

UC Irvine

UC Irvine Electronic Theses and Dissertations

Title

Bayesian State Space Models for Dynamic Functional Connectivity using fMRI Data

Permalink

<https://escholarship.org/uc/item/4gq7d5m0>

Author

Lee, Jaylen

Publication Date

2022

Copyright Information

This work is made available under the terms of a Creative Commons Attribution-NonCommercial License, available at <https://creativecommons.org/licenses/by-nc/4.0/>

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA,
IRVINE

Bayesian State Space Models for Dynamic Functional Connectivity using fMRI Data

DISSERTATION

submitted in partial satisfaction of the requirements
for the degree of

DOCTOR OF PHILOSOPHY

in Statistics

by

Jaylen Lee

Dissertation Committee:
Professor Veronica Berrocal, Chair
Professor Babak Shahbaba
Professor Zhaoxia Yu

2022

DEDICATION

To my parents, whose unwavering support let me reach this point. To Gabriela, without whom I would not have succeeded.

TABLE OF CONTENTS

	Page
LIST OF FIGURES	v
LIST OF TABLES	ix
LIST OF ALGORITHMS	x
ACKNOWLEDGMENTS	xi
VITA	xii
ABSTRACT OF THE DISSERTATION	xiv
1 Brain Imaging and Dynamic Functional Connectivity	1
1.1 Functional Brain Imaging	4
1.1.1 Functional Magnetic Resonance Imaging (fMRI)	5
1.1.2 The General Linear Model	7
1.2 Functional Connectivity	8
1.2.1 Linear correlation Based Metrics	9
1.2.2 Decomposition based Metrics	10
1.2.3 Clustering based metrics	12
1.2.4 Graph-based FC Summaries	14
1.3 Dynamic Functional Connectivity	15
1.3.1 Sliding Window Analysis	16
1.3.2 Change point Detection	18
1.3.3 The Hidden Markov Model	20
2 Functional Connectivity Dynamics modulated by Physiological Data	28
2.1 Methods	30
2.1.1 An HMM model for dynamic brain connectivity	30
2.1.2 Modeling connectivity transitions as a function of observed physiological factors	31
2.1.3 Modeling sparsity through a graphical horseshoe prior	34
2.1.4 Posterior Inference	36
2.1.5 Graph Selection	38
2.2 Simulation Study	40

2.3	Case Study	49
2.3.1	Experimental design and data collection	49
2.3.2	Model fitting	51
2.3.3	Inference results	53
2.4	Discussion	57
3	Time-Varying Psychophysiological Interactions	59
3.1	A Time-Varying PPI Model.	62
3.1.1	Shrinkage Prior for the Dynamic Physiological Effects	65
3.1.2	Shrinkage Prior for the PPI Effects.	67
3.2	Posterior Inference	68
3.2.1	Markov Chain Monte Carlo Algorithm	69
3.2.2	Selection of Non-Zero Components	70
3.2.3	Inference on Partial Correlations	74
3.3	Simulation	75
3.4	Case Study	77
3.5	Discussion	84
4	Dynamic Bi-Clustering of fMRI Data	87
4.1	DBC Model	90
4.1.1	Motivation	90
4.1.2	Dynamic bi-clustering	91
4.2	Posterior Inference	96
4.2.1	Updating the Covariate-Informed Model	98
4.3	Simulation	100
4.3.1	Simulation Setting 2	101
4.4	Case Study	102
4.4.1	Model Specification	103
4.4.2	Results	103
4.5	Discussion	104
5	Future Directions	113
5.1	Further Data Fusion	113
5.2	Neural Dynamic Connectivity Modeling	114
5.3	Scalability	115
	Bibliography	116
	Appendix A Appendix	131

LIST OF FIGURES

	Page
1.1 Graphical representation of the a simple HMM. The unobserved states $S(t)$ dynamically transition between a finite amount of states	21
2.1 For each value of τ_0 , I simulate 1000 undirected 100×100 graphs under the graphical horseshoe prior. Plotted above are the 2.5, 50, and 97.5 percentiles of the edge density as a function of τ_0 . A value of $\tau_0 = 1$ leads to approximately a 50% expected edge density, albeit with high variability, in the sampled graphs.	35
2.2 Graphical representation of the proposed PIBDFC. The data y_t^i are emissions from a distribution that is parameterized by a precision matrix $\Omega_{s_t^i}$, which encodes the functional connectivity and is determined by the state active at time t : $s_t^i \in \{1, \dots, S, t = 1, \dots, T, i = 1, \dots, N$. The probabilities of transitions from s_t^i to s_{t+1}^i are given by the (s_t^i, s_{t+1}^i) entry of the $S \times S$ matrix $Q_{\cdot, \cdot, t}^i$. This entry is modeled according to Equation 2.3.	36
2.3 Simulation Study 1: <i>Top</i> : The true partial correlation matrices for each state responsible for generating the simulation data in the Simulation Study 1. <i>Bottom</i> : The estimated partial correlation matrix from the proposed PIBDFC from a single repetition of the simulation. Each estimated partial correlation is the posterior mean of their respective distributions. Cells are set to 0 in post-hoc MCMC by controlling the BFDR at the 0.2 level. See Sections 2.1.5 and 3.3 for details.	41
2.4 Simulation Study 1: True Positive Rate, True Negative Rate, F1 Score, and state accuracy metrics for the PIBDFC, BDFC, and Tapered SW approaches over different assumptions for the signal-to-noise ratio in the correlation structure. Along each horizontal axis is the average strength of the non-zero partial correlations for each state.	43
2.5 Simulation Study 1: <i>Top</i> : The true state transition path for each subject (vertical axis) across each time point (horizontal axis). The color in each cell identifies which precision matrix in Figure 2.3 generated the simulated the data for each subject-time point pairs. <i>Bottom</i> : The <i>maximum a posteriori</i> estimated state trajectories from PIBDFC.	44

