently used for analysis (more details provided in Supplementary Materials). The SINGLE approach was implemented using the python implementation in *pySINGLE*, while all other methods were implemented in Matlab.

Performance metrics: We evaluate the performance of different approaches in terms of different metrics. First, we investigated the accuracy in recovering true connectivity change points at the network and edge level for each sample, using sensitivity (defined as the proportion of truly detected change points or true positives), as well as the number of falsely detected change points or false positives. In addition, the performance of the network connectivity change points at the cluster level was also evaluated by comparing the true connectivity change points for each sample within the cluster with the aggregated cluster level change points. We note that since there were variations in connectivity change points within each cluster, false positive change points are to be expected under any estimation approach; however our goal is to evaluate how well these false positives are controlled and the sensitivity in detecting true change points under different methods. In addition, we also evaluated accuracy in terms of estimating the strength of connections that is computed as a squared loss (MSE) between the estimated and the true edge-level pairwise correlations. The

pairwise correlations corresponding to dynamic precision matrix approaches for computing MSE were obtained by inverting the respective precision matrices.

In order to evaluate the accuracy in dynamic network estimation, we computed the F-1 score defined as 2(Precision×Recall)/(Precision + Recall), where Precision=$TP/(TP+FP)$ is defined as the true positive rate, and Recall=$TP/(TP+FN)$ represents the sensitivity in estimating the edges in the network. Here, TP, FP, FN, refer to the number of true positive, false positive, and false negative edges that are obtained via binary adjacency matrices derived by thresholding the estimated absolute partial correlations. We employed reasonable thresholds (0.05) that are commonly used in literature (Kundu et al., 2018). In contrast, it was not immediately clear how to choose such thresholds for pairwise correlations given the fact that they tend to be typically larger in magnitude and have greater variability. Hence, we did not report F-1 scores corresponding to pairwise correlations, although one could do so in principle by choosing suitable thresholds to obtain binary adjacency matrices. Finally, we also evaluated the clustering performance in terms of the clustering error (CE) and Variation of Information (VI). CE (Patrikainen and Meila, 2006) is defined as the maximum overlap between the estimated clustering with the true clustering, whereas VI (Meil$\check{a}$, 2007) calculates the entropy associated with different clustering configurations.

2.5.2 Results

The performance in terms of recovering the true clusters of subjects is provided in Table 2.2, in the presence of two covariates that are both related to the true connectivity changes. It is clear from the results that incorporating covariate information results in near perfect recovery of the clusters, in contrast to the covariate-naive version of the method. For $V = 100$, the dynamic pairwise correlation approach seems to have a slightly higher accuracy in terms of cluster recovery compared to the dynamic pre-

cision matrix approach when data were generated from a VAR model. However when covariates are not included, the BPMM-PR method has greater clustering accuracy compared to the BPMM-PC approach, since the former is able to pool information across the whole network to inform the clustering mechanism, in contrast to an edge-by-edge analysis under BPMM-PC. Table 2.3 reports the accuracy in recovering the true network-level change points under the proposed approaches at the level of the estimated clusters, as per discussions in Section 2.3. In this case, both idPAC and idPMAC methods are shown to have near perfect recovery of the true network connectivity change points when data were generated under GGM, and high sensitivity when data were generated under VAR. Moreover when using data from a VAR model, the idPAC method has a comparable or higher sensitivity but also higher false positives for $V = 100$ in terms of detecting connectivity change points at the cluster level, compared to the idPMAC method. We note that although all samples within a cluster had identical covariate information, the proposed approach was able to accommodate within cluster connectivity differences that is evident from low false positives and high sensitivity when estimating cluster level change points. Moreover as seen from Tables 2.4-2.5, the accuracy in recovering cluster level connectivity change points is considerably higher than the corresponding results at the level of individual networks. These results indicate the usefulness of aggregating information when it is reasonable to assume the existence of subgroups of individuals who share some similar facets of dynamic connectivity.

Table 2.4 reports the performance under pair-wise correlation based approaches, i.e. idPAC, BPMM-PC, SD, and CCPD. It is clear for the results that the proposed idPAC method has a near perfect sensitivity when data were generated under GGM, and a suitably high sensitivity under the VAR model, when estimating connectivity change points. The sensitivity for network and edge change point estimation, along with the MSE in estimating the pairwise correlations are significantly improved under

	idPAC				BPMM-PC			
	V=40		V=100		V=40		V=100	
	CE	VI	CE	VI	CE	VI	CE	VI
GGM+Erdos-Renyi	0	0	0	0	0.64	1.93	0.62	2.19
GGM+Small-world	0	0	0	0	0.57	1.92	0.71	2.23
GGM+Scale-free	0	0	0	0	0.63	2.01	0.66	2.19
VAR+Erdos-Renyi	0	0	0	0	0.61	1.93	0.67	1.97
VAR+Small-World	0	0	0	0	0.59	1.88	0.61	1.90
VAR+Scale-Free	0	0	0	0	0.61	1.78	0.61	1.93
	idPMAC				BPMM-PR			
	V=40		V=100		V=40		V=100	
GGM+Erdos-Renyi	0	0	0	0	0.43	1.41	0.54	1.59
GGM+Small-world	0	0	0	0	0.41	1.41	0.51	1.68
GGM+Scale-free	0	0	0	0	0.43	1.49	0.60	1.78
VAR+Erdos-Renyi	0.08	0.25	0.04	0.17	0.54	1.51	0.66	1.88
VAR+Small-World	0	0	0.03	0.14	0.48	1.47	0.58	1.91
VAR+Scale-Free	0	0	0.04	0.11	0.49	1.42	0.63	1.75

Table 2.2: Clustering performance under different network types. GGM implies that the Gaussian graphical model was used to generate temporally uncorrelated observations; VAR implies a vector autoregressive model that was used to generate temporally dependent observations. For the VAR case, the observations were pre-whitened before fitting the model.

	idPAC				idPMAC			
	V=40		V=100		V=40		V=100	
	sens	FP	sens	FP	sens	FP	sens	FP
GGM+Erdos-Renyi	1	2.15	0.99	1.58	0.97	3.94	0.99	3.18
GGM+Small-world	0.97	2.11	1	1.59	0.99	4.18	0.98	3.17
GGM+Scale-free	0.99	2.09	1	1.37	1	3.91	0.97	3.09
VAR+Erdos-Renyi	0.91	3.71	0.88	3.66	0.87	3.47	0.87	2.89
VAR+Small-world	0.84	3.44	0.8	3.09	0.82	3.45	0.81	2.98
VAR+Scale-free	0.88	3.29	0.84	3.68	0.85	3.3	0.81	3.01

Table 2.3: Cluster-based network change point estimation under the proposed approaches, assuming that all samples within a particular cluster have the same number and similar location of change points, with a limited degree of heterogeneity in functional connectivity. If this assumption holds, then the cluster level network change point estimation provides greater accuracy compared to the estimated change points at the level of individuals as reported in subsequent Tables.

Results for V=40	Network CP		Edge CP		MSE	Network CP		Edge CP		MSE
	sens	FP	sens	FP	MSE	sens	FP	sens	FP	MSE
	BPMM-PC					idPAC				
GGM+Erdos-Renyi	0.91	7.31	0.50	1.12	**0.1**	**1**	2.75	**0.92**	1.08	**0.09**
GGM+Small-world	0.92	5.99	0.47	1.03	0.12	**0.98**	2.77	**0.92**	1.01	**0.08**
GGM+Scale-free	0.91	7.29	0.49	1.19	0.12	**1**	2.81	**0.92**	1.1	**0.09**
	SD+GFL					CCPD				
GGM+Erdos-Renyi	0.3	3.13	0.09	2.97	0.29	0.92	2.15	0.31	4.1	0.16
GGM+Small-world	0.29	3.31	0.09	3.08	0.27	0.92	2.18	0.29	4.17	0.21
GGM+Scale-free	0.29	3.08	0.09	2.99	0.24	0.91	2.33	0.29	4.09	0.19
	BPMM-PC					idPAC				
VAR+Erdos-Renyi	0.68	6.55	0.43	1.08	0.2	**0.84**	5.57	**0.80**	1.06	**0.12**
VAR+Small-world	0.66	5.97	0.47	1.14	0.19	**0.77**	5.54	**0.74**	1.12	**0.09**
VAR+Scale-free	0.59	5.51	0.39	1.02	0.17	**0.78**	5.29	**0.73**	1.06	**0.09**
	SD+GFL					CCPD				
VAR+Erdos-Renyi	0.41	7.72	0.13	3.06	0.26	0.55	**1.12**	0.18	4.33	0.21
VAR+Small-world	0.56	6.29	0.14	2.98	0.19	0.64	**1.36**	0.17	3.47	0.23
VAR+Scale-free	0.42	6.99	0.17	3.13	0.22	0.58	**1.27**	0.19	3.29	0.2
Results for V=100	Network CP		Edge CP		MSE	Network CP		Edge CP		MSE
	sens	FP	sens	FP	MSE	sens	FP	sens	FP	MSE
	BPMM-PC					idPAC				
GGM+Erdos-Renyi	0.92	4.77	0.51	1.31	0.11	**1**	2.31	**0.83**	**1.16**	0.09
GGM+Small-world	0.91	4.69	0.49	1.33	0.1	**1**	2.37	**0.82**	**1.17**	0.09
GGM+Scale-free	0.91	4.71	0.50	1.31	0.11	**1**	2.29	**0.83**	**1.16**	0.09
	SD+GFL					CCPD				
GGM+Erdos-Renyi	0.3	3.13	0.09	2.97	0.29	0.9	**1.12**	0.29	4.6	0.18
GGM+Small-world	0.29	3.31	0.09	3.08	0.27	0.91	**1.18**	0.25	4.2	0.17
GGM+Scale-free	0.29	3.08	0.09	2.99	0.27	0.91	**1.02**	0.27	4.4	0.17
	BPMM-PC					idPAC				
VAR+Erods-Renyi	0.66	5.97	0.51	1.07	0.14	**0.82**	5.88	**0.81**	1.04	**0.11**
VAR+Small-world	0.59	6.03	0.41	1.02	0.14	**0.75**	5.44	**0.74**	1.05	**0.12**
VAR+Scale-free	0.62	5.49	0.44	0.99	0.15	**0.77**	5.51	**0.71**	1.11	**0.13**
	SD+GFL					CCPD				
VAR+Erdos-Renyi	0.37	8.03	0.1	3.14	0.15	0.55	**1.09**	0.17	3.75	0.22
VAR+Small-world	0.44	7.51	0.16	2.71	0.16	0.66	**1.44**	0.19	3.41	0.19
VAR+Scale-free	0.36	7.72	0.18	2.88	0.18	0.59	**1.31**	0.17	3.44	0.19

Table 2.4: Results under the dynamic pair-wise correlation approaches for network and edge-level connectivity change-point estimation (Edge CP) accuracy and network changepoint (Network CP) estimation accuracy for $V = 40, 100$. GGM and VAR correspond to data generated from Gaussian graphical models and vector autoregressive models. Significantly improved metrics among the four approaches corresponding to the GGM data and separately for the VAR data, are highlighted in bold font. The standard deviations corresponding to the reported metrics are presented in separate Tables in the Supplementary Materials.

Results for V=40	Network CP		Edge CP		MSE	F1	Network CP		Edge CP		MSE	F1
	sens	FP	sens	FP	MSE	F1	sens	FP	sens	FP	MSE	F1
	BPMM-PR						idPMAC					
GGM+Erdos-Renyi	0.85	6.99	0.32	1.04	0.1	0.79	**1**	**5.2**	**0.79**	**0.89**	0.08	**0.88**
GGM+Small-world	0.88	7.14	0.33	1.16	0.08	0.77	**1**	**5.11**	**0.81**	**0.91**	0.08	**0.9**
GGM+Scale-free	0.87	7.36	0.33	1.19	0.08	0.71	**0.97**	**5.6**	**0.77**	**0.92**	0.07	**0.89**
	DCR						SINGLE					
GGM+Erdos-Renyi	0.22	16.15	0.41	9.39	0.27	0.59	0.35	6.49	0.1	2.84	0.08	0.71
GGM+Small-world	0.19	11.83	0.49	9.66	0.22	0.61	0.32	6.55	0.09	2.88	0.07	0.77
GGM+Scale-free	0.21	10.92	0.49	9.058	0.23	0.62	0.33	6.01	0.09	2.94	0.07	0.69
	BPMM-PR						idPMAC					
VAR+Erdos-Renyi	0.66	**4.45**	0.29	1.16	0.10	0.77	**0.79**	4.81	0.68	1.22	0.09	0.81
VAR+Small-world	0.59	5.12	0.27	**1.03**	0.1	0.74	**0.78**	4.99	**0.69**	1.04	0.09	**0.79**
VAR+Scale-free	0.61	4.77	0.31	1.04	0.12	0.77	**0.76**	4.64	**0.71**	0.99	**0.09**	**0.82**
	DCR						SINGLE					
VAR+Erdos-Renyi	0.22	9.83	0.4	3.35	0.24	0.64	0.42	7.35	0.13	3.11	0.27	0.66
VAR+Small-world	0.24	10.14	0.33	3.61	0.23	0.63	0.44	7.12	0.17	3.04	0.26	0.62
VAR+Scale-free	0.21	9.98	0.32	3.61	0.22	0.59	0.38	6.77	0.21	3.36	0.23	0.6

Results for V=100	Network CP		Edge CP		MSE	F1	Network CP		Edge CP		MSE	F1
	sens	FP	sens	FP	MSE	F1	sens	FP	sens	FP	MSE	F1
	BPMM-PR						idPMAC					
GGM+Erdos-Renyi	0.92	6.83	0.28	1.09	0.08	0.83	**0.97**	**5.1**	**0.82**	**0.89**	0.08	**0.89**
GGM+Small-world	0.91	6.98	0.31	1.19	0.09	0.81	**0.97**	**5.44**	**0.81**	**0.99**	0.07	**0.87**
GGM+Scale-free	0.92	7.44	0.32	1.25	0.08	0.81	**0.96**	**5.6**	**0.79**	**0.94**	0.07	**0.87**
	DCR						SINGLE					
GGM+Erdos-Renyi	0.33	16.14	0.41	9.39	0.22	0.63	0.38	6.77	0.12	2.97	0.08	0.69
GGM+Small-world	0.31	15.88	0.4	9.66	0.27	0.59	0.35	6.48	0.12	3.02	0.08	0.71
GGM+Scale-free	0.34	16.82	0.39	10.08	0.27	0.64	0.35	7.02	0.11	2.97	0.08	0.70
	BPMM-PR						idPMAC					
VAR+Erdos-Renyi	0.73	4.41	0.29	1.18	0.14	0.77	**0.88**	4.22	**0.63**	**1.09**	0.13	**0.82**
VAR+Small-world	0.56	5.22	0.22	**0.91**	0.11	0.78	**0.72**	4.87	**0.61**	1.09	0.1	**0.81**
VAR+Scale-free	0.59	5.13	0.29	1.03	0.11	0.78	**0.77**	4.49	**0.65**	1.08	0.09	**0.81**
	DCR						SINGLE					
VAR+Erdos-Renyi	0.23	9.92	0.43	3.19	0.16	0.64	0.42	7.41	0.14	3.11	0.11	0.71
VAR+Small-world	0.31	10.23	0.37	3.37	0.19	0.67	0.47	7.66	0.13	3.28	0.12	0.69
VAR+Scale-free	0.25	10.23	0.38	3.61	0.18	0.65	0.44	7.59	0.13	3.19	0.11	0.66

Table 2.5: Results under the dynamic precision matrix estimation approaches for network and edge-level connectivity change-point estimation (Edge CP) accuracy and network changepoint (Network CP) estimation accuracy for $V = 40, 100$. GGM and VAR correspond to data generated from Gaussian graphical models and vector autoregressive models respectively. Significantly improved metrics among the four approaches corresponding to the GGM data and separately for the VAR data, are highlighted in bold font. The standard deviations corresponding to the reported metrics are presented in the Supplementary Materials.

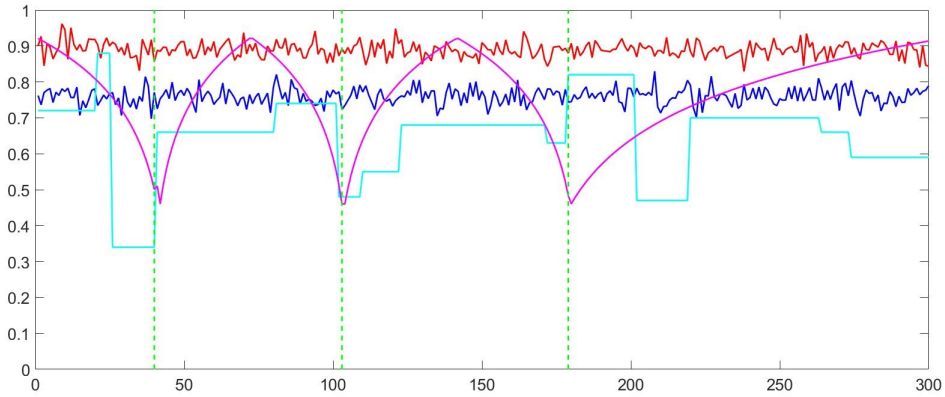

Figure 2.2: F1-score over time for one single subject under the case of dynamic partial correlation method. The vertical green lines are the true change points. Red line represents the proposed method with dynamic partial correlation (idPMAC), the cyan line represents the covariate-naive version (BPMM-PM), the blue line represents DCR, and the pink line represents SINGLE method.

idPAC compared to competing approaches in Table 2.4. The CCPD method is shown to have the lowest false positives when estimating the network level change points, but otherwise has poor sensitivity for change point estimation and high MSE, which is potentially due to the assumption of piecewise constant connectivity. The approach based on sliding window correlations has the poorest performance across all the reported metrics, which illustrates their drawback in estimating dynamic connectivity.

Table 2.5 reports the performance under precision matrix based approaches, i.e. idPMAC, BPMM-PR, SINGLE, and DCR. It is evident that the proposed idPMAC method has near-perfect or high sensitivity for detecting network level change points, corresponding to data generated under GGM and VAR models respectively. It also has a suitably high sensitivity for detecting edge level connectivity change points under both cases. Similarly, the MSE for edge strength estimation and the F-1 scores for network estimation accuracy are significantly improved under the proposed method in contrast to competing approaches. Figure 2.2 illustrates that the F-1 score over time under the proposed dynamic precision matrix method with covariates is almost always higher across almost all time scans compared to competing methods.

Moreover the DCR and SINGLE method have the least impressive performance in terms of connectivity change point estimation, which also translates to poor dynamic network estimation (low F-1 scores).

Our results clearly illustrate the advantages of the proposed methods over existing approaches that are not effective in leveraging information across samples. In addition, Tables 2.4-2.5 also illustrate the gains of incorporating covariate information under the proposed idPAC and idPMAC approaches over the covariate naive BPMM counterparts. It is interesting to note that the covariate naive BPMM still fares better than existing dynamic connectivity methods that fail to pool information across samples in a systematic manner. We also note that while the presence of false positive (FP) connectivity change points are expected due to the heterogeneity across samples, the proposed approaches provide desirable control of FP even while pooling information across samples with varying networks. In fact, the FP under the proposed method are lower than all competing methods except CCPD, whose performance is otherwise less impressive in terms of significantly lower sensitivity for change point detection, and inferior network estimation as reflected by poor MSE and F-1 scores.

When comparing the relative performance between idPAC and idPMAC methods, it is evident that the former has comparable or higher sensitivity but lower false positives in terms of estimating connectivity change points at the network level, when data are generated under a GGM. When data are generated under a VAR model, the idPAC method has higher sensitivity but also higher false positives compared to idPMAC, for estimating network connectivity change points. This is also true when estimating edge-level connectivity change points. In addition, since the idPMAC method estimates all edges simultaneously, the mean squared error for estimating edge strengths is often lower compared to the idPAC method. Moreover when the number of spurious covariates is increased, both these approaches experience a drop in performance (Figure 2.3), as expected. It is of note that the number of false positive

change points under the dynamic pairwise correlation approach increase minimally under the scale-free and small-world networks with an increase in the number of spurious covariates. However, this robust behavior was not replicated for network change points or other metrics of interest under the dynamic pairwise correlation method. In contrast, the recovery of the true clusters is shown to be more resilient under the dynamic precision matrix approach. This is evident from the top panels in Figure 2.3 that show a slower increase in the clustering error under the idPMAC method.

Method	v=20	V=40	V=100
BPMM-PC	21	80	321
BPNN-PR	25	92	348
idPAC	27	102	402
idPMAC	31	114	416
SD+GFL	3	9	44
CCPD	70	315	844
DCR	18	90	297

Table 2.6: Computation Time (in minutes) for simulation studies involving 300 time scans and 40 samples, under all approaches implemented via Matlab version R2017a.

The proposed approach is clearly scalable to higher dimensional networks, with the dynamic pairwise correlation method being slightly faster than the dynamic precision matrix estimation approach. We found the computation time under the proposed approaches to be slightly slower than existing dynamic connectivity methods such as DCR, which results from additional computations related to clustering and due to incorporation of covariates. Table 2.6 presents the computation time in minutes for all approaches implemented via Matlab on a personal desktop computer (Alienware Desktop) that had Intel(R) Core i7-4930 processor with 32GB RAM (SINGLE was implemented via Python and hence the computation time is not reported). We note that the total computation time under BPMM is expected to increase with V, T, N, which is true for most dynamic connectivity approaches.

Finally, we also conducted sensitivity analysis of the proposed approaches with

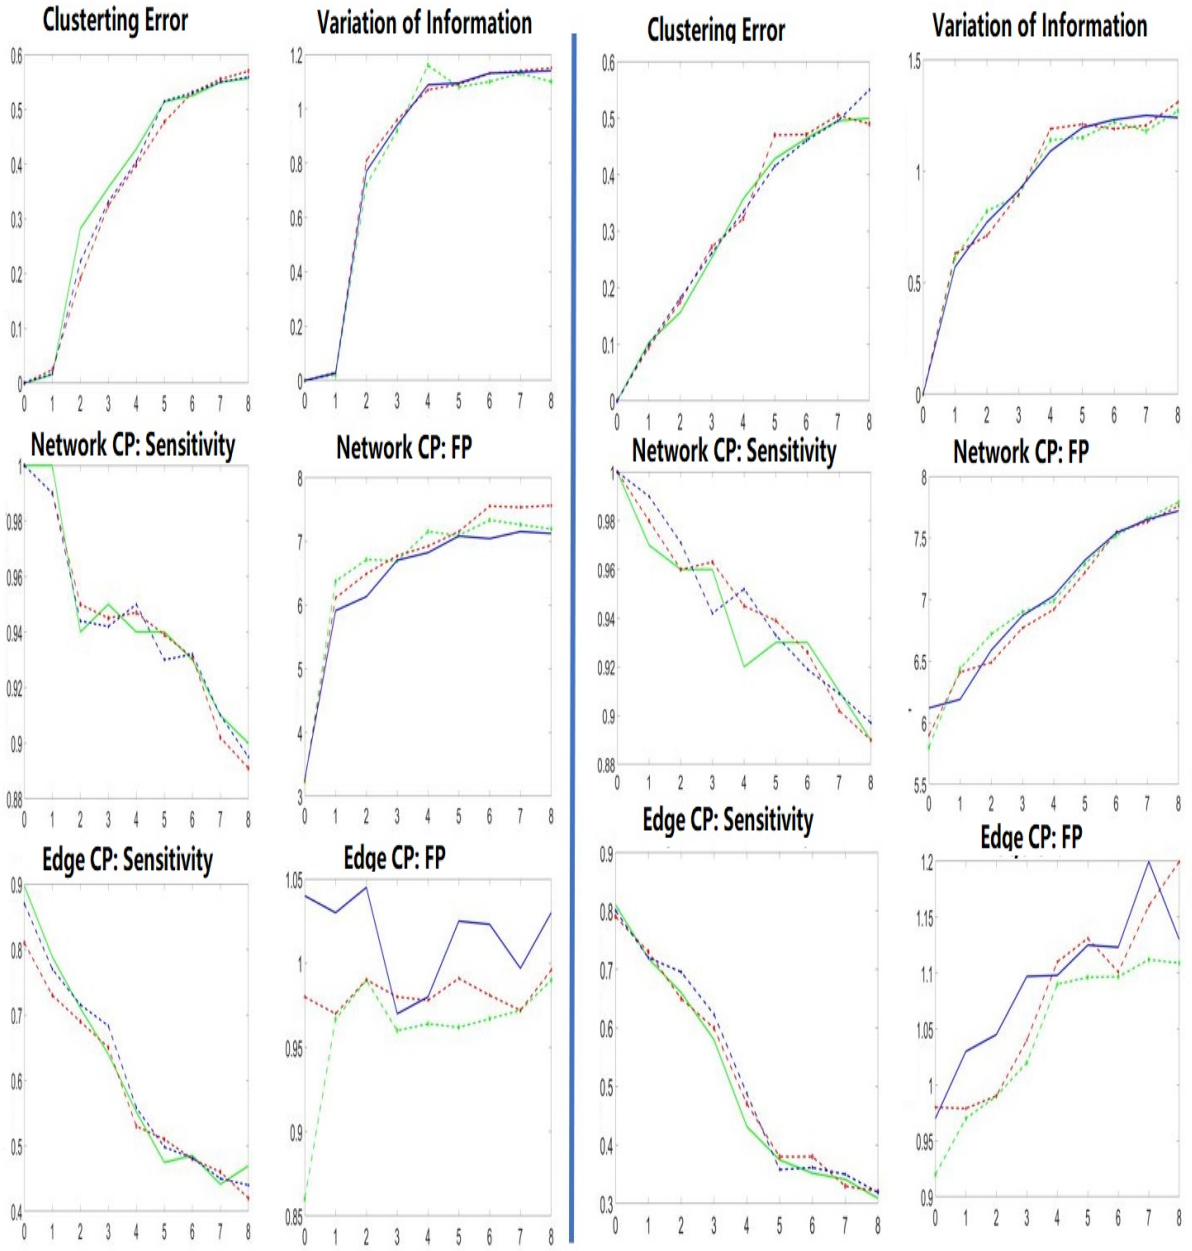

Figure 2.3: Performance of dynamic pairwise correlation (columns 1 and 2) and dynamic precision matrix (columns 3 and 4) methods under different number of spurious covariates represented by the X-axis. Lines with different color represent different network structure: Green (Erdos Renyi), Red (Small World), Blue (Scale Free). The top row provides the information of clustering performance (Clustering Error and Variation of Information), the middle row demonstrates the performance of network level change points estimation (sensitivity and number of False Positive estimations), and the performance of edge level change point estimation was provided in the bottle row.

respect to the model hyperparameters (see Tables 3-4 in Supplementary Materials). We found that moderate variations in the values of hyperparameters do not result in considerable changes in performance. Further, the clustering error is seen to be more sensitive to changes in hyperparameter values, while the metrics corresponding to dynamic network estimation seem more resilient to changes in hyperparameters, which suggests a degree of robustness for the estimated dynamic network with respect to the choice of hyperparameters. Although we do expect the performance of the methods to fluctuate to a greater degree for extreme choices of hyper-parameters, this is not of immediate concern to practitioners who use recommended values for the model hyper-parameters suggested in the manuscript.

2.6 Analysis of Task fMRI Data

2.6.1 Description of the study

We analyze a block task data involving a semantic verbal fluency at Veterans Affairs Center for Visual and Neurocognitive Rehabilitation, Atlanta. In a 12-week randomized controlled trail, 33 elderly individuals (aged 60-80, 11 males, 22 females) were assigned to two intervention groups: spin aerobic exercise group (14 participants) and the non-aerobic exercise control group (19 participants). During the intervention, individuals belonging to the aerobic spin group were required to do 20-45 minutes of spin aerobic exercise three times a week, led by a qualified instructor. For control group, participants were asked to do the same amount of non-aerobic exercise per week, such as group balance and light muscle toning exercise. A more detailed description of the data is available in Nocera et al. (2017).

For each participant, fMRI scans were conducted with 6 blocks of semantic verbal fluency (task) conditions with 8 scans, both pre- and post-intervention. The semantic verbal fluency task involved participants looking at different categories (e.g. “colors”)

at the center of video screen and they were asked to generate and speak 8 different objects associated with that category (e.g. "blue"). After task block, a rest block with 3-5 TRs would appear and participants were required to read the word "rest" out loud. A total of 74 brain scans were acquired using a 3T Siemens Trio scanner with a whole-brain, 1-shot gradient EPI scan (240mm FOV, 3.75 $\times$ 3.75 in-plane resolution, TR=5830ms, TA=1830ms, TE=25ms, flip angle (FA)=70). Analysis of Functional NeuroImages (AFNI) software and FMRIB Software Library (FSL) were used for pre-processing, as in Nocera et al. (2017). Slice-time corrections, linear trend removal, echo planar images alignment, and motion correction were performed as a part of the pre-processing pipeline. We used 18 brain regions for analysis that were shown to be differentially activated between the two intervention groups as described in Nocera et al. (2017). These regions are listed in Table 2.7 and comprise more regions in the right hemisphere due to decreased activity in that hemisphere in the spin group following the intervention, as compared to the control group. We note that since these regions corresponded to group differences due to spin exercise, they can not be described as "canonical" regions associated with semantic language function, which would also comprise some additional homologous regions in the left hemisphere. Since the purpose of the study was to investigate dynamic connectivity changes between brain regions due to the intervention, an analysis based on the selected 18 regions was undertaken instead of using canonical regions.

2.6.2 Analysis Outline

We performed the analysis separately for the pre-intervention and post-intervention data, under both the dynamic pairwise correlations and dynamic precision matrix estimation methods. We used age and gender as covariates for the pre-intervention dataset, while also using the type of intervention (spin or non-aerobic control) as an additional covariate for the post-intervention analysis. Our analysis is designed

ROI Number	Region name	Broadmann area	MNI coordinate
1	R Cerebullum 1	NA	(5,-62,-57)
2	R Inferior Temporal Gyrus	20	(41,-27,-30)
3	R Angular Gyrus	39	(44,-56,12)
4	R Middle Frontal Gyrus	10	(23,56,-6)
5	R Middle Temporal Gyrus 1	22	(53,-12,-9)
6	L Precuneus 1	7	(-9,-74,57)
7	L Cingulate Gyrus	NA	(-9,-33,39)
8	R Precuneus	7	(6,-80,48)
9	R Cerebellum 2	NA	(35,-53,-27)
10	R Middle Temporal Gyrus 2	21	(60,-45,-6)
11	R Inferior Frontal Gyrus/precentral gyrus	44	(59,9,9)
12	R Retrosplenial Area	30	(9,-47,18)
13	R Supramarginal Gyrus	40	(41,-36,33)
14	R Pars Triangularis/MFG	45	(47,47,-9)
15	L Precuneus 2	7	(-6,-71,45)
16	L Cuneus	19	(-15,-80,27)
17	L Superior Frontal Gyrus	6	(-17,-18,69)
18	R Middle Temporal Gyrus 3	22	(60,-36,0)

Table 2.7: Summary of brain regions used for analysis. R and L are abbreviations for right and left respectively.

to: (i) investigate the clustering behavior and inspect how these clusters differ with respect to demographics and the intervention type; (ii) investigate the cluster-level network differences using network summary measures; (iii) estimate the connectivity change points and examine how well they align with the changes dictated by the block task experiment; (iv) infer nodes and edges in the network with significantly different connectivity patterns between pre- and post-intervention.

Objective (i) enables us to characterize homogeneous dynamic connectivity patterns corresponding to clusters of samples in terms of their demographic and clinical characteristics; aim (ii) will be instrumental in interpreting the cluster-level network differences that will shed light on network variations across transient network states; aim (iii) will provide insights regarding the effectiveness of the proposed approaches in terms of recovering connectivity jumps where these changes are influenced by, but often not fully aligned with, the changes in the block task experimental design

(Hindriks et al., 2016; Kundu et al., 2018); and aim (iv) will inform investigators regarding dynamic connectivity differences that are associated with the type of intervention. For aim (ii), we were only able to report results under dynamic precision matrix estimation, since a graph theoretic framework is necessary to compute the network summary measures, which may not be feasible under a pairwise correlation analysis.

2.6.3 Results

Cluster analysis: As seen from Table 2.8, the analysis under both idPAC and idPMAC methods yielded 5 clusters consolidated over all time scans (using the K-means algorithm described in Section 2.2, although the size of the clusters were more equitable under the idPAC method. The pre-intervention analysis yielded clusters that were largely homogeneous with respect to gender. These clusters were also reasonably well-separated with respect to age under the idPAC analysis, whereas the age of the participants within clusters were more diverse under the idPMAC analysis. The post-intervention analysis yielded more heterogeneous clusters with respect to both age and gender, with only one cluster comprising all males under both the idPAC and idPMAC analyses. This suggests a realignment of the dynamic connectivity after the intervention is administered, such that individuals with similar genders and age-groups have synchronous dynamic connectivity patterns pre-intervention as identified via subgroups, but the subgroups and their composition with respect to age and gender change post-intervention. Our post-intervention analysis also suggests that the variability across clusters under the idPAC method can be largely explained via the intervention type.

Connectivity change point estimation: Table 2.8 illustrates the cluster level connectivity change point estimation. We observed that under both the idPAC and idPMAC methods, the estimated change points were consistent with 4 or more (out of

Method	idPAC					idPMAC				
Cluster index	1	2	3	4	5	1	2	3	4	5
Cluster features	Pre-intervention					Pre-intervention				
Size	8	6	8	7	4	3	5	17	6	2
% of females	0	100	0	14	100	0	100	0	100	0
Age (mean)	72.2	65.8	64.7	76.7	67.7	71.7	69	70.4	66.8	67
Age(range)	69-73	60-72	60-68	74-80	66-69	63-78	62-80	60-80	60-72	66-68
CP(Task-Rest)	6	3	4	4	4	4	5	5	4	3
CP(Rest-Task)	3	5	2	3	4	4	4	2	4	3
	Post-intervention					Post-intervention				
Size	8	4	7	11	3	3	4	9	11	6
% of females	63	75	0	18	33	67	100	0	9	67
Age (mean)	67.3	65	65.1	74.5.	73.7	73.7	69.3	68.6	73	62.7
Age(range)	62-70	60-71	60-68	71-80	68-78	67-80	68-72	63-78	68-80	60-66
CP(Task-Rest)	5	6	4	3	6	3	3	5	5	5
CP(Rest-Task)	3	5	2	2	4	2	5	2	4	2
Spin(%)	0	100	100	0	100	33	0	100	9	50

Table 2.8: Results for analysis of block task fMRI experiments. Size refers to the number of participants in each cluster, 'CP(Task-Rest)' and 'CP(Rest-Task)' denotes the cluster level connectivity change points that were detected within +/- 2 time scan of the change in experimental design from task to fixation, and from fixation to task, respectively. 'Spin' refers to the percentage of individuals assigned the Spin intervention belonging to each cluster.

6) changes in experimental design when transitioning from task to rest, except one cluster where 3 of the connectivity change points aligned with the experimental design. These patterns were consistent in both the pre- and post-intervention analysis; however the number of connectivity change points that were strongly aligned with changes in the experimental design were (on average) greater in the post-intervention analysis compared to the pre-intervention analysis. This suggests a learning effect of the task that was reflected in terms of higher concordance between the connectivity change points and the experimental design post-intervention. On the other hand, the cluster-level estimation of change points when transitioning from fixation to task was (on average) less aligned with the experimental design compared to the change points when transitioning from task to fixation, as seen in Table 2.8. This is somewhat

expected since there were only 3-5 time scans in each fixation block, which made it extremely challenging to detect connectivity changes when transitioning from fixation to task. However, the proposed approach was still able to detect at least two, and often 3 or more connectivity change points (out of 6) aligned with the experimental design that suggests a reasonable concordance between connectivity jumps and experimental transitions from fixation to task.

In contrast, the CCPD approach detected at most one or two connectivity change points, while the DCR method was not able to detect connectivity change points at all, which makes these results appear biologically impractical given the nature of the block task experiment. Although the changes in connectivity are not expected to be fully aligned with changes in the experimental design (Hindriks et al., 2016), one expects a certain degree of synchronicity between the two. Our results indicate that this is not captured at all via existing change point methods especially when there are rapidly occurring transitions in the experimental design, which highlights their limitations. Hence, our analysis clearly illustrates the advantages of pooling information across heterogeneous samples and incorporating covariate knowledge via a mixture modeling framework, which is simply not possible using existing approaches that rely on information from single subjects as in DCR, or that use empirical methods to pool information across individuals as in CCPD.

Cluster level network differences: In order to investigate the differences between the networks corresponding to the different clusters, we examined variations in dynamic network metrics that capture modes of information transmission in the brain. These network metrics include the characteristic path length (CPL) that measures the length of connections between nodes, and the mean clustering coefficient (MCC) that measures the clustering tendency averaged over all network nodes. Using permutation testing, we examined p-values to evaluate which pairs of clusters exhibited significantly different network summary measures. None of the clusters had significantly

different CPL values in the pre-intervention analysis, but several pairs of clusters exhibited significant CPL differences post-intervention. The CPL differences were particularly pronounced between the first and remaining clusters, as well as the last and remaining clusters in the post-intervention analysis. These two clusters also demonstrated the highest within cluster variability in CPL values amongst all clusters. Moreover, the number of pairs of clusters with significantly different MCC values increased from the pre-intervention to post-intervention analysis, with 8 out of 10 pairs of post-intervention clusters reporting significantly different MCC values compared to at least one other cluster. Hence, our results suggest greater variability in network organization between clusters in the post-intervention analysis compared to pre-intervention, which potentially reflects greater network heterogeneity after the 12 week intervention was administered.

Network differences pre- and post-intervention: We applied paired t-test with multiplicity adjustment (using Bonferroni correction) in order to infer which edges were significantly different between pre- and post-intervention at 5% level of significance, along with identifying which network nodes contained the greatest number of differential edges. Since the magnitude of the pairwise correlations and the corresponding edge strength differences were higher, we discovered higher number of edges with differential edge strengths under the idPAC analysis. For both the idPAC and idPMAC methods, the bulk of the pre- vs post-intervention connectivity differences were concentrated in individuals in the spin group exclusively that were not present in the control group. We obtained 57 significantly different edges under the idPAC analysis, and 38 significantly different edges under the idPMAC analysis, which were exclusive to the spin group - see Figure 2.4. In contrast, the number of significantly different edges between the pre- and post-intervention networks under the idPAC analysis were 20 corresponding to both the spin and control groups, and 7 corresponding to the control group only. Moreover the idPMAC analysis did not produce any significant

edge level differences between the pre- and post-intervention networks corresponding to both the intervention groups as well as for the control group only. Our results suggest a considerably strong realignment in dynamic connectivity after the 12-week intervention that were exclusive to the spin group, compared to negligible changes in the control group.

The changes between the pre-vs post intervention networks that occurred exclusively in the spin group under idPAC analysis were concentrated in the following brain regions: Right Angular Gyrus(8 edges), Left Precuneus(10 edges), Right Cerebellum(9 edges), Right Middle Temporal Gyrus(11 edges), and Right Middle Temporal Gyrus(8 edges). Similarly the following brain regions had the highest number of differential edges pre- vs post-intervention under the idPMAC analysis: Right Middle Frontal Gyrus(16 edges), Right Cerebellum(6 edges), Right Pars Triangularis/MFG(8 edges), and Right Middle Temporal Gyrus(7 edges). Two nodes, Right Cerebellum and Right Middle Temporal Gyrus had a large number of significantly differential edges under both idPAC and idPMAC analyses, while the right middle frontal gyrus had, by far, the largest number of differential edges (16) under the dynamic precision matrix analysis. In addition, we also observe that more nodes in right hemisphere of the brain have significantly differential connectivity, which is to be expected since the majority of the 18 brain regions being investigated lie in the right hemisphere.

The large number of differential connections with respect to the right cerebellum is believed to be attributable to the generation of internal models or context specific properties of an object (Moberget et al., 2014), and preferential activation during a semantic challenge (D'Mello et al., 2017). The connectivity between the right cerebellum and inferior frontal regions has been noted in earlier studies (Balsters et al., 2013), with the inferior frontal regions being responsible for ordering language and codifying the motor output for syntax (Balsters et al., 2013). Moreover, the differential connectivity in the right middle temporal gyrus is along the lines of earlier

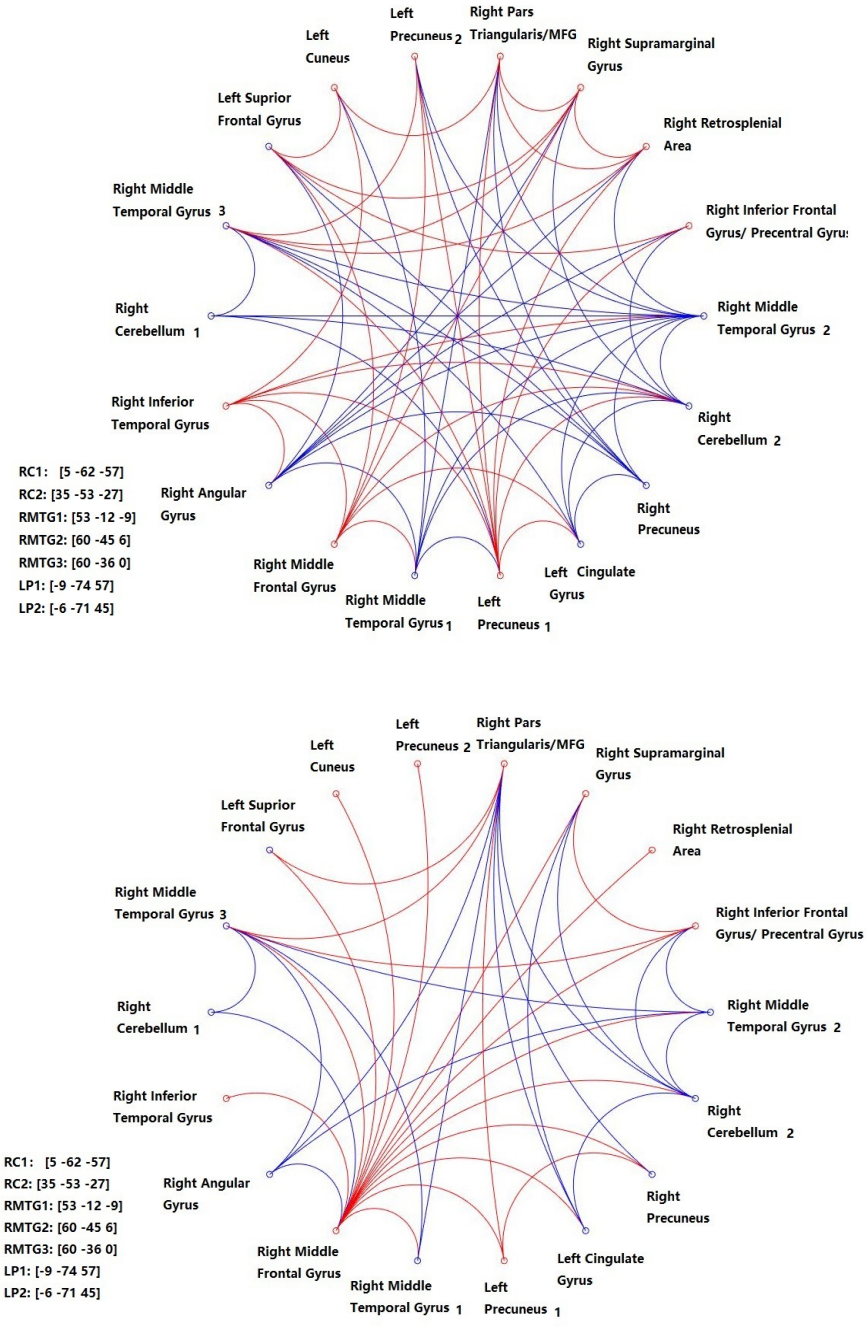

Figure 2.4: Circle plots for the edges that are significantly different pre- and post-intervention in spin group but not in the control group. The top and bottom panel correspond to the results under dynamic pairwise correlation and dynamic precision matrix estimation incorporating covariates, respectively. Red and blue lines correspond to lower or higher edge strengths in the pre-intervention network compared to post-intervention. RC1 and RC2 refer to the two brain regions in the right cerebellum; RMTG1-RMTG3 refer to the three brain regions in the right middle temporal gyrus; and LP1-LP2 refer to the two regions in the left precuneus. The MNI coordinates for these regions are provided in the Figure legend.

findings that illustrated the role of the left temporal gyrus as a hub for integration of sensory input into a transformation to semantic forms (Davey et al., 2016), and the corresponding connectivity differences in the right middle temporal gyrus may be attributable to a shift in laterality of involvement (Lacombe et al. 2015) due to aging. Finally, the large number of differential edges corresponding to the right middle frontal gyrus is potentially associated with semantic priming in older adults (Laufer et al., 2011). Given that this region is associated with executive function (Wang et al., 2019; Jolles et al., 2013) and is well characterized as being involved in working memory tasks, it is likely for connectivity differences to be focused on this region since the semantic task requires a continuous reference to working memory.

2.7 Discussion

In this article, we developed a novel approach that accurately estimates a population of subject-level dynamic networks by pooling information across multiple subjects in an unsupervised manner under a mixture modeling framework using covariates. The proposed approach, which is one of the first of its kind in dynamic connectivity literature, results in significant gains in dynamic network estimation accuracy, as illustrated via extensive numerical studies. The gains under the proposed method become particularly appealing compared to existing approaches in the presence of rapid transitions in connectivity as evident from our fMRI block task analysis. The proposed approach works best in fMRI task experiments involving a group of heterogeneous individuals executing the same task protocols, and in the presence of a carefully chosen set of covariates that are related to the dynamic network.

We also illustrate the robust performance of the proposed approach in the presence of a limited number of covariates that are not related to changes in connectivity, although the performance deteriorates as the number of spurious covariates increase.

In the presence of a large number of features that may not be necessarily related to dynamic connectivity, one can perform a screening step to exclude unimportant predictors from the analysis. This step will involve examining the associations between each covariate and the dynamic connectivity estimates obtained from the covariate naive BPMM approach, and subsequently only retaining the covariates with significant associations for analysis using the full model. This approach is expected to work well as long as the screening step does not exclude any important covariates and manages to largely filter out spurious covariates that are unrelated to the network. In future work, we plan to extend the proposed approach to incorporate feature selection that automatically identifies significant covariates that are related to the dynamic networks, and down-weights the contribution of unimportant covariates using Bayesian shrinkage priors. We note that although our analysis included covariates that do not vary with time, the proposed BPMM approach can be easily generalized to include time-varying covariates that are collected in-scanner (such as behavioral performance), when required.

Moreover for networks with higher densities or those with large variability in the values of the non-zero elements in the precision matrix, it is possible that the proposed approach may result in sub-optimal performance under a small number of mixture components. This is due to the symmetry constraint, which may potentially impose some restrictions on the clustering of the precision matrix elements when the number of clusters is small, and hence lead to inaccurate estimates. In such cases, a larger number of clusters would be required to fit the data well. Another avenue to tackle such potential restrictions is to generalize the proposed BPMM approach in equation (3) so as to specify independent mixture priors on each element of the upper triangular precision matrix subject to the constraint $\Omega_t^{(i)} \in M_V^+$, which will impose the same marginal distributions on elements (i, j) and (j, i) in the precision matrix. We plan to explore such generalizations in future work.

In addition to identifying important connectivity changes, during the fMRI block task experiment, our analysis conclusively established major changes between the pre- and post-intervention networks that were exclusive to the spin group. We note that existing literature has established the role of cardiovascular fitness in regulating aging related declines in both language and motor control (McGregor et al., 2011, 2013). However, much less is known about the effect of exercise intervention on dynamic connectivity, particularly in older adults. Because connectivity is a fundamental aspect of neuronal communication required for high-level cognitive processes, it is important to understand the potential impact of aging and/or aerobic exercise interventions in aging on changes in brain connectivity.

Further, our analysis also discovered subgroups of individuals with homologous dynamic connectivity, where the heterogeneity within these subgroups with respect to intervention was higher under the idPMAC method compared to the idPAC analysis. This indicates that dynamic pairwise correlations were more accurate in classifying participants in terms of the intervention administered. It is important to note that the separation of clusters with respect to intervention reflects the distinct patterns of dynamic connectivity between the 18 brain regions specified in our study that are known to be differentially activated in spin and control groups (Nocera et al., 2017). However, if additional regions are included that may not be necessarily associated with intervention type, it is entirely possible to obtain more heterogeneous clusters that have a more equitable composition with respect to intervention group. This is due to the presence of network edges between regions that are not necessarily associated with intervention and hence behave similarly in both the spin and control groups. Future work will focus on a more general analysis involving a larger number of cannonical regions known to be associated with the semantic language function.

Chapter 3

Non-parametric Bayesian Support Vector Machines for Brain Network-based Classification

3.1 Introduction

There is increasing evidence about the potential of resting state functional Magnetic Resonance Imaging (rs-fMRI) in terms of informing diagnosis of different disease conditions. In particular, the role of functional connectivity (FC) derived from rs-fMRI data has become increasingly prominent in terms of disease classification and prediction of disease severity, and there is a growing body of literature focusing on developing connectivity-based neuroimaging biomarkers for different diseases. This is consistent with the increasing evidence that different regions of the brain communicate with each other to drive cognitive process or disease etiology. There are several approaches for constructing networks for brain FC (Lukemire, Kundu, Pagioni, Guo, 2020; JASA). These approaches rely on selecting a pre-defined set of brain regions or nodes, and subsequently computing the correlations between them, which represents

the brain connectivity. When the correlations are time-averaged and stationary over the length of the fMRI session, these measures correspond to static connectivity; whereas more recent literature has also included dynamic connectivity models that allow time-varying correlations (Kundu, Ming, Nocera, McGregor, 2021). Moreover, the connectivity is typically computed in terms of pairwise correlations or partial correlations, with the latter being increasingly favored in the statistical and machine learning literature due to it's ability to regress out the effects of third party nodes that may otherwise induce spurious associations between regions (Smith et al., 2011). These correlations can be thresholded to infer edges or connections in the network. Each p-node network can also be represented as an $p \times p$ connection/adjacency matrix comprising binary off-diagonal elements that indicate the presence or absence of edges. In an undirected network that is of interest in this article, the adjacency matrices are symmetric, and the edge set can be fully represented by the off-diagonal elements.

Linear regression models for prediction, and linear classifiers such as the support vector machines (SVMs) have been the workhorse for the large majority of literature, although more recent methods have included non-linear regression models (Ma, Kundu, and Stevens 2022) as well as deep learning based approaches (Meng and Xiang, 2018). While deep learning methods are attractive in terms of providing an end-to-end pipeline, they often lack interpretability, and require a large training samples to adequately fit a massive number of embedded model parameters. Hence deep learning based methods may not perform well for small to moderate sample sizes due to overfitting or other related issues, and it may not be straightforward to infer the important network features associated with the phenotype. Similarly, there are a limited number of non-linear approaches such as Gaussian process models based on high-dimensional non-Euclidean network covariates in literature, and they require further investigation. In contrast, SVM-based linear classifiers assume linear relationships and hence are more interpretable, it is possible to infer important features

via appropriate priors in a Bayesian set-up, and SVM classifiers based on network connectivity features have been widely and successfully used in several disease areas including schizophrenia, bipolar disorder, autism spectrum disorder (ASD), attention deficit hyperactivity disorder (ADHD), Alzheimer's disease (AD) and mild cognitive impairment (MCI), and so on. See the review article by Du, Fu, and Calhoun (2018).

It is to be noted that the overwhelming majority of SVM-based classification approaches using FC as covariates rely on penalized or optimization methods to derive point estimates for model fitting. While penalized approaches are often useful in practical applications, there are several limitations. First, the results under penalized SVM methods are often sensitive to the choice of the penalty parameter that is often obtained by tuning or cross-validation. The problem is exacerbated in brain network applications, since the choice of the penalty parameter reflects an overall sparsity level for the model, but it may not be adaptive to different levels of shrinkage across different subsets of network edges in the model, resulting in inadequate performance. For example, in network applications, it is well-known that a large percentage of the network edges have zero or negligible effects, while the variability in the mental health outcomes are driven by a small proportion of significant edges with varying importance. Existing linear approaches involving L-1 penalties and regulated by a single penalty parameter, may not perform well in these settings since they are not readily adaptive to differential sparsity levels. For example, they are typically known to result in inflated models involving an overly large set of estimated non-zero features in linear regression settings (Chang, Kundu, Long, 2018), which often lead to poor feature selection performance and low accuracy. Second, penalized SVM approaches fail to report measures of uncertainty for the estimates, which is of paramount importance in high-dimensional neuroimaging applications. This is due to the fact that the brain network is derived from rs-fMRI data via suitable algorithms, and hence are subject to measurement error and mis-specifications. In these scenarios, a point

estimate may not suffice, and additional measures of uncertainty via credible intervals may be highly desirable for the parameter estimates. Third, inferring significant effects under existing penalized SVM methods often require computationally expensive procedures such as permutation tests (Lukemire, Kundu, Pagioni, Guo, 2020) that may be challenging to implement for high dimensional network applications. As a result, these methods may yield inaccurate feature selection results.

To address the above challenges in this article, we propose a novel non-parametric Bayesian support vector machine (BNP-SVM) approach for classification, based on high-dimensional brain networks. The key motivation behind the non-parametric Bayesian approach is to adaptively cluster the subsets of edges according to similarity in edge importance, with different clusters reflecting different degrees of importance of the edges in the SVM model. By pooling information across subsets of edges and allowing for different degrees of Bayesian shrinkage, the proposed approach is able to estimate the edge-specific importance in a robust and accurate manner, which results in considerable improvements over existing approaches. The number of clusters and the cluster memberships for the edge importance weights are determined in a data adaptive manner, under a fully non-parametric Bayesian set-up. The method is scalable to high-dimensional networks and performs well in terms of feature selection and classification performance in these high-dimensional settings. The proposed method overcomes the limitations of penalized or optimization based SVM methods described earlier. For example, it is naturally able to quantify uncertainty for the parameter estimates via credible intervals, and it provides an integrated framework for inference required for feature selection. In addition, it has considerable advantages over parametric Bayesian approaches in terms of ensuring model parsimony, greater computational efficiency, and superior numerical performance.

The main contributions of this article are as follows. First, the proposed approach is one of the first non-parametric Bayesian SVM based classification approach based

on static brain networks that is able to adaptively estimate edge importance under a Dirichlet process mixture (DPM) framework and infer significant network edges. This novel approach focuses on clustering edge importance weights and on binary labels, which is distinct from non-parametric models for multi-class classification (Nguyen et al., 2016) in machine learning literature. Second, given the recent success of dynamic networks in predicting intelligence (Sen and Parhi, 2021; Liu et al., 2018), we extend the BNP-SVM approach for classification (based on static networks) to incorporate dynamic networks, which is able to account for the temporal correlations of edges in the time evolving network and is able to naturally infer significant dynamic sub-networks. Our analysis provides a direct comparison regarding the classification performance of static and dynamic resting state networks, which is another important contribution of independent interest. Third, we extend the proposed method to integrate static network data from multiple fMRI sessions performed on the same set of individuals. There is very limited work in neuroimaging literature for such integrative analysis, although recent studies have indicated the benefits of performing multi-task and multi-session analysis (Xiao et al., 2020; Ma and Kundu, 2022). However, while the limited development in this area has primarily relied on regularization techniques that yield point estimates, we propose a non-parametric Bayesian approach for systematically integrating multi-session resting state fMRI (rs-fMRI) networks for improved classification performance. Fourth, we evaluate the effectiveness of our proposed framework via extensive numerical studies on various datasets that include rs-fMRI data from: (i) the Human connectome project (HCP) with the goal to classify fluid intelligence as well as crystallized intelligence based on resting state brain FC; and (ii) a ADHD dataset, where the goal is to classify individuals with Attention-Deficit/Hyperactivity Disorder (ADHD) from neurotypical controls (NC). To our knowledge, our paper is one of the first to establish the utility of sophisticated Bayesian SVMs based on brain networks as a powerful classification tool

in neuroimaging and mental health studies.

3.2 Methods and Materials

3.2.1 Description of Datasets

HCP study and pre-processing details

The Human Connectome Project (HCP) contains resting state as well as task fMRI scans for adults, along with a battery of cognitive measurements including fluid intelligence and crystallized intelligence. The HCP developed a novel multimodal parcellation that defines brain regions based on a combination of structural and functional features (Glasser et. al, 2016) having 360 nodes. The atlas includes in a larger number of regions (better resolution) than used in many other atlases, such as the 90 regions used in the structurally-based AAL atlas (Tzourio-Mazoyer et al., 2002) or the 264 regions in the functionally-based Power atlas (Power et. al, 2011). In the HCP data, each subject had two sessions of resting-state fMRI scans on the first day, where each scan was 14:33 minutes with TR=0.72, resulting in 1200 time points for each voxel. Further, these two scanning sessions were also repeated on a second follow-up day. We used the cortical surface data from the FIX pre-processed left-right phase-encoding scan for our analysis. For more details of pre-processing, see Smith et. al (2013). The pre-processing pipeline included slice timing correction, rigid body re- alignment, and normalization to the EPI version of the MNI template. The time courses were temporally detrended in order to remove gradual trends in the data. CompCorr (Behzadi et al., 2007) was performed to estimate and remove spatially coherent noise from the white matter and ventricles, along with the linearly detrended versions of the six rigid body realignment parameters and their first derivatives. From here, the data was spatially smoothed with a 6mm FWHM Gaussian kernel and bandpass filtered

be- tween 0.01-0.1 Hz. Finally, spike correction was performed using the AFNI package (Cox, 1996), as an alternative to motion scrubbing. The time series for each vertex was extracted, mean centered and normalized to unit variance, and then each time point averaged by parcel, resulting in a dataset with 360 nodes and 1200 time points for each subject.

ADHD study and pre-processing details

Pre-processed resting state fMRI data used was obtained from The Connectomics in NeuroImaging Transfer Learning Challenge (CNI-TLC) involving children. The CNI-TLC data was amassed retrospectively across multiple studies conducted by the Center for Neurodevelopmental and Imaging and Research (CNIR) at the Kennedy Krieger Institute (KKI) in Baltimore, Maryland. The overall cohort includes 145 children diagnosed with ADHD, 25 children with a primary diagnosis of Autism Spectrum Disorder (ASD) who also meet the diagnostic criteria for ADHD, and 170 neurotypical controls (NC). In our analysis we used 120 ADHD children after excluding those with the additional ASD diagnosis, and matched this cohort with another 120 NC subjects. The data was downloaded from the Github repositories for the CNI-TLC study that included 100 samples each from ADHD and HC cohorts for training (https://github.com/mdschirmer/2019_CNI_TrainingRelease/), and 20 samples from each cohort for validation (https://github.com/mdschirmer/2019_CNI_ValidationRelease). The rs-fMRI data used in this challenge was acquired on a Philips 3T Achieva scanner housed in the F.M. Kirby Research Center for Functional Brain Imaging at KKI. The acquisition protocol used a single shot, partially parallel gradient-recalled EPI sequence with TR/TE 2500/30ms, flip angle 70 degrees, and voxel resolution $3.05 \times 3.15 \times 3mm^3$. The scan duration was either 128 or 156 time samples. Children were instructed to relax with their eyes open and focus on a central cross-hair, while remaining still for the duration of the scan. All participants completed a mock

scanning session to habituate to the MRI environment. The rsfMRI data was pre-processed using a pipeline developed by CNIR and implemented in SPM-12 (Friston et al., 2007), and described in Schirmer et al. (2021). The data based on three standard parcellations were used : (1) the AAL atlas (Tzourio-Mazoyer et al., 2002), which consists of 90 cortical/subcortical regions and 26 cerebellar regions, (2) the Harvard- Oxford atlas (Desikan et al., 2006), which consists of 110 cerebral and cerebellar regions, and (3) the Craddock 200 atlas (Craddock et al., 2012), which is a finer parcellation of 200 regions. The choice of atlases enabled analysis at multiple spatial scales. Additional demographic variables including age, and sex, were also included in the analysis.

3.2.2 Classification Using Static Networks

Here, we develop a novel Bayesian SVM approach to efficiently classify subjects into two groups based on their static brain network and supplemental covariates. Towards this aim, we considered a generic dataset with n samples in the training set where the binary outcome $z_i \in \{-1, 1\}$, the corresponding full edge set $\mathbf{x_i}$, and denote the non-network related covariate information as $\mathbf{c_i}$ $(C \times 1)$ for subject i.

Static Functional Connectivity Estimation

Assume the total number of subjects is N, with T_i scans for the i-th subject. For subject i, the $V \times T_i$ fMRI data matrix is denoted as $Y^{(i)}$, which contains V regions of interest (ROI) and T_i time scans. In order to construct the binary network that is used as a predictor in our SVM model, we use existing approaches in literature to estimate a sparse inverse covariance or precision matrix. In particular, we first calculated the Pearson correlation matrix (denoted by Θ_i) for each subject using the observed fMRI data by averaging over all time points, and subsequently applied the graphical lasso algorithm (Friedman et al., 2008) to compute a sparse inverse

covariance matrix separately for each sample, although other algorithms can also be used. Numerically, the graphical lasso algorithm was implemented via the QUIC method proposed by Hsieh et al., (2013) that relies on quadratic approximations. The off-diagonal elements in the sparse precision matrix Ω_i corresponding to the ith sample denotes the strength of the network edges in terms of the conditional dependency between two nodes in the network given all other nodes. For example, a zero off-diagonal element implies conditional independence between the corresponding nodes. The amount of structural zeros or the sparsity level of the network is controlled via a tuning or penalty parameter (λ_{gl}) in the QUIC algorithm, with a larger λ_{gl} value resulting in sparser networks. We describe a method to choose a suitable λ_{gl} value in the sequel.

Preliminary edge screening: We propose a model-free approach for initial screening. We screen out those edges with minimal variations in edge strengths across samples in order to exclude edges that are not instrumental in explaining between-subject differences from the analysis. Such a screening step is able to reduce the dimension of the explanatory variables that are included in the classification model without compromising accuracy, while speeding up computations. For example, those edges that are present (absent) in all or almost all samples are not instrumental for differentiating between groups of individuals and will be screened out. We exclude all edges for which the variability in edge strengths across all sample was limited and did not exceed a certain threshold. We use the threshold of 0.1 for our analysis, although more stricter thresholds were also considered (0.05) but without any distinct improvements in accuracy. We will denote the screened set of edges used in the final classification model as $\mathbf{x}_i^*$, having dimension $Q \times 1$. We note that although the set of screened edges represent a viable set of candidate network connections that may drive differences across subgroups or classes, the final set of significant edges related to differences in class labels is inferred from amongst this set of initial screened edges

under the fully Bayesian SVM approach, as elaborated in the sequel.

Bayesian SVM via Pseudo-Likelihood

We will concatenate the edge set and supplementary covariates to form an augmented feature vector set, which is denoted as $\mathbf{u_i} = [\mathbf{x_i^*}; \mathbf{c_i}]$. We propose a linear classifier using a linear function of the features, i.e. $f(\mathbf{u_i}) = \mathbf{u_i}^{\mathbf{T}}\beta^* = \mathbf{x_i}^{*\mathbf{T}}\beta + \mathbf{c_i}^{\mathbf{T}}\gamma$ for binary classification. Here β is a $Q \times 1$ vector of edge-specific effects, γ is $C \times 1$, and β^* is of dimension $P \times 1$, where $P = Q + C$. Using the derivations in Polson and Scott (2011), the Bayesian SVM pseudo-likelihood can be represented as follows

$$
\begin{aligned}
L(z|u, \beta^*, \gamma, \sigma_\epsilon^2) &= \prod_{i=1}^{N} \frac{1}{\sigma_\epsilon^2} e^{-\frac{2\max(1-z_i {u_i}^T \beta^*, 0)}{\sigma_\epsilon^2}} \\
&= \prod_{i=1}^{N} \int_0^\infty \frac{1}{\sigma_\epsilon \sqrt{2\pi\rho_i}} exp\big(- \frac{(\rho_i + 1 - z_i u_i^T \beta^*)^2}{2\rho_i \sigma_\epsilon^2}\big) d\rho_i, \qquad (3.1)
\end{aligned}
$$

where, σ_ϵ^2 represents the scale parameter in the likelihood, and a latent parameter ρ is introduced to facilitate posterior computation. We note that the pseudo-likelihood is directly related to the hinge loss that is routinely used in SVM models, having the form $\max(1 - z_i {u_i}^T \beta^*, 0)$, where smaller value of the loss function implies large values of the pseudo-likelihood. Moreover, large values of σ_ϵ^2 results in sharper changes in the pseudo-likelihood with changes in the values of the linear mean function, as evident from Figure 3.1.

Non-parametric Priors for Sparse Learning

In practice, one expects only a subset of edges to be instrumental for classifying the binary outcome variable. Further, different edges may have different degrees of importance in the model. In order to ensure appropriate shrinkage of the edge effect sizes that allows for differential sparsity levels, we specify a Laplace or Double

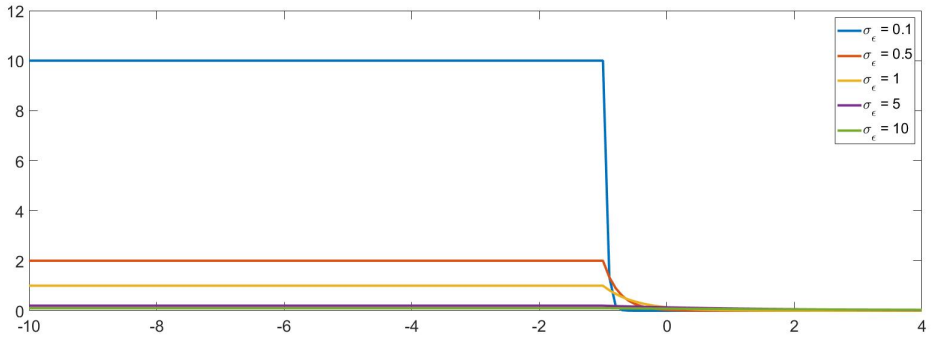

Figure 3.1: Plot of hinge loss under different values of σ_ϵ. $z_i = 1$, X-axis represents different value of ${u_i}^T\beta^*$.

exponential (DE) prior on $\beta^* = (\beta_1^*, \ldots, \beta_P^*)^T$ as:

$$\pi(\beta^*) = \prod_{p=1}^{P} \frac{\lambda_p}{2\sigma_\epsilon} \exp^{-\frac{\lambda_p}{\sigma_\epsilon}|\beta_p|}, \lambda_p^2 \sim P, P \sim DP(Mf_0), \tag{3.2}$$

which is modulated by the local shrinkage parameter $\lambda_p > 0$, that follows a non-parametric Dirichlet process (DP) prior with precision parameter M, and base measure f_0 that is chosen as $f_0 := Gamma(r, \delta)$ for ease in posterior calculations. Under a Bayesian specification, the prior on the scale parameter is specified as an inverse-Gamma prior with shape and scale parameters being a_1, b_1, i.e. $\sigma_\epsilon \sim IG(a_1, b_1)$. The overall specification results in a novel Dirichlet process mixture of Laplace distributions on the effect sizes β^*, which results in a novel Bayesian non-parametric SVM (BNP-SVM) approach.

The proposed approach is founded on the Bayesian Lasso method proposed by Park and Cassella (2008), where a large value of λ_p implies greater prior shrinkage towards zero for β_p (see Figure 3.2). However unlike in Park and Cassella (2008) that used global shrinkage, our approach is distinct in allowing for feature-specific shrinkage parameters. The novel specification in (3.2) results in clustering of shrinkage parameters ($\lambda' s$) across edges, where each cluster represents a distinct degree of shirnkage or importance of the coefficients in the SVM model. Moreover, the number

of clusters and the cluster memberships are unknown and learnt in an unsupervised manner. The precision parameter M controls the prior number of clusters (larger M implies more clusters). The non-parametric prior defines a class of prior distributions on the set of densities for λ, where the prior guess of the density is given by the base measure f_0. The following Lemma captures the closed form of the non-parametric mixture prior on the coefficients.

Lemma 1: *The prior distribution on the coefficients can be written as* $f(\beta_p^*) = \int \frac{\lambda}{2} \exp\{-\lambda|\beta_p^*|\} dP(\lambda) = \sum_{l=1}^{\infty} \pi_h f_h(\beta_p^*; \lambda_h)$, *where* $f_h(\beta_p^*; \lambda_h) = \frac{\lambda_h}{2} \exp\{-\lambda_h|\beta_p^*|\}$, $p = 1, \ldots, P$, *and* $\pi_h = \nu_h \prod_{h'=1}^{h-1}(1-\nu_{h'}), \nu_h \sim Be(1, M)$.

Lemma 1 summarizes the fact the the prior distribution can be represented as an infinite mixture distribution with stick-breaking weights, which follows directly from the results in Sethuraman (1994). The infinite mixture representation highlights the non-parametric nature of the proposed prior. Further, using results in Antoniak (1974), the number of distinct clusters Δ for the shrinkage parameters λ is guaranteed to increase with P as well as the parameter M, under the rule $E(\Delta) = \sum_{m=1}^{P} M/(M+m-1) \approx M \log((P+M)/M)$, where $E(\Delta)$ represents the prior expectation of Δ.

The proposed approach results in sharp improvements over parametric approaches that are not necessarily equipped to combine information across edges, by learning differential sparsity levels (represented by λ) across subsets of edges in an unsupervised manner. For example, the traditional Bayesian lasso approach (Park and Cassella, 2008) specifies a global shrinkage parameter λ across all edges, which is an extreme case of the proposed model with $\Delta = 1$, but can not capture differential sparsity patterns across features. An alternative method is to force unique shrinkage parameters across all edges, which is also a special case of the proposed approach with $\Delta = P$. However, this approach is expected to result in poor accuracy due to the inability to robustly estimate the importance weights by pooling information across edges. The

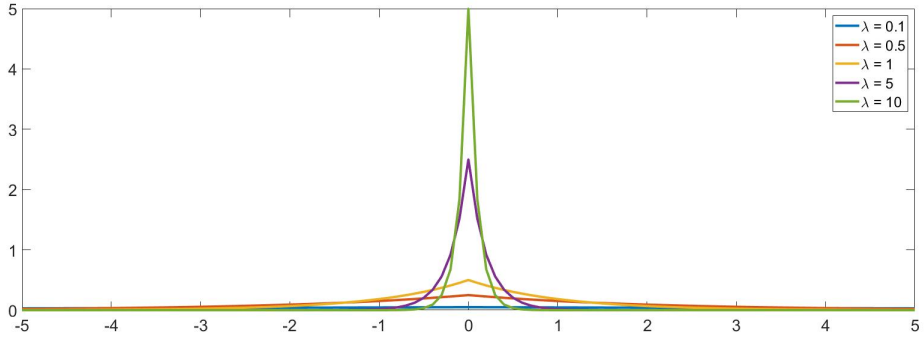

Figure 3.2: Plot of Laplace Prior under different values of β_p^*. X-axis represents different value of β_p^*

proposed approach provides a desirable balance between these two extreme scenarios, and results in data-adaptive learning of differential sparsity levels across unknown subsets of edges.

Posterior Computation

In order to facilitate posterior calculations, we re-express the Laplace distribution as a scale-mixture of normals with an exponential mixing density as $\beta_p^*|\sigma^2_{\beta_p} \sim N(0, \sigma^2_{\beta_p}),\ \sigma^2_{\beta_p} \sim \frac{\lambda_p^2}{2} exp\big(-\frac{\lambda_p^2}{2}\sigma^2_{\beta_p}\big)$, where $\lambda_p^2 \sim P,\ P \sim DP(Mf_0)$. The model parameters are updated via a Markov Chain Monte Carlo sampling (MCMC) sampling scheme. A fully Gibbs sampler was used for posterior computation, which cycles through the following update steps.

MCMC Steps: Algorithm I

(i) Use the following conditional distribution to sample σ_ϵ^2: $\sigma_\epsilon^2|\beta^*, u, z, \rho \sim IG\Bigg(a_1 + \frac{3N}{2}, b_1 + \sum_{i=1}^{N} \frac{(\rho_i + 1 - z_i {u_i}^T \beta^*)^2}{2\rho_i}\Bigg)$.

(ii) ρ_i is sampled using the inverse Gaussian conditional distribution $\pi(\rho_i^{-1}|-) = IN\big(|1 - z_i {u_i}^T \beta^*|^{-1}, \frac{1}{\sigma_\epsilon^2}\big), i = 1, \ldots, N$.

(iii) Use a multivariate normal distribution with mean μ_β and variance Σ_β to sample β^*, where $\Sigma_\beta = \big(1 + D_{\sigma_\beta}^{-1}\big)^{-1},\ \mu_\beta = \Sigma_\beta u^T z$, and D_{σ_β} represents the $P \times P$ diagonal

matrix with entries $(\sigma^2_{\beta_1} \dots \sigma^2_{\beta_P})$.

(iv) Use $\pi(\sigma^{-2}_{\beta_p}|-) = IN(|\frac{\lambda_p}{\beta^*_p}|, \lambda^2_p)$ to update $\sigma^2_{\beta_p}$.

(v) The conditional distribution of λ^2_p could be updated via slice sample for Dirichlet process atoms

Choice of the network sparsity:

The proposed BNP-SVM approach may be sensitive to the choice of the estimated static network that is used as the input feature vector for classification. In order to select the optimal network sparsity level, we fit the graphical lasso algorithm over a grid of shrinkage parameter values, and fit the proposed BNP-SVM based on the corresponding edge set from the estimated static network. For a given network, we calculate the classification accuracy in a validation sample. Subsequently, we choose the network sparsity level that results in the best validation accuracy, and use this choice for out-of-sample classification for test sample. Such a validation scheme provides a key understanding of network sparsity levels yielding optimal classification.

3.2.3 Extension to Classification Using Dynamic Networks

It is of interest to evaluate how the resting state static FC compares with resting state dynamic FC in terms of classification accuracy. To this end, we extend the Bayesian SVM to include dynamic connectivity features. Dynamic FC involves a time-varying connections for each edge, which need to be integrated into the model, thus posing methodological challenges. We use a set of extracted features from the connectivity time-series for each edge for our model, where these extracted features can be automated or can be manually specified based on subject matter expertise.

Dynamic Connectivity Estimation: Unlike static FC that used the precision matrix to capture the network edge strengths averaged over all time points, the dynamic FC approach calculates a connectivity matrix at each time point using a sliding window

approach, denoted as $\tilde{\Sigma}_2, \ldots, \tilde{\Sigma}_{T-1}$, which capture the time varying pairwise correlations. Instead of using the entire timecourse of pairwise correlations represented by the above sequence $\{\tilde{\Sigma}_{t'}\}_{t'=2}^{T-1}$, we extract summary features for each edge that reduces dimensionality with minimal loss of information. In particular, consider the extracted features $\{u_{kl,1}, \ldots, u_{kl,R}\}$ from the time-series of sliding window correlations $\{\tilde{\sigma}_{kl,2}, \ldots, \tilde{\sigma}_{kl,T-1}\}$ corresponding to edge (k, l), where the number of extracted features R is pre-specified. For subject i, let $u_i^r, r = 1 \ldots R$ be the $Q \times 1$ vector containing the subset of screened edges for feature r.

Proposed Model

The Bayesian SVM pseudo-likelihood based on dynamic FC may be represented as $L(z|u, \beta^*, \gamma, \sigma_\epsilon^2) = \prod_{i=1}^{N} \frac{1}{\sigma_\epsilon^2} \exp\{-\frac{2}{\sigma_\epsilon^2} \max(1 - z_i \sum_{r=1}^{R} {u_i}^{rT} \beta_r^* - {\mathbf{c_i}}^T \gamma, 0)\}$. Noting that the effects for the extracted features corresponding to the same edge are expected to be correlated, we propose a structured regression coefficient as $\beta_{q,r}^* = \beta_q^* \eta_r, r = 1 \ldots R, 1 \leq q \leq Q$, where $\{\eta_1, \ldots, \eta_R\}$ represent feature-specific correlated effects that are considered invariant across edges that is modeled as $\eta \sim N(0, \Sigma_\eta)$, with $\Sigma_\eta \sim IW(b, D), b = R, D = d^* \times I_R$, $d^* \sim IG(c, d)$, such that Σ_η captures the correlations within η. On the other hand, β_q^* reflects the global edge-specific contribution of the extracted features in the SVM model, with $\beta_q^* = 0$ implying no effect corresponding to the dynamic connection of edge q. The edge-specific parameters $(\beta_1^*, \ldots, \beta_Q^*)$ are assigned the same DP mixtures of Laplace prior as in (3.2). We will denote this model as dynamic BNP-SVM or dyBNP-SVM.

We note that while the structured coefficients in the dyBNP-SVM model $\beta_{q,r} = \beta_q^* \eta_r (r = 1 \ldots R, 1 \leq q \leq Q)$ can be generalized to accommodate greater flexibility, such modeling assumptions can lead to a dramatic increase in the number of parameters and slow down computations. For example, one can instead specify $\beta_{q,r} = \beta_q^* \eta_{q,r}$, which allows for edge-specific and feature-specific variations. However, it is clear that

the overall number of parameters for such a specification becomes exceedingly large with the increase in the network size, resulting in computational bottlenecks and potential overfitting issues. Hence we do not consider such generalizations further.

Posterior Computation

The form of the pseudo-likelihood using dynamic connectivity features lends itself to a similar treatment for MCMC sampling as in model (3.1) with static connectivity features, with some additional steps needed to sample the η effects and associated hyperparameters.

MCMC Steps based for dyBNP-SVM: Algorithm II

The posterior distributions of σ_ϵ^2 and ρ_i^{-1} are given similarly as in steps (i) and (ii) in Algorithm I involving static connectivity, with $\mathbf{u}_i$ replaced by ${u_{\eta,i}}^*$. Similarly the updates for $\sigma_{\beta_p}^2$ proceeds similarly as in Algorithm I. Moreover, the posterior distributions for β^* follows the Multivariate Normal distribution with mean μ_β and variance Σ_β with $\Sigma_\beta = \left(1 + D_{\sigma_\beta}^{-1}\right)^{-1}$, $\mu_\beta = \Sigma_\beta {u_\eta}^{*T} z$, where D_{σ_β} represents the $P \times P$ diagonal matrix with entries $(\sigma_{\beta_1}^2 \ldots \sigma_{\beta_P}^2)$. Similarly, the posterior distributions of η follows the Multivariate Normal distribution with mean μ_η and variance Σ_η^*, with $\Sigma_\eta^* = \sum_{i=1}^N \left\{ \widehat{u_{\eta,i}}^T \beta^* \beta^{*T} \widehat{u_{\eta,i}} \right\} + \Sigma_\eta$, $\mu_\eta = \sum_{i=1}^N \left\{ \frac{\beta^{*T} \widehat{u_{\eta,i}} (\rho_i + 1)}{\rho_i \sigma_\epsilon^2} \right\} \Sigma_\eta^*$, where $\widehat{u_{\eta,i}}$ is $R \times Q$ matrix with the $r-th$ row as $u_{ir}, r = 1 \ldots R$. The posterior distribution for updating Σ_η is given as $\Sigma_\eta \sim IW(b+1, D + \eta^T \eta)$, while that for updating d^* is given by $Inv - Gamma(c + Rb/2, d + 0.5 Tr(\Sigma_\eta^{-1}))$, where $Tr(\cdot)$ represents the Trace operator. The conditional distribution of λ_p^2 is given similarly as that in the posterior update in step (v) in Algorithm I.

Choice of window length

A common question in sliding window based approaches is how to optimally select the window length. In our approach focused on classification, we fit the proposed dyBNP-SVM model separately over a grid of window length values, and compute the classification performance in a validation sample. Subsequently, we choose the window length that results in the lowest mis-classification accuracy in the validation sample, and use this choice for out-of-sample classification for the test sample.

3.2.4 Classification via Integrating Multiple fMRI Sessions

In the HCP study, data was collected for two sessions (left-right and right-left phase encoding) on the first day (denoted as LR1 and RL1) as well as on another follow-up visit (denoted as LR2 and RL2) for each individual. One major question of interest in the neuroimaging community, is whether it is possible to develop classification approaches with higher accuracy via integrating information across multiple sessions. To address this question, we used static networks computed using resting state fMRI data from LR1 and LR2 sessions together in the proposed approach based on static connectivity. In other words, the edge sets for both the LR1 and LR2 networks were included jointly in the Bayesian SVM model for classification, where the network density for these networks was determined from the single-session static connectivity analysis as in Section B.5. A preliminary edge screening was performed as in Section B.1. for each of the two scans, and subsequently this subset of edges was used for classification.

3.2.5 Feature Selection

The fully Bayesian framework seamlessly enables one to perform inference. For example, testing for significant coefficients is by constructing $100(1-\alpha)\%$ credible intervals

based on the posterior distributions of $\beta_1, \ldots, \beta_Q$. Here α is the level of the credible set, which is usually selected as 5%, and it can be adjusted to account for multiplicity corrections. The resulting approach enables one to identify the significant network edges and other covariates that play an important role in classification under the SVM model. In order to determine which edges are consistently significant in the SVM model, we divide the overall sample into two halves, of roughly equal demographic compositions. Subsequently, we fit the Bayesian model and infer the significant edges separately for each half of the data. Then, those network edges that show up as significant for both slices of the data are considered as consistently significant and reproducible. Moreover, we also evaluate the concordance of feature selection corresponding to the two data slices via the Intraclass Correlation Coefficient (ICC), as a measure of feature selection reproducibility. Given that we are able to achieve almost perfect classification for the top and bottom 10% fluid intelligence groups corresponding to the LR1 scan (see results in the sequel), we focus on this subset of individuals for feature selection analysis. We do not consider larger subsets for feature selection, given that the reproducibility of results is expected to deteriorate for such subsets involving greater heterogeneity leading to lower classification accuracy.

3.3 Numerical Results Using rs-fMRI Data

3.3.1 Study Objectives

Human Connectome Project Data

Our goal is to use resting state functional connectivity to classify individuals belonging to various strata in the intelligence spectrum. To this end, we consider classifying individuals belonging to the top and bottom ζ% intelligence groups, where ζ is chosen

	Fluid Intelligence				Crystallized Intelligence			
	Bottom		Top		Bottom		Top	
	Age	Male	Age	Male	Age	Male	Age	Male
10%	28.4	46.5	28.7	52.5	27.6	45.5	29.1	50.5
12%	28.1	45.4	28.7	52.1	27.7	46.2	29.3	51.3
15%	28.1	47.7	28.6	50.3	27.5	45.6	29.3	53.0
18%	28.2	46.9	28.6	49.4	27.7	48.0	28.2	55.3
20%	28.3	47.9	28.6	49.0	27.8	48.5	29.3	55.6
25%	28.5	50.1	28.7	48.0	27.9	49.6	29.2	54.8

Table 3.1: Demographics information of HCP data

to be $10\%, 13\%, 15\%, 18\%, 20\%, 25\%$, to reflect different levels of intelligence. The demographics for these different population strata are given in Table 3.1. We perform the analysis separately for both fluid intelligence and crystallized intelligence, with the goal being to classify them into high or low intelligence groups. We note that while the classification accuracy is expected to be high when classifying the top and bottom 10% intelligence levels, it is likely to deteriorate for larger sub-populations with lesser separation between the low and high intelligence groups. Age and gender are included as explanatory variables, along with the brain network that is calculated using the multimodal Glasser atlas with 360 nodes or regions of interest that was developed by the HCP study. We use rs-fMRI data from both LR1 and LR2 scans for analysis.

We perform three different streams of analysis for each intelligence strata. First we use static FC to evaluate classification performance using data from one scanning session, separately for the LR1 and LR2 scans from HCP. Subsequently, we repeat this analysis with dynamic connectivity for these scans. Finally, we combine the resting state static networks from LR1 and LR2 scans to perform an integrative classification analysis. The three separate analysis provides us with an understanding of which types of network-based features (single-session or multi-session analysis based on static networks, or dynamic networks) derived from resting-state fMRI data is expected to yield superior classification accuracy. For the static connectivity analysis using one scanning session, we computed the network across varying sparsity levels

by tuning the λ_{gl} parameter over a grid for each individual, and we chose the network density for all individuals that corresponded to a pre-determined sparsity level. We considered a range of pre-determined network sparsity values varying from 10% to 25%. The classifier is fit using the same degree of network sparsity for all individuals, and the optimal network sparsity is selected based on the best classification accuracy. The pre-screening procedure as described previously was used to select a subset of edges to be used in the classifier.

For the multi-session analysis, we used the network sparsity level that was deemed optimal under the single session analysis for LR1 and LR2 scans. Moreover, all edges that were included after pre-screening for the LR1 and LR2 scans were combined to obtain the pre-screened edge set for the classification analysis involving the multi-session data. Finally, for classification based on dynamic functional connectivity, we considered different methods for extracting the features from the time-varying sliding window correlations for each edge that would be subsequently used in the classification analysis. These methods included (i) the three manually specified features used in Liu et al., (2019) that involved the mean FC along with measures of variation and stability of the time-varying FC; (ii) features extracted from a principal component analysis (PCA) of the time-varying functional connectivity, that retained at least 95% variability; and (iii) features extracted from independent component analysis (ICA), where the number of components was chosen to be the same as in the PCA analysis. Hence, our analysis is able to evaluate which set of features extracted from the dynamic connectivity profiles (manually selected or data-adaptive) yields the best classification performance.

ADHD Data

In this analysis, we used resting state fMRI data from the CNI transfer learning challenge (CNI-TLC), to classify individuals with ADHD, and neurotypical controls

(TC), based on resting state FC. Te full details for this dataset can be found in Schirmer et al. (Medical Image Analysis, 2021, 70). The training data comprised 100 ADHD and 100 TC individuals, and the testing data had 20 individuals in each cohort. The mean age for the ADHD and TC groups were 10.4(1.5) and 10.3(1.2) respectively, and the gender distribution for these groups comprised 31.0% boys for ADHD and 30.0% boys for TC. The SVM model was trained separately on data coming from the three parcellations AAL (116 ROIs), Harvard Oxford (110 ROIs) and Craddock200 (200 ROIs), and the relative performance over the three parcellations was compared. Coupled with the HCP analysis, this second analysis highlights the generalizabilty of the proposed BNP-SVM approach for different disease areas, and different choices of brain parcellations, that goes beyond intelligence classification.

3.3.2 Comparison Methods and Metrics

We evaluate the performance of the proposed method in the context of several state-of-the-art competing methods for classification. The first competing method is penalized SVM (Becker et al., 2009), which use the combination of SCAD and ridge penalties as the penalty function of SVM to overcome the limitations of each penalty applied separately. As the performance of SVM based model is highly associated with the choice of tuning parameters, the penalized SVM was trained using different tuning parameters and the cross-validation was used to select the best tuning parameter. The second competing method is penalized logistic regression, where the elastic net penalty (the combination of L-1 and L-2 penalties) was used. Again, a grid of tuning parameters were tested and the combination with best performance based on cross-validation was used in the final fitted model.

We evaluate the performance of different approaches using out-of-sample classification accuracy of the test samples, via several metrics popularly used in literature (Schirmer at al. 2021). In particular, we report the mis-classification rate defined as

the ratio of incorrect classification labels over the total number of instances evaluated, the $F1$ score that is computed as the harmonic mean between recall (sensitivity) and precision, and the informedness that is also known as Youden's J statistic and summarizes the true positive and true negative rates of a classifier. We note that precision is defined as the fraction of correctly classified positive samples in relation to the number of the total positive classified samples, while sensitivity is defined to be the fraction of positive samples that are correctly classified. As a classification method, classification accuracy is the most direct measure of the performance and is considered superior for a method with low mis-classification accuracy and higher $F1$ score, geometric mean and Youden's J statistic.

3.3.3 Analysis Results for HCP Data

Results based on Static Connectivity

As seen from Figure 3.3, the proposed DP mixture of Laplace approach consistently resulted in superior classification accuracy across all the network sparsity levels compared to the different competing methods. From amongst the different network sparsity levels over which we conducted our analysis on, the best classification accuracy for both fluid as well as crystallized intelligence was obtained with network density as 0.2 under the LR1 scan, and 0.18 for the LR2 scan under the proposed approach. This finding is of independent interest, and can potentially shed light on the optimal network densities required to obtain good classification accuracy for cognition/intelligence based on static connectivity obtained via sparse precision matrices. The results for classification performance using static connectivity for the HCP data under all approaches are reported in Table 3.2 under these optimal network sparsity levels, corresponding to both fluid intelligence and crystallized intelligence. The same network sparsity levels are used to report results under all approaches in Table 3.2 for the sake of consistency, and given the consistently superior performance of the

proposed approach over varying sparsity levels.

Several aspects are clear from the reported results. First, the classification accuracy under all approaches was the highest when classifying the subgroups corresponding to the top and bottom 10% intelligence groups, but the accuracy deteriorates as the size of these subgroups are increased gradually to the top and bottom 25% brackets, which is expected given the decreasing separation between high and low intelligence groups for more heterogeneous sub-populations. Second, the proposed DP mixture of Laplace approach has statistically significant improvements in terms of F-1 score, informedness, and the geometric mean metrics for almost all cases, barring a few exceptions. The mis-classification rate is also significantly lower under the proposed method for the overwhelming majority of cases when classifying fluid intelligence, but the penalized SVM approach is seen to have comparable mis-classification rates under the penalized SVM approach for some scenarios involving crystallized intelligence classification. Unfortunately, all competing approaches consistently had inferior classification performance compared to the proposed method even for subpopulations with lesser separation between the high and low intelligence groups, due to their inability to adaptively pool information across features in order to determine differential sparsity levels. These results clearly illustrate the prowess of the proposed Bayesian approach across varying intelligence spectrum, which highlights it's utility as a classification method for connectome fingerprinting. Third, by including age and gender along with the static network in the model, it is possible to obtain improved classification accuracy compared to using only the network in the SVM model (results not reported here). This illustrates the importance of including demographic features (particularly age) in classification models for intelligence and cognition.

Fourth, we discover that the classification performance is generally improved under the LR1 scan compared to the LR2 scan under all methods for fluid intelligence classification, but on the other hand, classification accuracy is often higher corre-

sponding to the LR2 scan under all approaches. These results are consistent with the definitions of crystallized intelligence that is based on accumulated knowledge from the past, meaning that individuals were likely to be increasingly familiar with the experience with the resting state fMRI experiment during the LR 1 scan, that resulted in higher separation (in the form of increased classification accuracy) between the high and low intelligence groups. In contrast, since fluid intelligence that measures intuitive and spontaneous reasoning, the separation between the top and the bottom fluid intelligence sub-groups ($10\% - 15\%$) was greater during the first LR1 scan where all individuals participated in the fMRI experiment for the first time.

Results based on Dynamic Connectivity

Table 3.3 provides the result of fluid intelligence classification based on dynamic functional connectivity. From amongst the three types of features extracted from the sliding window correlations, the manually extracted features defined in Liu et al., (2019) yielded the best classification accuracy across all approaches. Hence, our analysis is able to validate the utility of the manually crafted metrics for dynamic functional connectivity proposed in Liu et al. (2019) and is consistent in terms of reporting poor classification accuracy under ICA and PCA decomposition as also reported in Sen and Parhi (2020).

As for the static connectivity case, the proposed approach is shown to have higher classification accuracy compared to the competing methods. Moreover, the dynamic connectivity based classification results are less accurate compared to resting state static connectivity results across all methods. This result illustrates that static FC derived from resting state fMRI is better suited for classifying intelligence levels compared to resting state dynamic connectivity, which is potentially due to the fact that temporal fluctuations in connectivity during the scanning session that are captured via dynamic FC may not be directly related to variations in intelligence. This finding

	Fluid Intelligence Classification								
	LR1			LR2			Multi-session		
10%	MC	F1	inf	MC	F1	inf	MC	F1	inf
DPL-SVM	0.00	0.99	0.99	0.05	**0.97**	**0.93**	0.00	0.99	0.99
L1-SVM	0.00	0.99	0.99	0.06	0.93	0.88	0.00	0.99	0.99
L1-Log	0.00	0.99	0.99	0.06	0.93	0.88	0.00	0.99	0.99
12%	MC	F1	inf	MC	F1	inf	MC	F1	inf
DPL-SVM	**0.09**	**0.92**	**0.84**	**0.14**	**0.87**	**0.72**	0.07	0.94	0.87
Pen-SVM	0.12	0.87	0.75	0.17	0.83	0.68	0.07	0.93	0.87
Pen-Log	0.12	0.87	0.75	0.18	0.81	0.63	0.08	0.91	0.83
15%	MC	F1	inf	MC	F1	inf	MC	F1	inf
DPL-SVM	**0.17**	**0.85**	**0.68**	0.25	**0.78**	**0.54**	0.14	0.87	0.73
Pen-SVM	0.19	0.80	0.62	0.26	0.73	0.48	0.15	0.84	0.70
Pen-Log	0.21	0.80	0.61	0.27	0.72	0.46	0.15	0.85	0.71
18%	MC	F1	inf	MC	F1	inf	MC	F1	inf
DPL-SVM	0.25	**0.77**	**0.51**	0.25	**0.78**	**0.53**	0.23	0.79	0.56
Pen-SVM	0.26	0.73	0.48	0.26	0.73	0.47	0.25	0.74	0.50
Pen-Log	0.27	0.72	0.46	0.27	0.71	0.43	0.26	0.73	0.48
20%	MC	F1	inf	MC	F1	inf	MC	F1	inf
DPL-SVM	0.30	**0.71**	**0.40**	0.30	**0.73**	**0.43**	0.25	0.78	0.54
Pen-SVM	0.30	0.69	0.40	0.31	0.69	0.39	0.26	0.76	0.51
Pen-Log	0.32	0.67	0.36	0.32	0.68	0.37	0.28	0.69	0.44
	Crystallized Intelligence Classification								
10%	MC	F1	inf	MC	F1	inf	MC	F1	inf
DPL-SVM	0.13	0.88	0.75	0.11	**0.91**	0.81	0.09	0.92	0.84
Pen-SVM	0.14	0.86	0.73	0.11	0.89	0.79	0.09	0.91	0.83
Pen-Log	0.15	0.85	0.70	0.12	0.88	0.75	0.10	0.91	0.82
12%	MC	F1	inf	MC	F1	inf	MC	F1	inf
DPL-SVM	0.17	0.85	0.68	0.14	0.86	0.71	0.12	0.89	0.77
Pen-SVM	0.18	0.82	0.68	0.15	0.84	0.70	0.13	0.86	0.72
Pen-Log	0.18	0.81	0.64	0.17	0.82	0.67	0.14	0.84	0.70
15%	MC	F1	inf	MC	F1	inf	MC	F1	inf
DPL-SVM	0.24	0.77	0.52	0.23	0.78	0.55	0.20	0.81	0.61
Pen-SVM	0.24	0.75	0.52	0.24	0.76	0.54	0.21	0.79	0.60
Pen-Log	0.24	0.74	0.51	0.26	0.75	0.52	0.23	0.75	0.51
18%	MC	F1	inf	MC	F1	inf	MC	F1	inf
DPL-SVM	**0.26**	**0.78**	0.53	0.25	**0.77**	0.51	0.23	0.75	0.56
Pen-SVM	0.27	0.73	0.49	0.25	0.74	0.49	0.24	0.74	0.50
Pen-Log	0.28	0.73	0.48	0.27	0.71	0.44	0.26	0.73	0.48
20%	MC	F1	inf	MC	F1	inf	MC	F1	inf
DPL-SVM	0.29	0.74	0.46	0.31	0.71	0.40	0.26	0.73	0.48
Pen-SVM	0.29	0.72	0.45	0.32	0.67	0.36	0.27	0.70	0.46
Pen-Log	0.31	0.69	0.40	0.33	0.65	0.33	0.28	0.69	0.44

Table 3.2: Classification results for HCP data based on static connectivity via partial correlation, using single- and multi-session analysis. Results are reported under varying sub-populations reflect the top and bottom spectrums of intelligence. MC, F1, and Inf refer to the mis-classification rate, the F1 score and the informedness metrics.

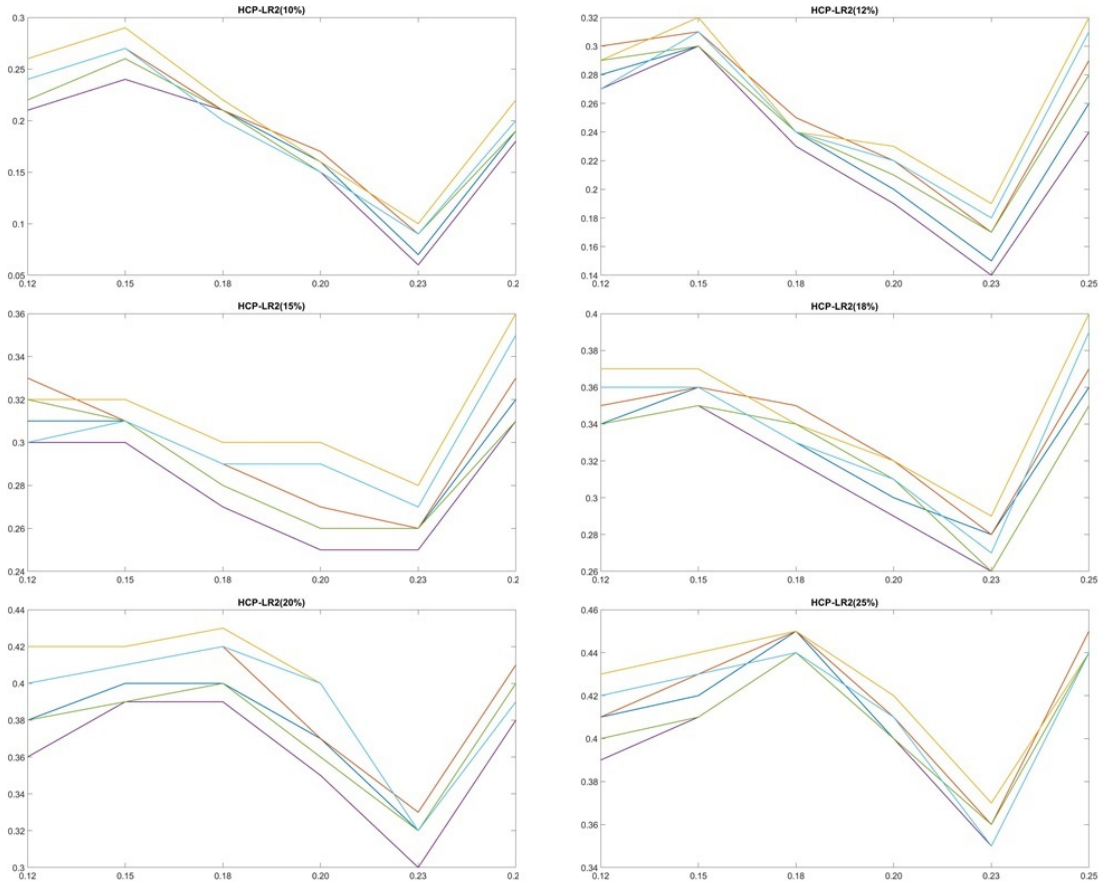

Figure 3.3: Misclassification Rate of fluid intelligence using static functional connectivity based on HCP-LR1 with different percentage of subjects and different network density. Blue line: BNP-SVM (network only), Red line: Penalized SVM (Network only), Yellow line: Penalized Logistic Regression (Network only), Purple line: BNP-SVM (Network+Age+Gender), Green line: Penalized SVM (Network+Age+Gender), Cyan line: Penalized Logistic Regression (Network+Age+Gender)

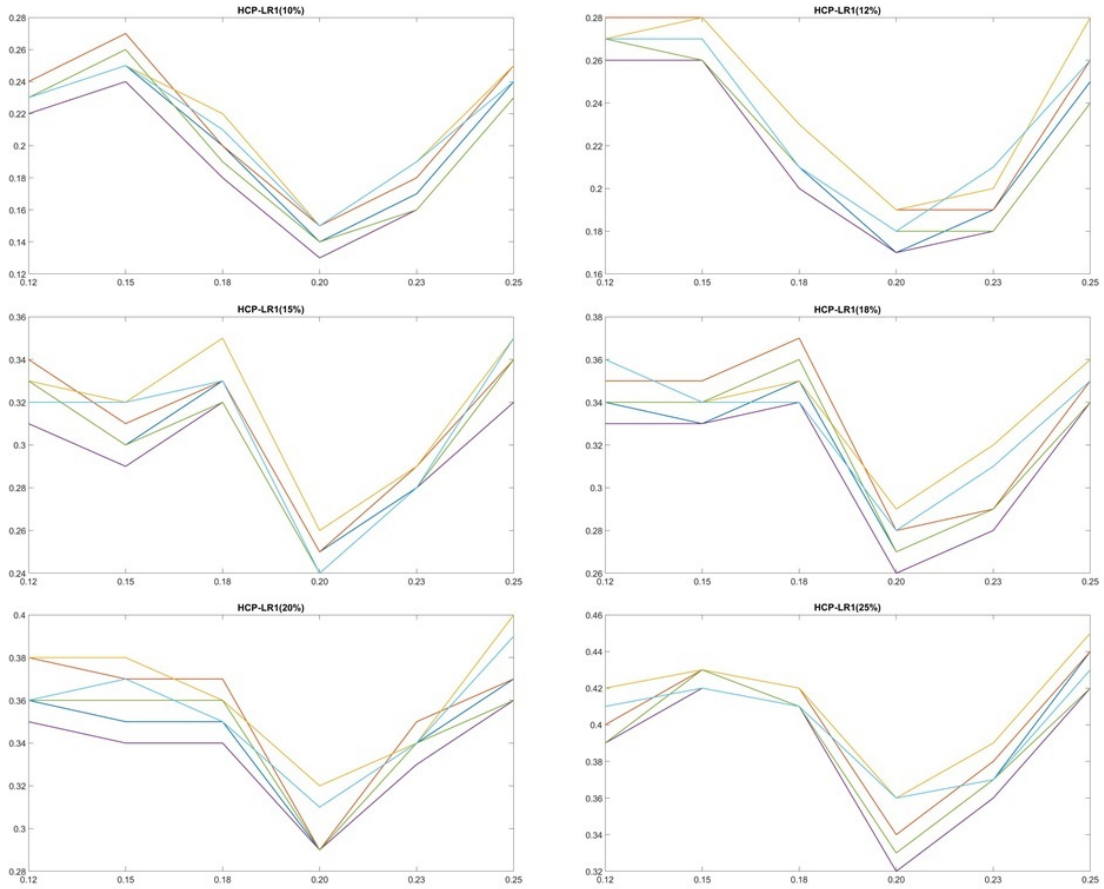

Figure 3.4: Misclassification Rate of crystallized intelligence using static functional connectivity based on HCP-LR1 with different percentage of subjects and different network density. Blue line: BNP-SVM (network only), Red line: Penalized SVM (Network only), Yellow line: Penalized Logistic Regression (Network only), Purple line: BNP-SVM (Network+Age+Gender), Green line: Penalized SVM (Network+Age+Gender), Cyan line: Penalized Logistic Regression (Network+Age+Gender)

Gf Classification	MC	F1	inform	MC	F1	inform
Sliding Window	HCP-LR1			HCP-LR2		
DPL-SVM	0.06	0.95	0.89	0.20	0.80	0.60
Pen-SVM	0.07	0.93	0.87	0.21	0.79	0.59
Pen-Log	0.07	0.92	0.86	0.22	0.78	0.57
PCA	HCP-LR1			HCP-LR2		
DPL-SVM	0.36	0.68	0.33	0.36	0.67	0.31
Pen-SVM	0.37	0.64	0.25	0.39	0.63	0.25
Pen-Log	0.42	0.58	0.19	0.42	0.58	0.19
ICA	HCP-LR1			HCP-LR2		
DPL-SVM	0.33	0.71	0.39	0.35	0.67	0.32
Pen-SVM	0.35	0.65	0.30	0.36	0.63	0.24
Pen-Log	0.41	0.60	0.22	0.43	0.57	0.18

Table 3.3: Classification accuracy for fluid intelligence based on dynamic connectivity using pairwise correlation

Gf Classification	CNI-CC200			CNI-AAL			CNI-HO		
Without	MC	F1	inform	MC	F1	inform	MC	F1	inform
Age and Gender	Network(0.18)			Network(0.14)			Network(0.20)		
DPL-SVM	0.14	0.87	0.72	0.12	0.89	0.78	0.15	0.87	0.74
Pen-SVM	0.16	0.84	0.70	0.13	0.87	0.76	0.17	0.86	0.73
Pen-Log	0.18	0.82	0.64	0.17	0.83	0.67	0.18	0.81	0.64
With Age and Gender	Network(0.18)			Network(0.14)			Network(0.20)		
DPL-SVM	0.14	0.87	0.73	0.12	0.88	0.76	0.14	0.87	0.74
Pen-SVM	0.15	0.85	0.70	0.12	0.87	0.75	0.16	0.86	0.70
Pen-Log	0.17	0.83	0.67	0.17	0.83	0.69	0.17	0.82	0.67

Table 3.4: Classification accuracy for ADHD vs control using using partial correlation with different network density based on CNI data

is complimentary to recent results in (Sen and Parhi, 2020) that shows good prediction accuracy for intelligence based on dynamic connectivity based on task data in the Human Connectome Project. Moreover, the classification accuracy for fluid intelligence is higher under the LR1 scan compared to the LR2 scan, while for crystallized intelligence classification, ...

3.3.4 Analysis Results for CNI Data

Results based on Static Connectivity

According to Table 3, the results for the CNI dataset analysis using static connectivity show more improved mis-specification rates and higher sensitivity under the proposed model compared to alternative methods, although it can result in slightly lower specificity. This is true across all the three brain parcellation schemes, although the network density that provides the optimal classification performance is seen to vary according to the parcellation scheme. This illustrates that the parcellation scheme can have an effect on the manner in which the network features are associated with the clinical outcome.

In addition, the mis-classification rate reported under the proposed approach based on partial correlations are seen to be improved when compared to the results obtained for the same dataset in Schirmer et al (2021), although that paper used a training sample of 100 in contrast to our training sample of 80. We also see improved performance in terms of other metrics such as F1 score and informedness.

Results based on Dynamic Connectivity

As for the HCP analysis, the classification performance under dynamic functional connectivity based on sliding windows appear weaker compared to the classification results using static connectivity obtained via sparse precision matrices ...

3.4 Discussion

In this paper, we proposed one non-parametric Bayesian SVM which combined Dirichlet Process framework to overcome the limitation of traditional linear classifiers such as penalized SVM when applying to brain network. We have also applied the proposed method into two different data set: HCP and ADHD. Results showing that sparse precision matrix corresponding to the static network using resting fMRI provides the best classification performance. The network density resulting in optimal classification performance vary depending several factors, such as percellation, sample size, and binary phenotype. Dynamic connectivity based on sliding window correlation using resting state fMRI has a worse performance compared to static connectivity.

Chapter 4

Estimating Dynamic Connectivity Correlates Of PTSD Resilience Using MultiModal Imaging

4.1 Introduction

Over the last decade, numerous advances have been made in developing neuroimaging biomarkers for mental illnesses that offer tremendous versatility in terms of understanding and targeting pathophysiological mechanisms such as structural decline (e.g. loss in volume, cortical thinning), functional changes (e.g. fMRI hyperactivity, altered network connectivity), white matter decline (e.g. white matter integrity), pathology aggregation (e.g. amyloid levels), and so on. Of these, static functional connectivity has emerged as one of the most promising biomarkers capable of classifying and predicting mental disorders (see Du, et al., 2018 for a recent review). However, it is increasingly recognized that the brain connectivity may not remain constant over time and is likely to exhibit dynamic variations that may be linked to changes in vigilance (Thompson, et al., 2013), arousal (Change, et al., 2013), emotional state (Cribben, et

al., 2012), behavioral performance (Jia, et al., 2014), disease status (Jin, et al., 2017), and so on. Most existing methods use the sliding window approach (Sakoğlu, et al., 2010) to compute dynamic connectivity, and then use it to classify disease phenotypes. While the sliding window is one of the most popular approaches for dynamic connectivity, it has some known drawbacks including the choice of window length (Kundu, et al., 2018). Further, (Laumann, et al., 2017) suggested that correlations measured by resting-state BOLD are relatively stable over short timescales and may not reflect instantaneous changes in cognition, which suggests the need for alternative approaches to estimate dynamic connectivity. An alternative to sliding window analysis is change point estimation methods that allows for unsupervised learning of changes in connectivity states characterized by a collection of discrete jumps with the connectivity being constant between two consecutive jump points (Kundu, et al., 2018). However, single subject change point based approaches for dynamic connectivity needs further improvements in order to develop it as an effective neuroimaging biomarker, which is a central issue that we investigate in this article.

There has been limited literature on disease phenotypic classification using dynamic FC (Du, et al., 2018) that highlights some initial potential as a neuroimaging biomarker. In particular, dynamic network differences have been discovered between PTSD and non-PTSD groups (Jin, et al., 2017; Fu, et al., 2019). However, disease phenotype classification, while a worthy objective, may be of secondary interest in heterogeneous mental disorders where there is no gold standard for classification, such as our motivating PTSD application. For such disorders, experts may prefer modeling continuous measures of disease progression as a clinical outcome of interest. Unfortunately, there have been limited efforts for using dynamic functional connectivity to model continuous clinical measures (Haslam, 2003; Widiger, et al., 2005). Hence, considerable efforts needs to be made in this direction. Some challenges include: (a) developing reliable measures of dynamic connectivity based on single subject data;

(b) discovering localized brain regions whose dynamic connectivity signatures are directly related to behavior; (c) developing predictive models for continuous clinical outcomes based on dynamic networks, that goes beyond disease phenotype classification; and (d) developing quantifiable dynamic network summaries that provide the highest differentiating power with respect to the clinical outcomes of interest. For our PTSD application, the continuous clinical outcome needs to be chosen carefully so that it accounts for varying levels of trauma exposure that represents a major source of heterogeneity for individuals with and without PTSD.

In this article, we seek to address the above issues using data on PTSD symptom severity among individuals obtained from the Grady Trauma Project. Our results show that network fluctuations in terms of the frequency of switching between different states is weakly related to the disease phenotype, but the fluctuations in the connectivity strength is related to a continuous metric of disease severity called resilience that is defined as the ability to maintain high levels of emotional and social functioning in the face of stress exposure. To address (a), we develop a novel approach to estimate dynamic FC based on resting state fMRI (Rs-fMRI) data that is guided by brain structural connectivity (SC) information obtained via diffusion tensor imaging (DTI) data. The estimates of the dynamic network obtained via the multimodal approach is shown to be more sensitive to network changes and provides greater predictive power when modeling trauma resilience. The multimodal dynamic FC approach is also shown to recover the true dynamic network accurately via extensive validation studies. Our analysis helps discover edges or connections between brain regions whose variations in connectivity strengths are directly related to trauma resilience, as well as edges that are more stable in the dynamic network, which provides insights into question (b). In order to address questions (c) and (d) we apply state of the art scalar-on-function regression approaches to model trauma resilience based on dynamic connectivity, and identify time-varying global network summaries

encoding patterns of information transmission in the brain that provide the highest predictive accuracy under the proposed multimodal approach compared to static network models as well as dynamic network models using fMRI only. We also identify functional modules in the network where the local patterns of dynamic information transmission are most predictive of trauma resilience with superior accuracy under the proposed approach. A graph theoretic approach is used to model the network, where brain regions are perceived as nodes and edges describe functional or structural connectivity depending on the imaging modality.

To our knowledge, we are one of the first to develop and investigate multimodal dynamic FC as a neuroimaging biomarker for predictive analysis of disease severity in mental illness. Of course, the motivation for structurally guided dynamic functional connectivity comes from a well-established literature illustrating the relationship between static FC and SC. In particular, there is strong evidence that white matter fiber tracts regulate static FC (Damoiseaux et al., 2009; Sporns, 2013). Based on such evidence, there has been some limited development of static FC approaches guided by SC knowledge (Hinne, et al., 2014; Messé, et al., 2014) using multi-subject data as well as data for single individuals (Higgins, et al., 2018). Since static FC can be interpreted as an average of dynamic FC values over time, it is natural to conjecture that dynamic FC is also regulated by brain SC to some degree. Hence, it is of interest to develop data-adaptive approaches to compute multimodal dynamic FC, especially to improve dynamic network estimates for single subject analysis. Since it is non-trivial to generalize the existing approaches to our settings of interest involving the estimation of dynamic brain networks guided by SC information, our efforts represent major methodological and scientific innovations with the ultimate goal of developing neuroimaging biomarkers based on dynamic FC.

4.2 Materials and Methods

4.2.1 Description of Grady Trauma Project Data

Our study involves female African-American participants from the Grady Trauma Project (GTP). These participants were recruited from primary care clinics at Grady Health System, a publicly funded, tertiary care center serving a predominantly socioeconomically disadvantaged inner-city population in Atlanta, Georgia. A majority of these participants have experienced significant psychological trauma of various types (Gillespie, et al., 2009). Imaging modalities including Rs-fMRI and DTI data were collected for each individual - the details for the pre-processing steps are provided in the Supplementary Materials. We focused on a moderate subset of 24 individuals who were aged below 30 years and hence we expected to have more homogeneous brain function and structure. The participants were all female and African-American (age 25.8 ± 3.1) and did not have any disability. Out of these participants 1 was Hispanic and 23 were non-Hispanic, 12 subjects were diagnosed as PTSD (resilience score, 13.1 ± 8.2) and the other 12 subjects were not (resilience score, -6.9 ± 5.3). Our clinical outcome of interest is the resilience score that is measured via the Connor-Davidson Resilience Scale (CD-RISC) (Connor, et al., 2003). Resilience is a transdiagnostic indicator of mental health in the face of adversity, and is highly relevant to groups of individuals who have experienced high levels of trauma exposure and other forms of stress, as in the Grady Trauma Project sample. High resilience scores can be interpreted as a direct indicator of PTSD severity, and all individuals with high resilience had PTSD positive status. We note that although our sample size is moderate, we expect to obtain accurate inferences under the proposed approach that involves the integration of functional and structural imaging to obtain the dynamic FC, which is subsequently used to model resilience using state-of-the art statistical techniques. Moreover, using the continuous resilience score as our clinical outcome

of interest allows us to pool information across varying levels of disease severity to inform our analyses, which provides greater power to detect important signals compared to an alternate analysis that compares distinct clinical phenotypes represented by smaller subgroups.

Power Atlas and Functional Modules

We use a whole brain parcellation corresponding to the fMRI time courses from the 264 ROIs under the Power system (Power, et al., 2011) to perform our network analysis. Although the network analysis involved 264 nodes, we further group these ROIs into 10 functional modules as identified by (Cole, et al., 2013), which better characterize resting state functional networks. These modules included sensory/somatomotor, cingulo-opercular (CON), salience (SAL), auditory (AUD), subcortical (SCOR), default mode network (DMN), visual, fronto-parietal (FPL), ventral-attention network (VAN), and dorsal-attention network (DAN). The coordinates for the ROIs and their allocation to these modules in presented in a Table in Supplementary Materials. We note that 37 ROIs were excluded from our analysis based on (Cole, et al., 2013), either because they were located in cerebellum or they were identified with unknown functionality.

4.2.2 Validation Studies Using Simulated Data

We conducted extensive simulation studies to evaluate the performance of the proposed approach, under different network structures. In the first set of simulations (Scenario I), we generated data from an underlying change point model with three change points (i.e. four state phases). A network and the corresponding precision matrix was constructed at each time point (described below), and these were constant within each state phase. The measurements were generated under a Gaussian distribution characterized by the time-varying precision matrix. We generated data for

$V = 20, 50$ regions and with $T = 300, 500$ time points. In a second set of simulations (Scenario II), we allow the network to change more slowly over time, that is more consistent with the timescale of the haemodynamic activity. In particular, instead of three change points as in the first scenario, we now have three transition periods, each comprising seven consecutive time points. A certain percentage of the edges are flipped from the network at the previous time point to obtain the modified network for the next time point within each transition period. The network is assumed to be constant between two consecutive transition periods. This scenario is more challenging since the network changes multiple times over the course of the experiment. The goal of this experiment is to investigate if the proposed approach can detect the transition periods and whether it can approximate the true dynamic network sufficiently well when the underlying assumptions of the proposed model may not hold.

We generate data using three different types of networks (details provided in Supplementary Materials). In addition to Scenarios I and II, we also reported results under Erdos-Renyi networks for 100 nodes with three true change points, to test the performance in higher dimensions, and investigated the scenario involving a large number of change points (10) with 50 nodes. These challenging settings help us evaluate the performance of the methods for high dimensions and large number of fluctuations in the dynamic network. For each simulation setting, 25 replicates were used.

4.2.3 Overview of Statistical Approach

Our goal is to develop an approach for estimating dynamic connectivity based on Gaussian graphical models that guided by SC knowledge. We denote our approach as multimodal dynamic FC (mDFC) through the rest of the paper. GGMs assume that the fMRI measurements are normally distributed and are characterized by a sparse inverse covariance or precision matrix that has zero off-diagonals corresponding to

absent edges in the network. Moreover, the non-zero elements of the precision matrix encode the strength of the important edges. The dynamic GGM approach proposed in this paper combines the concept of change point estimation with GGMs, where the fMRI measurements are assumed to come from a Gaussian distribution having a sparse inverse covariance matrix that is constant within a state phase, but changes between state phases. Each state phase denotes the period between two consecutive change points, whose number and locations are unknown. We note that in a dynamic GGM, the pattern of zeros in the precision matrix at each time point essentially provides all the necessary information about the time-varying network.

Denote $\mathbf{y}_t$ as the $V \times 1$ vector of spatially distributed fMRI measurements over V voxels or regions of interest (ROI), at the t-th time point ($t = 1, \ldots, T$). Denote the SC probability corresponding to the edge (j, l) as p_{jl} where $j \neq l, j, l = 1, \ldots, V$, and denote the corresponding SC probability matrix as $\mathcal{P}$. These SC probabilities are obtained from DTI data, and are made symmetric (i.e. $p_{jl} = p_{lj}$) as in (Higgins, et al., 2018). We specify the following dynamic GGM

$$\mathbf{y}_t \sim N(_t, \Omega^{-1}_{\mathcal{P},G_t}),\ t = 1, \ldots, T, \tag{4.1}$$

where $N(, \Omega^{-1})$ refers to a multivariate Gaussian distribution with mean and covariance matrix Ω^{-1}, $\Omega_{\mathcal{P},G_t}$ denotes the inverse covariance or precision matrix at time point t that depends on the time-varying network G_t characterized by the vertex set $\mathcal{V}$ and edge set E_t, as well as brain SC information $\mathcal{P}$. The vertex set $\mathcal{V} = \{1, \ldots, V\}$ consists of a set of pre-defined voxels/ROIs or nodes, the edge set E_t contains the set of all edges present in G_t, and $\{\Omega_{\mathcal{P},G_1}, \ldots \Omega_{\mathcal{P},G_T}\}$ encodes the strength of the time varying FC. The goal of change point modeling is to develop an algorithm that is able to learn the best fitting partition of the time course defined by the change points $0 = a_1 < a_2 < \ldots < a_K < a_{K+1} = T$ so that G_t remains constant for consecutive

time points except for discrete jumps at the change points. These change points are unknown for our problems of interest, and estimated in an unsupervised and data-adaptive manner. For conciseness, we denote $\Omega_{\mathcal{P},k}$ as the constant precision matrix for the k-th state phase corresponding to the interval $(a_{k-1}, a_k]$. throughout the article. A diagrammatic illustration for our approach is provided in Figure 4.1. More details regarding the approach including the structurally informed precision matrix estimation, and change point estimation using the subnetwork sampling scheme are provided in the Supplementary Materials.

4.2.4 Analysis Outline

Using the mDFC approach, we first compute the dynamic resting state network and subsequently investigate whether or not the temporal changes in the dynamic network can explain variations in disease severity measured via resilience. Figure 4.2 provides a visual depiction of the different steps in our analysis, which are described in detail below.

Detecting Temporal Connectivity Fluctuations Related to Resilience

We would like to investigate whether temporal variability is related to disease severity. To evaluate this, we (a) examined whether temporal variability of an edge in terms of the number of fluctuations between the present and absent states (measure via change points) is significantly different between individuals with and without PTSD (or equivalently high and low resilience groups); and (b) investigated whether the temporal variability of the connectivity strength (measured via partial correlations) is related to the continuous resilience measure. For (b), we performed a regression analysis to identify those edges whose standard deviation of the partial correlations are significantly related to the resilience score. Due to a large number of edges involved, an univariate analysis was performed for one edge at a time, and significant effects

were identified after multiplicity adjustment for p-values using Benjamini-Hochberg corrections (Benjamini, et al., 1995). Edges having a significantly positive association under this analysis will imply connections where greater variability in connectivity strength enhances resilience and vice-versa, whereas significant negative associations will imply connections where greater temporal fluctuations lead to decrease in resilience and vice-versa.

Prediction Based on Dynamic FC via Scalar on Function Regression

We use scalar on function regression to predict the resilience score using the dynamic functional connections as a function of time. The scalar-on-function regression (Ramsay, et al., 2005) can be expressed as $z_i = \int_0^T \hat{\rho}_i(t)\beta(t)dt + \epsilon_i$, where ϵ_i denotes the residual that is assigned a Gaussian distribution, z_i represents the scalar clinical outcome and $\rho_i(t)$ denotes some network summary measure of the estimated time varying functional connectivity for the i-th individual derived under the dynamic network (see the sequel for more details), and $\beta(t)$ denotes the time-varying coefficient function that weights the dynamic connection over time in order to model the outcome and can be interpreted as a dynamic analogue of regression coefficients in usual linear regression models. The model fitting proceeds by expressing the time varying quantities using known basis functions and unknown wavelet coefficients, and thereafter using standard techniques in linear regression models to estimate the unknown wavelet coefficients and thereby estimate the time varying regression coefficient.

We used the R package FDBoost (Brockhaus, et al., 2018) for implementing the scalar-on-function regression. The predictive accuracy of the scalar-on-function regression was assessed using out of sample mean squared error that calculates the averaged squared difference between the observed and predictive values in the test sample. A lower error implies higher predictive accuracy and a better model fit. We randomly split the sample into training and test sets, fitted the scalar on regression

model on the training data, and computed the out of sample MSE on the test sample. We then repeated the above operation for 100 random splits and obtained a boxplot of the our of sample MSE. We did not include gender, race or age in our regression model since the entire sample comprised African-American females between 19-30 years.

Dynamic Network Summaries for Prediction of Disease Severity

Edge level analyses, although more easily interpretable, may often be subject to greater levels of noise and may be less reproducible across studies. Hence, instead of using edge level features, we investigated the predictive ability of global dynamic network summaries such as small-worldedness, global efficiency and global clustering coefficients, as well as local clustering coefficient and local efficiency corresponding to some local functional modules. These include Visual, Salience, Subcortical, VAN, and DAN modules that were identified as regions with the highest FC changes for trauma exposed individuals under our dynamic network analysis (see Figure 4.4). We also reported the prediction results corresponding to the small-worldedness for all modules. Clustering coefficient and small worldedness were chosen based on recent findings of differences in these metrics in static networks between individuals with and without PTSD (Rowland, et al., 2018), whereas the global and local efficiency are additional network metrics that we chose to investigate in the context of predicting resilience. We note that all the network metrics change across state phases under the dynamic network, and hence are time-dependent. The network metrics represent the efficiency of information transmission at a global or local level in the brain, and were computed at each time point and for each subject using the Matlab toolbox Brain Connectivity Toolbox (Rubinov, et al., 2010). They are described in detail in Supplementary Materials.

Comparison with Existing Methods

We compare our method with two competing approaches: a hierarchical Bayesian structurally informed Gaussian graphical model (siGGM) by (Higgins, et al., 2018) which estimates the static network over the experimental time course while accounting for the SC information, and dynamic connectivity regression (DCR) for single participants proposed by (Cribben, et al., 2013) in validation studies. For the analysis of the data from the PTSD study we replace DCR with the SC naive version of the proposed approach that is similar to DCR, since the DCR approach is not applicable to high-dimensional networks used in our analysis involving 264 nodes or ROIs. We used default set-up in the (Cribben, et al., 2012) paper to implement the DCR method for validation studies, i.e. minimum block size of 30, significance level for permutation test of 0.05 and the number of bootstrap is set to be 1000. While both the change point detection and the network estimation performance were reported under the DCR approach and the SC naive version of the proposed method, the siGGM only reports the network estimation performance, since it does not account for dynamic changes. The performance metrics used for comparison are described in detail in Supplementary Materials.

4.3 Results

4.3.1 Findings in PTSD Data Analysis

Multimodal Approach is More Sensitive to Dynamic Network Changes

Our first aim is to compute dynamic FC guided by brain SC information for each participant, and investigate brain regions that exhibit connections with the greatest and least temporal variability. We calculated networks which were approximately around 15 percent density for all participants, which seem to reflect an acceptable

sparsity level in brain connectivity studies. Figure 4.3 provides the histogram for the number of change points for the proposed approach and the SC naive version of the method. Our method detected 5 change points on average across all participants, with the number of change points ranging from 3 to 7. On the other hand, the SC naive version of the method registers only 1 change point for a large majority of participants, and only one subject has 3 change points. Given 146 brain volumes in the fMRI time series and recent findings that some brain networks may change within as little as 30–60s (Sakoğlu, et al., 2010; Shirer, et al., 2012), the number of change points under the SC naive version seem to be unrealistic whereas the number of change points under the proposed method appears more practical and supported by previous evidence. Our findings reveal that the brain network computed via the proposed multimodal FC method using brain SC information is more sensitive to temporal fluctuations in the dynamic network to a far greater degree than the SC naive version. Hence incorporating brain SC knowledge provides greater power to detect dynamic changes in the network (also see validation experiments).

The edges with the largest temporal fluctuations are identified as those that consistently switch over the different state phases averaged across individuals, as reported in Figures 4.4(A)-(C). In particular, this Figure depicts the proportion of times each edge flips (changes from present to absent state from one time point to the next, and vice-versa) over the course of the scanning session averaged over all individuals (Figure 4.4(A)), and separately for the PTSD (Figure 4.4(B)) and non-PTSD (Figure 4.4(C)) cohorts, where a higher proportion implies greater temporal variability. From Figure 4.4 (A), we observe that connectivity within and between certain modules such as the Visual, FPL, DAN, and VAN, exhibit the highest temporal fluctuations in connectivity. Interestingly, the highest temporal fluctuations in connectivity occur between the different functional networks, whereas the temporal fluctuations within most of these functional modules is somewhat limited. On the other hand, edges connected to nodes

in Sensory, CON, AUD, and DMN have the least temporal variability both within and between modules, which implies stable connections that can be characterized by static connectivity. Hence, our approach identifies certain concentrated regions in the brain with greatest temporal fluctuations in connectivity, and also identifies several other regions can exhibit static connectivity in terms having negligible frequency of state phase changes.

We observed that connections between the Visual-FPL modules, Visual- VAN module, Visual- DAN module, Visual- SAL network module, and the Visual - SCOR regions have the highest temporal fluctuations in the dynamic network. The Visual network seems be most active in terms of temporal fluctuations - some of the regions that had the highest fluctuations in the Visual network were located in left Occipital Inf, left Lingual, right inferior temporal gyrus, right Cuneus, and left middle Occipital lobes. These regions have been shown to be involved in upregulated perception of environmental stimuli that could be arousal-mediated (Mueller-Pfeiffer, et al., 2013). The SAL and SCOR modules have previously demonstrated altered resting state connectivity (Rabinak, et al., 2011; Brown, et al., 2014), and grey matter alterations were also discovered in subcortical areas for PTSD individuals (O'Doherty, et al., 2017). The highest temporal fluctuations in the Salience network came from regions located in the middle frontal gyrus and superior frontal gyrus (Brodmann Area 10) that have been known to be associated with cognitive control (Nicholson, et al., 2017). Moreover, resting state connectivity differences in prefrontal cortex (Kennis, et al., 2015) and hypoactivity in medial prefrontal cortex (Koenigs, et al., 2009) have been observed in individuals with varying PTSD severity. Additionally, the areas with the highest temporal network fluctuations in the subcortical module were located in the Putamen region close to the white matter just lateral of caudate, that is known to be associated with motivated behavior/approach (Elman, et al., 2009). Finally, increased resting state FC has also been discovered in DAN and VAN modules (Block,

et al., 2017).

Do Temporal Network Fluctuations Drive Disease Severity?

To investigate this question, we first compared the temporal fluctuations for individuals with and without PTSD in terms of the frequency of state changes in the network in Figures 4.4 (A)-(C). Figures 4.4 (B)-(C) clearly illustrate very limited differences in temporal variability patterns in individuals with and without PTSD, although individuals with PTSD tend to have slightly higher number of fluctuations in brain network state phases. In particular, only 213 edges (0.61% of all possible edges in the network) have significantly different frequency of state phase changes between the PTSD and non-PTSD groups at 5% level of significance after adjusting for family-wise error rate over all edges using Bon-ferroni corrections. Almost all of these edges correspond to a higher frequency of changes under the PTSD group. Overall, this part of our analysis point to the limited ability of the frequency of state phase changes to distinguish between the PTSD and non-PTSD groups. Hence it is imperative to develop alternate measures of dynamic connectivity that provides a greater distinction with respect to disease severity.

We seek to develop a measure for temporal changes in connectivity that is more effective in terms of showing greater associations with disease severity. Hence, we investigated the ability of temporal variability of edge-wise connectivity strengths (measured as the standard deviation of partial correlations) to predict resilience scores that is a more continuous representation of the disease spectrum. Figure 4.4(D) illustrates those dynamic FC that have significant associations between temporal variability in connectivity strength and resilience. The significant edges were detected using univariate analysis at 5% level of significance and using multiplicity corrections for controlling family-wise error rate over all edges via Bon-ferroni corrections. The overwhelming number of edges between the Visual and Sensory functional modules have

fluctuations in connectivity strength which are significantly related to resilience. The temporal fluctuations of connectivity strengths between and within the Visual, FPL, VAN and DAN modules are also see to be significantly associated with resilience. The functional modules with the smallest proportion of edges whose temporal variability is related to resilience include the Salience and Subcortical modules, and the DMN. Clearly, our findings regarding significant dynamic network differences driving disease severity is stronger and more informative compared to the findings in (Fu, et al., 2019), who discovered some preliminary dynamic FC differences between PTSD individuals and healthy controls, but were unable to recover dynamic FC features that are significantly associated with PTSD status.

The overwhelming majority the significant associations (96.1%) are negative, which means increased temporal fluctuations for these connections lead to decreased resilience, and vice-versa. A small number of positive associations between the temporal variability of the connection strengths and resilience can be found for edges corresponding to nodes within the DMN, and between DMN and the Visual module and FPL modules. Since the connections in DMN were shown to relatively stable over time (Figures 4.4 (A)-(C)), significant positive associations between resilience and edges within the DMN implies lower resilience resulting from stable connections within DMN. Similarly, the increased temporal connectivity fluctuations within and between the Visual, FPL, VAN and DAN modules (Figures 4.4 (A)-(C)), combined with the largely negative associations between temporal variations of connectivity strengths and resilience for these edges implies a decrease of resilience due to increased dynamic connectivity in these regions. The same conclusion also holds for edges between the Visual and Salience modules, and those between the Visual and Subcortical modules. On the other hand, negligible temporal fluctuations in other regions of the brain (Figures 4.4 (A)-(C)), coupled with largely negative associations between edges in these regions and resilience (Figure 4.4 (D)), point to an increase in

resilience due to more stable associations in these regions.

Dynamic Networks Provide Higher Predictive Accuracy for Trauma Resilience

Another important objective is to investigate the ability of the dynamic FC to differentiate individuals with distinct resilience levels and check whether dynamic networks estimated by fusing functional and structural imaging data can provide significantly greater predictive power compared to standard network modeling approaches. The results of our predictive analysis are reported via boxplots for out of sample MSE corresponding to the different network metrics are presented in Figure 4.5. An additional Figure (4.6) also demonstrates the predictive accuracy using small-worldedness derived from different localized functional modules.

A multiplicity adjusted permutation test revealed that the improvements in predictive accuracy were significantly better under mDFC when using dynamic global efficiency, clustering coefficient, and small-worldedness. It is also clear that when focusing on localized functional modules, the proposed approach has a lower MSE compared to the SC naive dynamic connectivity for most cases, and a comparable MSE for the remaining cases. Moreover, the MSE under mDFC derived from localized functional modules is significantly lower compared to siGGM that computes static networks incorporating brain SC knowledge, which clearly highlight the predictive advantages of dynamic FC over static FC in PTSD. We note that the predictive MSE using small-worldedness from the SAL network is significantly lower compared to the other two methods, and lower or comparable for the other localized functional modules. Further, the mDFC approach shows significant improvements in predictive accuracy when using the following local time varying network features: the dynamic clustering coefficient for nodes in DAN and VIS modules, as well as those in SAL and VIS modules, nodes in VAN and VIS modules, and finally nodes in the SCOR

and VIS modules. Superior predictive accuracy under mDFC is also obtained when using dynamic efficiency computed using nodes in the DAN module, as well as nodes combined from the DAN and VIS modules, with the predicting accuracy under the latter being significantly higher. This could potentially imply an important role of between module connections for predicting resilience. These findings may indicate modules in the brain where dynamic connectivity changes have the greatest differentiating power with respect to resilience, and help localize the brain regions where the temporal network fluctuations have the strongest predictive power with respect to resilience.

We note that our findings go further than the results in (Rowland, et al., 2018) who illustrated differences between individuals with and without PTSD based on small-worldedness and global clustering coefficient derived from static networks. Our investigation not only reveals the importance of global small-worldedness, clustering coefficients, and efficiency computed using dynamic networks for modeling disease severity, but also reveals more localized functional modules whose clustering coefficient and efficiency based on the dynamic network are related to resilience.

4.3.2 Results from Validation Studies

Figures 4.7, 4.8 and 4.9 report the results for the validation studies corresponding to Scenarios I and II. We note that we could not report results under DCR for the 100 node scenario since the DCR is not scalable to high dimensional networks.

Multimodal Dynamic FC Provides Near Perfect Estimation for Network Changes

Figure 4.7(a)-(b) provides a visual illustration of the change point detection performance under the first scenario, which plots the frequency with which each time point is detected as a change point under the sub-network sampling scheme. In this Figure,

the peaks in frequency are seen to concentrate around the true jumps which illustrate the ability to accurately detect change points. The results in Figures 4.8 and 4.9 suggest that the proposed approach is able to detect essentially almost all the true change points under both scenarios, as computed via the proportion of change points detected. Moreover, the proportion of false change points detected (computed as the number of falsely detected change points over replicates) is close to 0 or negligible. In contrast, the detection of false change points is much higher under DCR, and it has poor performance in terms of detecting the true change points. In fact, the multimodal dynamic FC approach has a significantly higher proportion of change points detected, and significantly lower number of false positives. Additional experiments (not presented here) reveal that the performance of DCR improves when the total number of time points in the experiment, along with the distance between consecutive change points is increased. However, for practical experiments with a few hundred time points, the DCR approach seems to fail in terms of change point estimation. The proposed approach, which incorporates SC knowledge and espouses a novel sub-network sampling scheme, performs considerably better in terms of detecting the true change points while incurring minimal false positives.

Figure 4.7(b) presents the change point detection results for the high-dimensional case of $V = 100$ for data simulated under the Erdos-Renyi network with three jump points. The Figure clearly depicts high frequencies for change point detection around the true change points, thereby suggesting that the proposed approach could successfully detect true change points based on the sub-network sampling mechanism. Finally, the results for the case involving a higher number of true change points (10) with $V = 50$ regions and $T = 500$ time points, is presented in Figure 4.7(c), which clearly shows peaks under the sub-network sampling scheme around all of the 10 true change points, thereby indicating the power of the proposed approach in detecting discrete jumps. In order to accurately detect all the true change points, we increased the

number of sub-network samples to 100 for this case. For this case, the DCR approach fails to detect an overwhelming majority of the jumps (results not reported).

In addition, we assessed the performance of the proposed approach when both the number of change points and ROIs increases. Simulation results (not presented here) shows that as long as the number of sub-networks is large enough and true change points are not exceedingly close together, our method successfully detects peaks around the true change points, under the sub-network sampling scheme. The computation time for each sub-network is reasonable as long as the number of nodes in the sub-network is moderate - see Figure 4.7(d) for more details. We conjecture that an increasing number of sub-networks will be required for a good performance as the number of nodes in the network, as well as the number of true change points is increased. In summary, using a combination of a powerful sub-sampling scheme and incorporating prior SC knowledge, the proposed method is shown to provide vast improvements over existing approaches in literature in terms of detecting underlying true network changes.

Multimodal Dynamic FC Results in Higher Network Estimation Accuracy

The results of graph estimation are reported in Figures 4.8 and 4.9. Compared to the competing methods, the proposed method consistently has a significantly higher area under the Receiver Operating Characteristic Curve (AUC) value for network estimation. In addition, it has a higher or comparable sensitivity and consistently has a significantly higher specificity, implying lower false positives in graph estimation and suitable power to detect true positives. In several cases, both the sensitivity and specificity under the multimodal dynamic FC were higher compared to the other two methods. A higher AUC, along with higher specificity and comparable sensitivity, illustrates the ability of the multimodal dynamic FC approach to better control for false positives while having a similar or higher power to detect true signals. We

note that the accurate estimation of the dynamic network under the multimodal FC approach illustrates the importance of using brain SC to guide the estimation of dynamic FC.

4.4 Discussion

In the current study we developed a novel method (mDFC) for estimating dynamic changes in fMRI resting state functional connectivity guided by SC information in the brain. The approach can be scaled to a large number of change points as well as nodes, and is applicable to diverse settings. Our findings suggest that (1) estimating dynamic network connectivity models can be improved with the addition of DTI-based structural constraints; and (2) including metrics of dynamic change in resting networks will improve models for predicting psychiatric risk and resilience to trauma and stress. For (1), we see that SC guided dynamic FC estimation results in higher sensitivity in detecting network changes, which suggests that dynamic FC changes are potentially correlated with brain SC. One possible explanation for the association between SC and network changes is that the dynamic FC can be considered as distinct manifestations of an underlying intrinsic network that is associated with the brain SC, in a manner that is similar to the association between static FC and SC. Additional work is needed to investigate the above conjecture; but if true, then this could potentially be a novel finding with considerable implications. For (2), our PTSD data analysis discovers resting state network alterations among participants exposed to varying degrees of trauma based on dynamic connectivity, illustrates the direct link between temporal fluctuations of connections with PTSD resilience which implies that the dynamic network is related to mental health. The ability of small-worldedness, clustering coefficients and efficiency computed from the dynamic network to accurately predict trauma resilience points to the potential of these dynamic

network metrics as neuroimaging biomarkers in trauma resilience studies. Moreover, the localized brain regions identified as directly related to resilience may have strong clinical interpretations in terms of PTSD diagnosis and treatment. In particular, the local clustering coefficient and local efficiency computed using nodes in the Visual and Dorsal attention network both result in highest improvements in predictive accuracy, which points to the importance of these regions with respect to trauma resilience in a civilian, highly traumatized sample of African American women. These findings stand in contrast to current neural circuit-based models of trauma resilience. Previous models of the neural correlates of resilience have often focused on populations of treatment-seeking patients with either depression or PTSD, whereas in the current study we recruited broadly among a sample at high risk for trauma but without respect to any mental health diagnosis or complaint. Furthermore, while the majority of resting state network modeling in psychiatry has focused on static functional connectivity, our novel analyses that are highly sensitive to network dynamics allowed us to provide a model of the neural correlates of resilience with much lower error than a static connectivity-based comparison model (refer to Figures 4.5- 4.6).

These findings build upon traditional neural circuit models of trauma-related psychopathology and resilience, which have focused primarily on individual differences in fear learning and extinction, supported by plasticity within an amygdala-mPFC-hippocampal circuit (Liberzon, et al., 2007; Johnson, et al., 2012). Recent findings suggest that stress resilience also depends on the instantiation of fear memories throughout a broader network of regions, including the primary sensory cortex through which initial information about the trauma or threat was gathered (Ressler, 2020) Network-based resting state functional connectivity findings have supported this idea in PTSD. For example, visual and sensorimotor networks were implicated in identifying subtypes of PTSD among war-exposed male military veterans (Maron-Katz, et al., 2019) . A subgroup with greater re-experiencing symptoms showed

hyper-connectivity between visual and sensorimotor networks, and hypo-connectivity between sensorimotor network and prefrontal regions. Similarly, in our dynamic connectivity analyses, we identified that greater temporal fluctuation in the visual and sensorimotor networks were negatively associated with resilience in the context of a sample of female civilian trauma survivors with histories of multiple trauma spread over the lifespan, supporting the importance of sensory resting state nodes to trauma-related pathology across multiple populations. Furthermore, these results find parallels in preclinical rodent models of fear learning in which fear-related behavior to auditory tones that had previously been associated with footshocks was dependent on plasticity in auditory cortex in addition to associations formed in the amygdala (Keifer, et al., 2015; Banerjee, et al., 2017; Lai, et al., 2018). In humans exposed to trauma, plasticity in the visual and sensorimotor cortices, and their interplay with the salience network, which is hypothesized to direct sensory attention resources to salient environmental stimuli (Menon, et al., 2010), may critically determine how trauma memories are encoded and stabilized over the long term. Further work linking resting-state network dynamics with trauma and stress-related psychopathology is likely to provide new mechanistic insights about the neural processes contributing to psychiatric resilience and vulnerability following a major stressor. The current findings may be consistent with an "overconsolidation" hypothesis (Pitman, 1989; Yehuda, et al., 2007) such that individuals at risk for PTSD or other forms of psychopathology following trauma may have very efficient communication between various aspects of the network supporting fear memory, and thus encode and stabilize a very strong fear memory when exposed to trauma. However, currently, there is very little theoretical work to explain how network dynamics might contribute to individual differences in behavior or psychiatric symptoms. This is an important area of development for future efforts in neural circuit modeling in psychiatry.

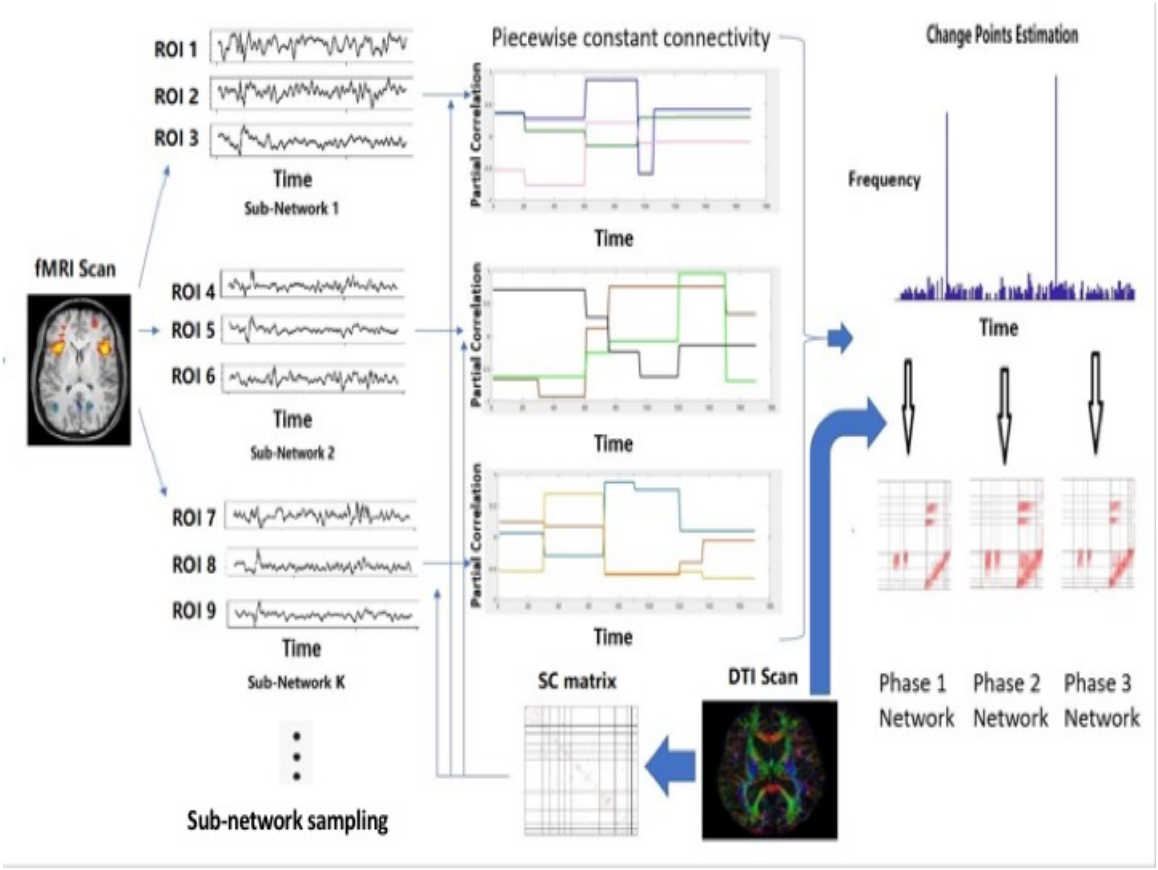

Figure 4.1: A diagrammatic illustration of our novel multimodal dynamic FC approach using Rs-fMRI data that is guided by brain SC information computed from DTI data. Given a set of nodes in the network, the approach is able to learn change points or jumps in the network in an unsupervised manner, where the number and locations of the change points are unknown and the network is assumed to remain constant within a state phase defined as the time interval between two consecutive change points. The greedy partitioning scheme used to compute change points uses state phase-specific networks that are computed after incorporating brain SC knowledge - in this manner, the change point estimation procedure is influenced by the given brain SC information. In order to scale up the multimodal dynamic FC (mDFC) approach to high-dimensional networks, we propose a sub-network sampling scheme where we use the mDFC approach to compute change points using several smaller subsets of nodes or sub-networks. This process is applied repeatedly for a large number of sub-networks, and the set of change points for each sub-network is recorded. The sub-network sampling scheme yields a frequency or score for each time point to be identified as a network level change point, and a systematic data-adaptive thresholding strategy to determine frequency cut-offs that can be used to determine network level change points that are consistently identified across most sub-networks. Conditional on the estimated network level change points, the structurally informed precision matrix estimation is applied once again to compute a distinct sparse inverse covariance matrix encoding the network separately for each state phase. The state phase specific networks are computed by integrating brain SC information that encourages greater weights for FC corresponding to those edges with strong SC under a Gaussian graphical model.

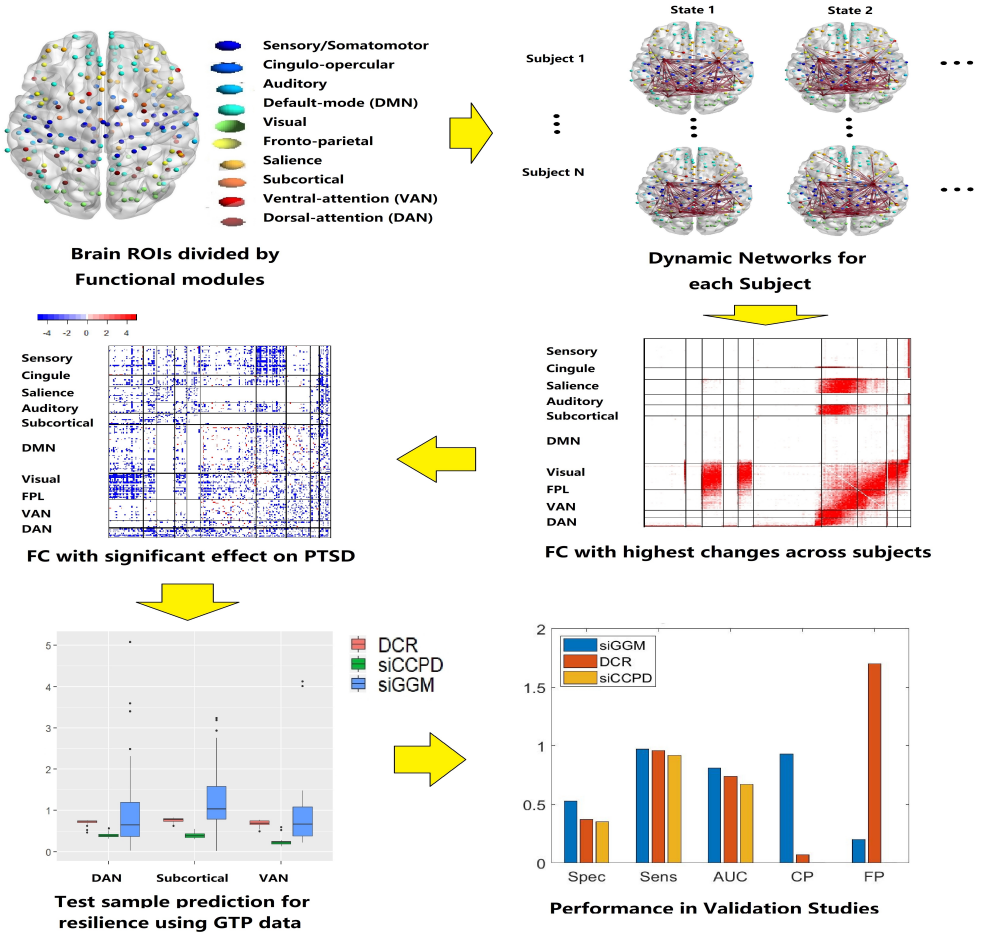

Figure 4.2: The proposed analysis pipeline. Panel (a) illustrates the nodes used in brain functional connectivity that are distinguished based on the known functional modules. Panel (b) illustrates the computed dynamic functional connectivity separately for each individual, the method for which is detailed in the Methods section and Figure 4.1. Panel (c) depicts a heatmaps with summary measures that reflect the degree of temporal variation for edges across all the individuals. Panel (d) illustrates our discovery regarding the edges whose temporal fluctuations are directly related to trauma resilience. Panel (e) provides boxplots for out of sample prediction accuracy using the edge-wise dynamic connections to predict the continuous clinical outcome, via the scalar-on-function statistical methodology. Panel (f) provides a visual depiction of the performance metrics from our extensive validation studies comparing the proposed approach with alternative methods.

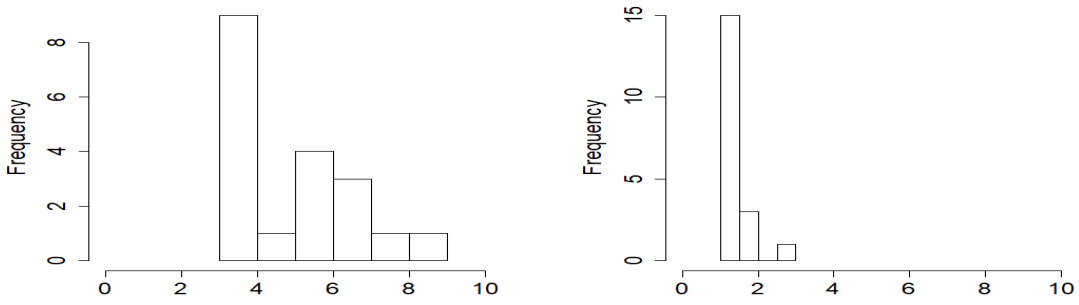

Figure 4.3: Histogram for the number of FC change points detected in the PTSD data analysis. The left and right panels depict the results under the proposed approach and under the SC naive version of the method. The multimodal dynamic FC approach seems to be more sensitive to network changes.

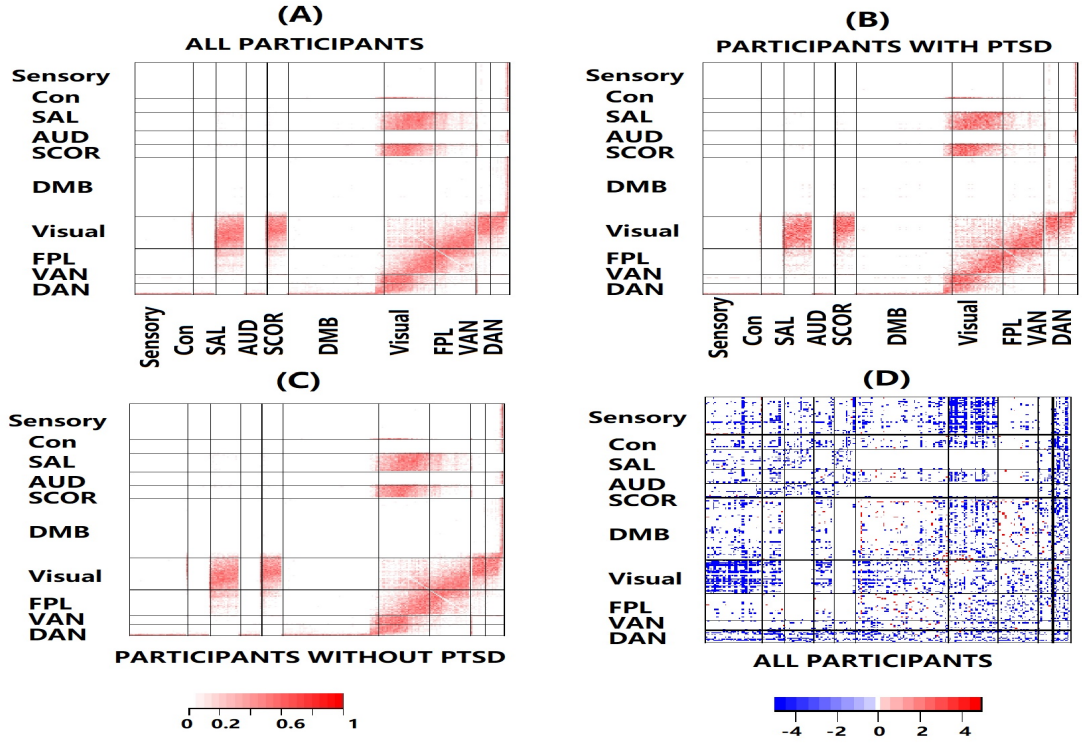

Figure 4.4: Panels (A)-(C) illustrates the edge-wise temporal variation averaged over all individuals, individuals with PTSD and those without PTSD, respectively. Here the temporal variation for an edge was calculated as the ratio of the number of state changes for that edge divided by the number of state changes in the network. Edges in several modules including Visual, SAL, SCOR, VAN and DAN show strong temporal fluctuations resulting from frequent state changes. On the other hand, the Sensory, Cingulomotor, and DMN register the fewest temporal fluctuations over time. Only 213 edges illustrated significant differences in terms of the proportion of edge-specific state changes between the PTSD and non-PTSD groups, which suggests the inadequacy of this measure to distinguish disease severity. Out of these 213 edges, almost all had a higher frequency corresponding to the PTSD group, illustrating higher temporal fluctuations in this cohort. Panel (D) illustrates edges whose fluctuations in terms of the edge strength (measured via edge-specific standard deviations for partial correlations over time) are significantly related to PTSD resilience (multiplicity adjusted). Most of these edges lie between functional modules and are contained between the Visual and other modules, as well as between the DAN and other modules. Blue and red colors imply a negative and positive association with PTSD resilience respectively. It is clear that an increase in temporal edge strength fluctuations in most edges leads to decrease in resilience and vice-versa. However, a small number of connections within DMN and between DMN and other modules lead to increased resilience corresponding to higher fluctuations in edge strength and vice-versa.

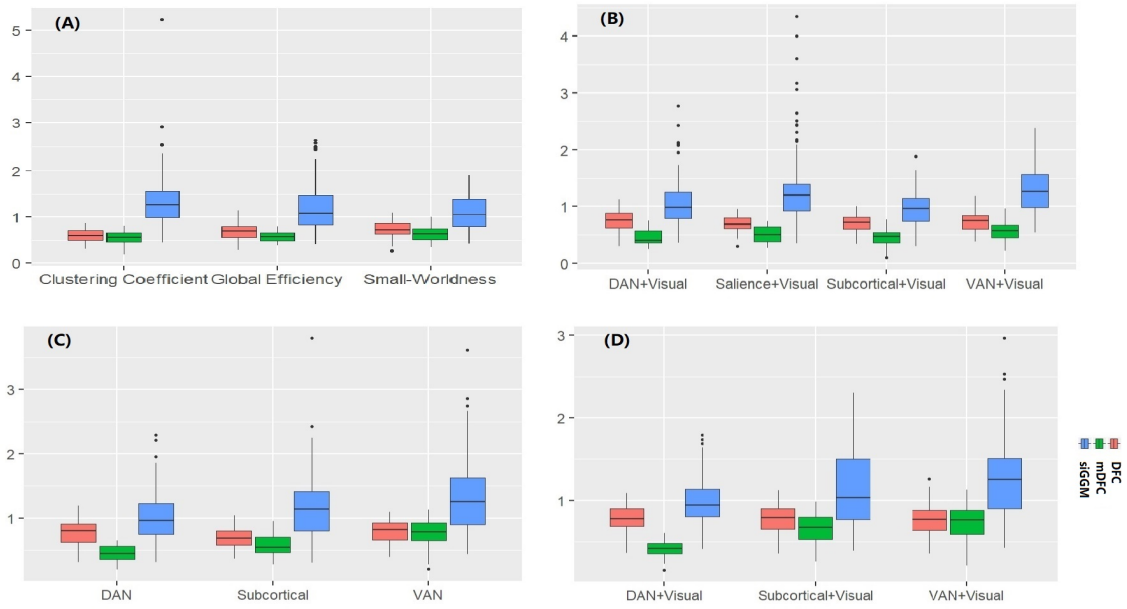

Figure 4.5: Prediction performance in terms of Mean squared error or MSE when using dynamic network metrics for predicting resilience score under mDFC, SC naive version of mDFC (denoted as DFC), and siGGM. The subplots indicate MSE values when using the following time-varying explanatory variables in scalar-on-function regression (A): global clustering coefficient and global efficiency; (B) local clustering coefficient for DAN+VIS, SAL+VIS, SCOR+VIS, VAN+VIS functional modules (C) local efficiency in DAN, SCOR, and VAN functional modules; and (D) local efficiency in DAN+VIS, SCOR+VIS, VAN+VIS functional modules. mDFC has lower or comparable MSE in all cases, and significantly lower MSE (higher prediction accuracy) when using dynamic global efficiency and clustering coefficient; dynamic clustering coefficient using nodes in DAN+VIS modules, SAL+VIS modules, VAN+VIS modules and SCOR+VIS modules; dynamic local efficiency when using nodes from DAN module, as well as DAN+VIS modules.

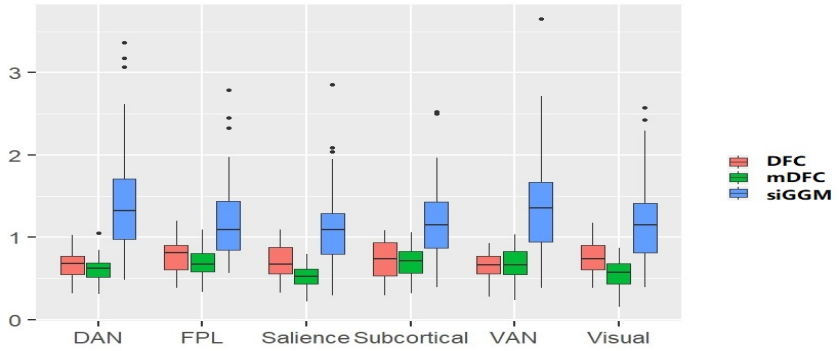

Figure 4.6: Prediction performance in terms of Mean squared error or MSE when using small-worldedness derived from localized functional modules for predicting resilience scores. Results are reported for multimodal dynamic connectivity (mDFC) and the SC-naive version of the method (denoted as DFC), along with the siGGM approach that computes static networks. The results indicate MSE values are lower or comparable under the mDFC method across all local functional modules, with significant improvements corresponding to the salience network.

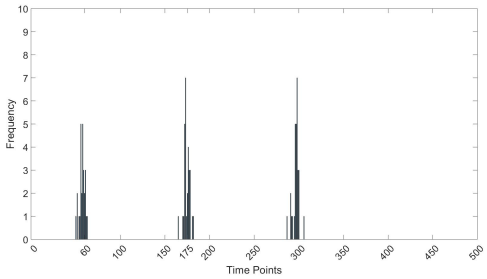

Figure 7(a)

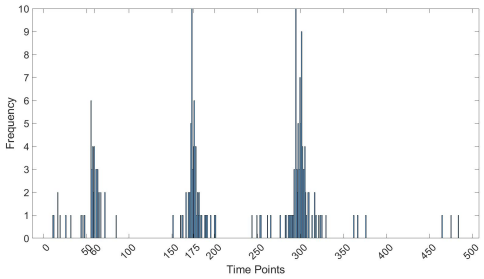

Figure 7(b)

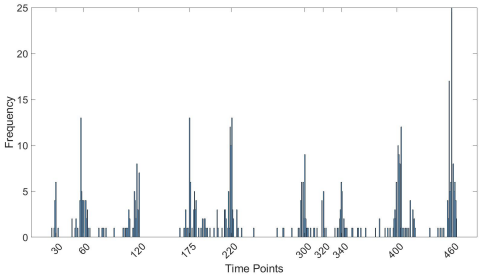

Figure 7(c)

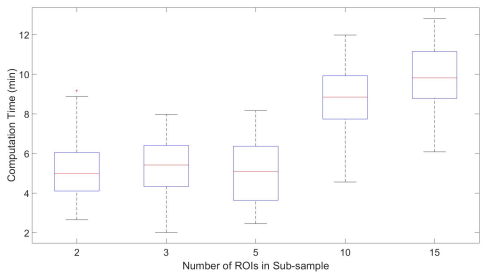

Figure 7(d)

Figure 4.7: Figures 7(a)-7(c) denote frequency plots for change point estimation. Figures 7(a) and 7(b) correspond to the case of $V = 20$ and $V = 100$ nodes respectively, with the true change points being located at 60, 165, and 300. Figure 7(c) corresponds to the case of 10 true change points which are labeled on the X-axis. The histograms show a strong clustering around true change points. Although there exist some loosely grouped frequencies corresponding to spurious change points, they are almost always eliminated through sub-network sampling mechanism. Figure 7(d) depicts the computation time as the sub-network size is varied.

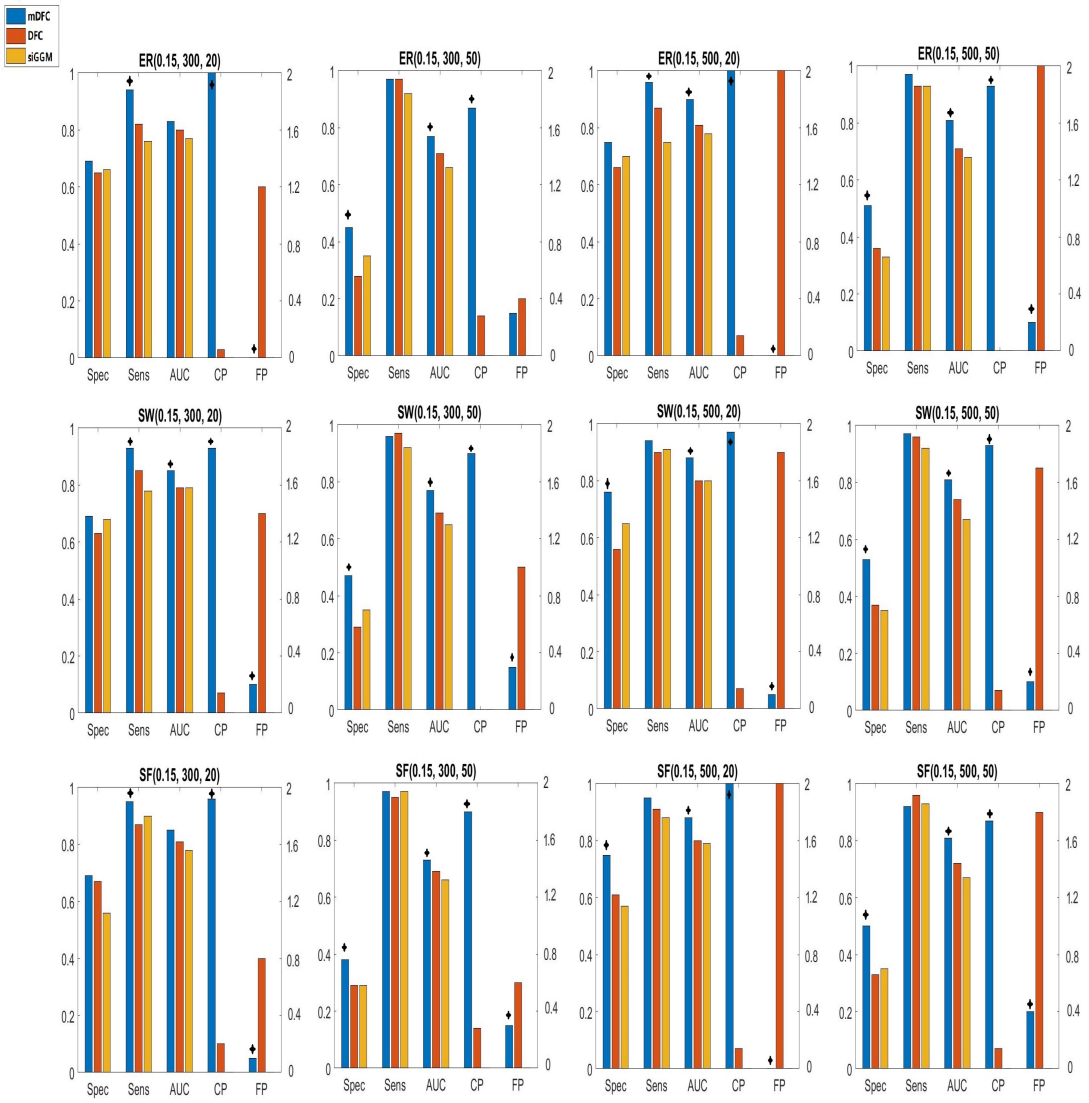

Figure 4.8: Simulation results corresponding to true networks with discrete jumps at change points. There are total 3 true change points for each simulation. The first column denotes different simulation scenarios: ER, SW, and SF, denote Erdos-Renyi, small world network, and scale-free networks respectively. The numbers within the parenthesis denote the network density, number of nodes, and number of time points respectively. CP is the percentage of estimated true change points. FP is the average number of false estimated change points. mDFC has Strong power to detect all true change points without and FP. In terms of graph estimation, mDFC has significant higher AUC compared with siGGM as well as DCR that is denoted as DFC in this Figure. The significantly improved metrics are highlighted in bold.

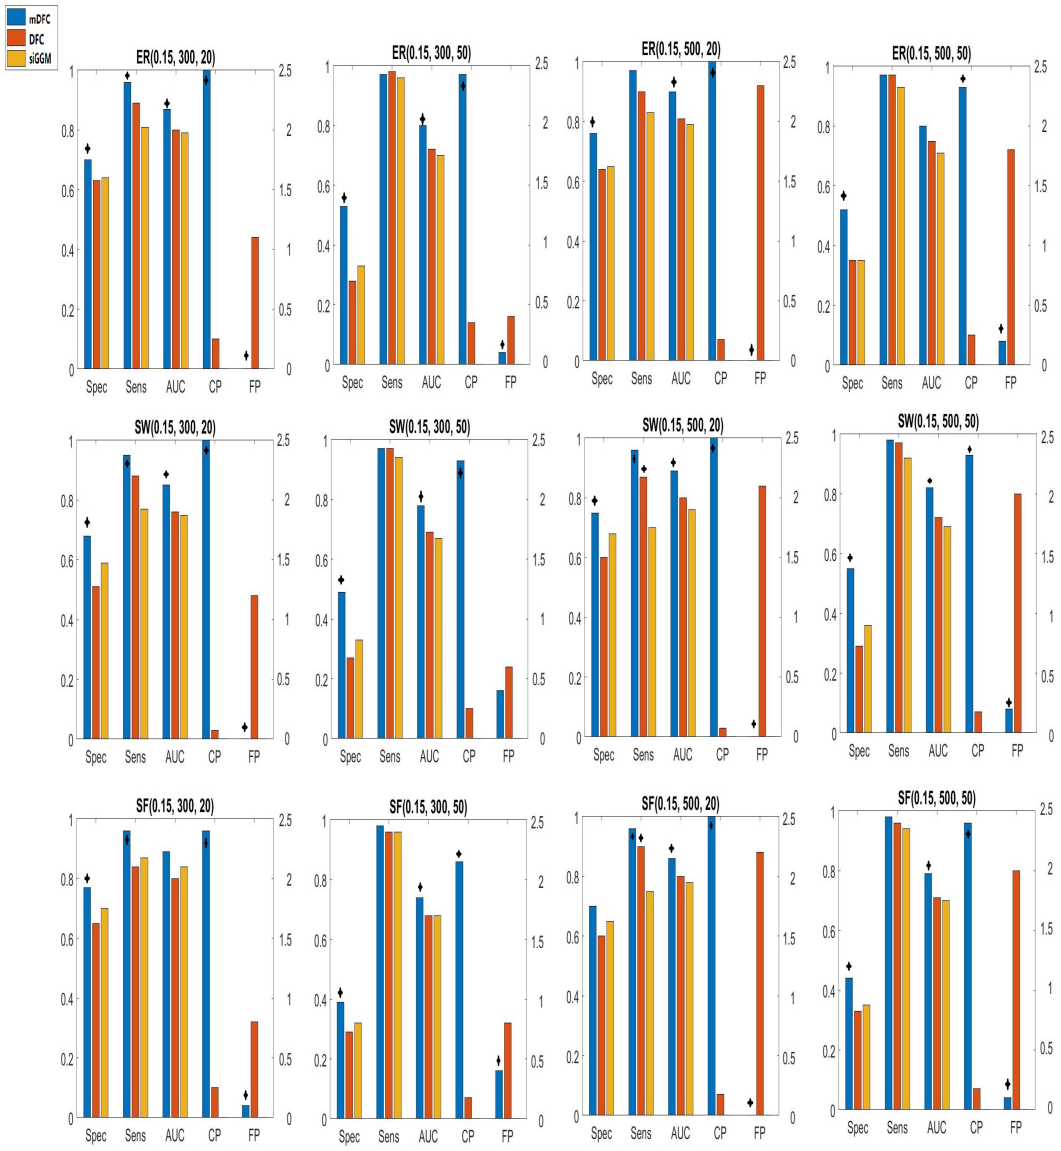

Figure 4.9: Simulation results corresponding to true dynamic networks with three transition periods instead of discrete jumps. The first column denotes different simulation scenarios: ER, SW, and SF, denote Erdos-Renyi, small world network, and scale-free networks respectively. The numbers within the parenthesis denote the network density, number of nodes, and number of time points respectively. CP is the percentage of estimated true change points. FP is the average number of false estimated change points. In terms of change points detection, mDFC performs better than DCR which is denoted as DFC in this Figure. mDFC also has higher AUC compared with siGGM and DCR.

Chapter 5

Summary and Future Work

Neuroimaging, especially fMRI, generate rich data to study the human brain. However, it is also challenging to analysis brain imaging data because of its high dimension and complexity. In addition, the high noise-to-signal ratio of fMRI data and heterogeneous among individuals make it even harder to estimate the subject-based network. Most fMRI study also collect demographic information and behaviour data of participants. But none of the existing method has tried to incorporate this useful information to better estimate the dynamic functional connectivity of a subject. The main objective of our research is to propose a way to discover the temporal relationship between different ROIs of the brain, or dynamic functional connectivity, by incorporating the covaraites (both time-independent and time-dependent).

Our first proposed method is related with time-independent covariates. We propose a fundamentally novel hierarchical Bayesian mixture modeling approach incorporating covariates for estimating a population of individual-specific dynamic FC using heterogeneous multi-subject fMRI data. The covariates are modeled as the weight being used in the mixture model. By doing so, we could borrow the information from other subjects with same or similar covariates, which are supposed to have similar brain network over time compared with subjects with totally different covariates. In

addition, our proposed method has two version, one based on pairwise correlation, and another one based on precision matrix. The simulation studies demonstrate that our two models both estimate accurate network change points with high sensitivity and comparable low false positives. Our method provides a way to inference edge-level change points in addition to the network-level. Using our proposed method, we analyze fMRI data from a study of the block task data involving a semantic verbal fluency. In conclusion, by incorporating covariates, our proposed method provide a better way to cluster subjects into different sub-group based on their network feature. In addition, it has ability to detect rapid change for fMRI data, which is a limitation for most existing change point methods. It also provide a way to inference the edge-level change, which might be more useful in the future brain study.

Our second proposed method is using a novel semi-parametric Bayesian Support Vector Machine (SVM) approach that incorporates high-dimensional networks as covariates and is able to assign varying levels of shrinkage to the coefficients in an unsupervised manner via a Dirichlet process mixture of double exponential priors. Although SVM-based methods are heavily used in classifying mental disorders, there are few, if any, semi-parametric Bayesian SVM approaches for classification based on high-dimensional brain networks that naturally provides the ability to conduct inferences. We apply the approach to a connectome fingerprinting problem using the Human Connectome Project (HCP) data as well as a second application involving classification of individuals with attention deficiency hyperactivity disorder (ADHD) and showcase the superior classification accuracy of the proposed approach.

In chapter three, we examine the potential of multimodal dynamic FC, computed by fusing functional magnetic resonance imaging (fMRI) and diffusion tensor imaging data, in terms of predicting continuous clinical measures of disease severity. We develop concrete measures of temporal network variability that are directly linked with disease severity and identify regions whose temporal connectivity fluctuations

are significantly related to the disease. Our results illustrate the distinct advantages of prediction of disease severity compared to the usual analysis based on disease phenotype categories, it shows that the multimodal approach is more sensitive to connectivity changes and highlights the predictive prowess of multimodal dynamic FC over existing static and dynamic network models.

Appendix A

Computation Detail for Topic 1

A.1 Computational Details for Change Point Estimation

Generalizing the approach in Tibshirani and Wang (2007) to the multivariate case, we apply lowess independently to each time-series of pair-wise correlations as a first step. Denote the smoothed fit as $\bar{\rho}_q = (\bar{\rho}_{q1}, \ldots, \bar{\rho}_{qT})$ for the q-th pair-wise correlation profile, $q = 1, \ldots, p(p-1)/2$. The fraction parameter in the lowess fit which controls the smoothness level, is chosen to be small so as to avoid oversmoothing which will cause difficulty in detecting potential change points. Then for each time series, we compute the first order differences $\delta_{qt} = \bar{\rho}_{qt} - \bar{\rho}_{q,t+1}, t = 1, \ldots, T-1$, followed by the median of $(\delta_{q1}, \ldots, \delta_{q,T-1})$ denoted as $\mu_q, q = 1, \ldots, p(p-1)/2$. Next, we compute the median of the absolute deviations $\{|\delta_{q1} - \mu_q|, \ldots, |\delta_{q,T-1} - \mu_q|\}$, and denote it as $\Delta_q, q = 1, \ldots, p(p-1)/2$. Finally, we note that equation (2) can be expressed as $\min_{\mathbf{u} \in \Re^{p(p-1)/2}} \sum_{t=1}^{T} ||\tilde{\mathbf{r}}_t - \mathbf{u}_t||^2$ subject to $\sum_{t=1}^{T-1} ||\mathbf{u}_{t+1} - \mathbf{u}_t|| \leq s_2$, where $\lambda \propto 1/s_2$, with a smaller value of s_2 implying a lesser number of increments. We choose the

threshold parameter as

$$s_2 = \max_{q=1,\ldots,p(p-1)/2} \{2\Delta_q + \sum_{t=1}^{T-1} |\delta_{qt}| 1(|\delta_{qt}| > 4\Delta_q)\}, \tag{A.1}$$

where the expression inside the parenthesis corresponds to the threshold for the q-th individual pairwise connectivity time series and is motivated by the choice in Tibshirani and Wang (2007), and where $1(\cdot)$ is the indicator function. The above expression (A.1) specifies the threshold for the increment term in the fused lasso as the maximum of the thresholds for the individual pairwise connectivity time series, in order to ensure that no true change point is omitted in the first step. Equation A.4 assumes that first order differences with absolute values greater than $4\Delta_q$ corresponds to a change in connectivity, and the $2\Delta_q$ term ensures that the threshold is not very small and is able to capture all the true change points accurately with no omissions, even at the cost of detecting spurious change points. Using the choice of the threshold in (A.1), we then re-fit the fused lasso (2.2) to obtain an initial estimate of the number of change points ($Kmax$) and their locations $^* = (t_1^*, \ldots, t_{Kmax}^*)$. The initial estimate for the number of change points is potentially inflated due to the choice of a small value of the lowess fraction parameter, and the manner in which λ was computed in (A.1). This is to ensure that we do not exclude true change points in the initial fit. In the next step we propose a screening criteria to exclude false change points in *.

In particular, the screening criteria involves a post-processing step as follows. For each given subset of $k < Kmax$ change-points, we approximate the signals between the successive change-points with the mean value of the points in that interval. Subsequently, we calculate the total sum of squared errors (SSE) between the set of real signals and these piecewise constant approximations to them. Though it may appear computationally intensive to do this for all subsets of $k < Kmax$ change-points, a dynamic programming strategy (Picard et al., 2005) enables us to compute the

subset of k change points having the minimum SSE from among all possible sets of $k < Kmax$ change points in $O(Kmax^2)$ time. The dynamic programming ensures that the computation is tractable even for a large number of nodes and change points. An exhaustive search becomes impossible for large $Kmax$ since the number of partitions of a set with T elements into $Kmax$ segments is $\binom{T-1}{Kmax-1}$. The dynamic programming strategy reduces the computation time from $O(T^{Kmax})$ to $O(T^2)$ using the additive property of SSE. Let $\beta_{k+1}(i,j)$ denote the SSE corresponding to the best partition of the data between the i-th and j-th time points into $k+1$ segments, noting that $\beta_{k+1}(0,n) = SSE(k+1)$. The recursive algorithm is as follows

$$\begin{aligned} k &= 0,\ \forall 0 \le i < j \le T,\ \beta_1(i,j) = \sum_{t=i+1}^{j} ||\tilde{\mathbf{r}}_t - \bar{\mathbf{r}}_{ij}||^2, \\ \forall k &\in [1, Kmax],\ \beta_{k+1}(i,j) = \min_h \{\beta_k(1,h) + \beta_1(h+1,j)\}, \end{aligned} \tag{A.2}$$

where $\bar{\mathbf{r}}_{ij} = \frac{1}{j-i}\sum_{t=i}^{j} \tilde{\mathbf{r}}_t$ is the mean for the vector of sample pairwise correlations between time points i and j. The above dynamic programming takes advantage of the additivity of the SSE, considering that a partition of the data into $k+1$ segments is a union of a partition into k optimal segments and a set containing 1 segment, which enables us to efficiently compute the best partition of the data into $k+1$ segments, $k = 1, \ldots, Kmax$. Once these optimal partitions corresponding to ${}^*_k, k = 1, \ldots, K_{max}$ have been computed via the dynamic programming strategy, one can eliminate spurious change points from this set using the curvature approach described below.

We determine the optimal number of change points by examining the curvature of the SSE curve as follows. Denote the minimum SSE obtained from the subset of all possible k change points derived from the set * as SSE(k), and denote the corresponding locations of the change points as ${}^*_k$. Clearly, SSE($k+1$), will be smaller than SSE(k); however, after a certain point, adding an extra change-point will have a

negligible impact on the SSE reduction. First we normalize the minimum SSE scores $SSE(1), \ldots, SSE(Kmax)$ as $J_k = \frac{SSE(Kmax)-SSE(k)}{SSE(Kmax)-SSE(1)}(Kmax-1)+1$, where $J(1) = Kmax$, and $J(Kmax) = 1$. Then we compute the curvature of these normalized scores via discrete second derivatives as $\nabla_k = J_{k-1} - 2J_k + J_{k+1}$, and select the number of change points as $K = \max\{1 < k < Kmax : \nabla_k > 0.5\}$ such that the second derivative does not rise above a certain threshold on addition of further change points. The choice of the threshold 0.5 is recommended in earlier papers on change point estimation (Picard et al., 2005), and worked adequately well for our applications. The idea behind this approach is that if the curvature of the normalized SSE scores does not change beyond a certain value by adding an extra change point, then that will imply that the SSE does not reduce significantly, suggesting that no additional change points are required. In our experience based on extensive numerical studies, this approach is able to eliminate any false change points included in the initial set of change points * and produces reliable estimates.

Appendix B

Reference

1. Allen, E.A., Damaraju, E., Plis, S.M., Erhardt, E.B., Eichele, T., Calhoun, V.D., 2014. Tracking whole brain connectivity dynamics in the resting state. Cerebral Cortex. 24(3), 663-676.

2. Baker, J. T., Holmes, A. J., Masters, G. A., Yeo, B. T., Krienen, F., Buckner, R. L., & Ongur, D. (2014). Disruption of cortical association networks in schizophrenia and psychotic bipolar disorder. JAMA psychiatry, 71(2), 109-118.

3. Bi, X., Zhao, X., Huang, H., Chen, D., & Ma, Y. (2020). Functional brain network classification for Alzheimer's disease detection with deep features and extreme learning machine. Cognitive Computation, 12(3), 513-527.

4. Bleakley, K., & Vert, J. P. (2011). The group fused lasso for multiple change-point detection. arXiv preprint arXiv:1106.4199.

5. Bressler, S. L., Menon, V. (2010). Large-scale brain networks in cognition: emerging methods and principles. Trends in cognitive sciences, 14(6), 277-290.

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

7. Chang, C., & Glover, G. H. (2010). Time–frequency dynamics of resting-state brain connectivity measured with fMRI. Neuroimage, 50(1), 81-98.

8. Chen, P. Y., Chen, C. L., Hsu, Y. C., & Tseng, W. Y. I. (2020). Fluid intelligence is associated with cortical volume and white matter tract integrity within multiple-demand system across adult lifespan. NeuroImage, 212, 116576.

9. Cribben, I., Wager, T., & Lindquist, M. (2013). Detecting functional connectivity change points for single-subject fMRI data. Frontiers in computational neuroscience, 7, 143.

10. Cribben, I., Haraldsdottir, R., Atlas, L. Y., Wager, T. D., & Lindquist, M. A. (2012). Dynamic connectivity regression: determining state-related changes in brain connectivity. Neuroimage, 61(4), 907-920.

11. Dubois, J., Galdi, P., Paul, L. K., & Adolphs, R. (2018). A distributed brain network predicts general intelligence from resting-state human neuroimaging data. Philosophical Transactions of the Royal Society B: Biological Sciences, 373(1756), 20170284.

12. Fan, L., Su, J., Qin, J., Hu, D., & Shen, H. (2020). A deep network model on dynamic functional connectivity with applications to gender classification and intelligence prediction. Frontiers in neuroscience, 881.

13. Fan, W., Yang, L., Bouguila, N., & Chen, Y. (2020). Sequentially spherical data modeling with hidden Markov models and its application to fMRI data analysis. Knowledge-Based Systems, 206, 106341.

14. Ferguson, M. A., Anderson, J. S., & Spreng, R. N. (2017). Fluid and flexible minds: Intelligence reflects synchrony in the brain's intrinsic network architecture. Network Neuroscience, 1(2), 192-207.

15. Friedman, J., Hastie, T., & Tibshirani, R. (2008). Sparse inverse covariance estimation with the graphical lasso. Biostatistics, 9(3), 432-441.

16. Friedman, J., Hastie, T., & Tibshirani, R. (2010). Regularization paths for generalized linear models via coordinate descent. Journal of statistical software, 33(1), 1.

17. Friston, K. J. (2011). Functional and effective connectivity: a review. Brain connectivity, 1(1), 13-36.

18. Fu, S., Ma, X., Wu, Y., Bai, Z., Yi, Y., Liu, M., ... & Jiang, G. (2019). Altered local and large-scale dynamic functional connectivity variability in posttraumatic stress disorder: a resting-state fMRI study. Frontiers in psychiatry, 10, 234.

19. Greicius, M. D., Flores, B. H., Menon, V., Glover, G. H., Solvason, H. B., Kenna, H., ... & Schatzberg, A. F. (2007). Resting-state functional connectivity in major depression: abnormally increased contributions from subgenual cingulate cortex and thalamus. Biological psychiatry, 62(5), 429-437.

20. Hindriks, R., Adhikari, M. H., Murayama, Y., Ganzetti, M., Mantini, D., Logothetis, N. K., & Deco, G. (2016). Can sliding-window correlations reveal dynamic functional connectivity in resting-state fMRI?. Neuroimage, 127, 242-256.

21. Hutchison RM, Womelsdorf T, Allen EA, et al. Dynamic functional connectivity: promise, issues, and interpretations. Neuroimage. 2013;80:360-378. doi:10.1016/j.neuroimage.2013.05.079

22. Kim, J., Wozniak, J. R., Mueller, B. A., & Pan, W. (2015). Testing group differences in brain functional connectivity: using correlations or partial correlations?. Brain connectivity, 5(4), 214-231.

23. Kundu, S., Ming, J., Pierce, J., McDowell, J., & Guo, Y. (2018). Estimating dynamic brain functional networks using multi-subject fMRI data. NeuroImage, 183, 635-649.

24. Kyllonen, P., & Kell, H. (2017). What is fluid intelligence? Can it be improved?. In Cognitive abilities and educational outcomes (pp. 15-37). Springer, Cham.

25. Le Bihan, D., Mangin, J. F., Poupon, C., Clark, C. A., Pappata, S., Molko, N., & 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.

26. Lindquist, M. A., Xu, Y., Nebel, M. B., Caffo, B. S. (2014). Evaluating dynamic bivariate correlations in resting-state fMRI: a comparison study and a new approach. NeuroImage, 101, 531-546.

27. Lukemire, J., Wang, Y., Verma, A., & Guo, Y. (2020). HINT: A Hierarchical Independent Component Analysis Toolbox for Investigating Brain Functional Networks using Neuroimaging Data. Journal of Neuroscience Methods, 108726.

28. Picard, F., Robin, S., Lavielle, M., Vaisse, C., Daudin, J. J. (2005). A statistical approach for array CGH data analysis. BMC bioinformatics, 6(1), 27.

29. Quinn AJ, Vidaurre D, Abeysuriya R, Becker R, Nobre AC, Woolrich MW. Task-Evoked Dynamic Network Analysis Through Hidden Markov Modeling. Front Neurosci. 2018;12:603.

30. Raichle, M. E., MacLeod, A. M., Snyder, A. Z., Powers, W. J., Gusnard, D. A., & Shulman, G. L. (2001). A default mode of brain function. Proceedings of the National Academy of Sciences, 98(2), 676-682.

31. Sakoglu, U., Pearlson, G. D., Kiehl, K. A., Wang, Y. M., Michael, A. M., Calhoun, V. D. (2010). A method for evaluating dynamic functional network connectivity and task-modulation: application to schizophrenia. Magnetic Resonance Materials in Physics, Biology and Medicine, 23(5-6), 351-366.

32. Sporns, O. (2013). Structure and function of complex brain networks. Dialogues in clinical neuroscience, 15(3), 247.

33. Tibshirani, R., & Wang, P. (2008). Spatial smoothing and hot spot detection for CGH data using the fused lasso. Biostatistics, 9(1), 18-29.

34. Vidaurre, D., Quinn, A. J., Baker, A. P., Dupret, D., Tejero-Cantero, A., & Woolrich, M. W. (2016). Spectrally resolved fast transient brain states in electrophysiological data. Neuroimage, 126, 81-95.

35. Warnick R, Guindani M, Erhardt E, Allen E, Calhoun V, Vannucci M. A Bayesian Approach for Estimating Dynamic Functional Network Connectivity in fMRI Data. J Am Stat Assoc. 2018;113(521):134-151.

36. Widiger, T. A. & Samuel, D. B. Diagnostic categories or dimensions? A question for the Diagnostic and statistical manualof mental disorders. Journal of abnormal psychology 2005;114,494.

37. Woodward, N. D., Karbasforoushan, H., & Heckers, S. (2012). Thalamocortical dysconnectivity in schizophrenia. American Journal of Psychiatry, 169(10), 1092-1099.

38. Yuan, P., Voelkle, M. C., & Raz, N. (2018). Fluid intelligence and gross structural properties of the cerebral cortex in middle-aged and older adults: A multi-occasion longitudinal study. Neuroimage, 172, 21-30.

39. Zeng, L. L., Shen, H., Liu, L., Wang, L., Li, B., Fang, P., ... & Hu, D. (2012). Identifying major depression using whole-brain functional connectivity: a multivariate pattern analysis. Brain, 135(5), 1498-1507.