2.6	Simulation Study 1: Estimation of the connectivity change points in a representative subject. The horizontal axis indicates the time-points while the vertical axis reports the posterior probability $P(s_t^1 \neq s_{t-1}^1 Y_{1:T}^i)$. The posterior probabilities of a change point are in red, whereas the black spikes represent the true change points for the subject. I also display a horizontal dotted line at 0.95 to reflect the informal rule of declaring a change-point if $P(s_t^1 \neq s_{t-1}^1 Y_{1:T}^i) > 0.95$	45
2.7	Real Data Analysis: the estimated connectivity networks for the ROIs. Nodes represent ROIs and the edges denote the partial correlations between the connected nodes. The edge colors correspond to the directionality of the partial correlations and the width corresponds to the magnitude. Node colors identify clusters of regions into <i>a priori</i> defined networks. See Section 2.3.3 and Table A.1 in the Appendix	52
2.8	Real Data Analysis: Estimated states' transition path for each subject. The horizontal axis indicates the TR with vertical dotted lines indicating portions where the subject raises their arm. Subject sequences are aligned so that the left half of the plot shows sequences from the sham condition and the right half shows sequences from the active condition. The vertical axis display the subject indices, ordered by similarity in state trajectory according to a hierarchical Euclidean distance based clustering of their MAP transition behavior.	55
2.9	Real Data Analysis: The posterior distribution of the group effect of pupillary dilation e^n (left), and individual effects of pupillary dilation e^p . Rows indicate the propensity for transitioning into states 2 and 3 respectively. For the individual effects, subjects are identically ordered as in Figure 2.8. The horizontal dotted line is the posterior mean for the group level effects, $\eta_2 = 0.687$ and $\eta_3 = 0.651$ respectively.	56
3.1	<i>Illustration of the proposed time-varying PPI paradigm</i> : similarly to the Generalized PPI model, the proposed model allows to capture stronger dependence between two realized time-series due to stimulus effects (top). In addition, it allows to capture smooth changes in the correlation over time, independent of the stimulus effect (bottom). For more details, see Section 3.1.	64
3.2	<i>Choice of the global shrinkage parameter κ_B</i> : Simulations from $P(\theta \zeta)$ where the blue line is a smoother estimating $E(\theta \xi)$ (left) and from $P(\zeta \kappa_B, a_\zeta = 0.1)$ (right) in (3.4). The conditional distribution of $\theta \zeta$ is heteroskedastic with increasing mean and variance as a function of ζ . Larger and more varying values of ζ are associated with smaller values of κ_B	67
3.3	The average partial correlation matrices across each condition for a single simulation dataset. (Top row) The average estimated partial correlation matrices across the rest, stimulus 1, and stimulus 2 conditions respectively. (Bottom row) The average data-generating partial correlation matrices across the rest, stimulus 1, and stimulus 2 conditions respectively. We see there are some false positive connections among the 4,5,6 block where	78

3.4	The paths of the estimated partial correlations for the (1,2), (4,5), and (10,11) pairs from a single simulation dataset. The solid black line shows the true partial correlations for the respective region pairs over time. The dashed and dotted lines are the estimates from dppi and gppi respectively, constructed using the method outline in Section 3.2.3. dppi is able to recover the underlying dynamic correlations while also accommodating the effects of stimulus 2 modifying the partial correlations in the (1,2) and (10,11) pairs. With the (4,5) row, we show dppi is also able to perform similarly to gppi in the absence of resting state dynamic partial correlations, showing the benefits of the shrinkage prior on $\theta_d(t)$	79
3.5	The estimated partial correlation between right side Caudate and ACC by the entropy of the slow learning rate process.	84
3.6	The estimated partial correlation between right side Caudate and Putaman by the entropy of the slow learning rate process.	85
3.7	The estimated partial correlation between right side Parahippocampus and the Hippocampus by the entropy of the slow learning rate process.	85
4.1	Toy example of data and objectives considered by our model, with 4 units, 4 signal sources and 3 time points (i.e. $N = 4, R = 4, T = 3$). (a) BOLD signals (represented by blue bars) recorded in multiple units (in this case subjects, represented by rows) at multiple sources (in this case brain regions, represented by the location of a bar) and time points (represented by columns). (b) The objectives of the model are (i) to identify time-specific clusters of signal sources that generate similar signals in a given time point (here 3 clusters, represented by the color orange, green or magenta of a bar, i.e. $c_{rt}^{(s_i)} \in \{1, 2, 3\}$); and (ii) to learn time-invariant clusters of units that exhibit the same clustering of signal sources across time points (here 2 clusters, represented by the color red or blue of a row label, i.e. $s_i \in \{1, 2\}$).	91
4.2	Illustration of the dynamic biclustering model with data simulated with 3 (time invariant) unit clusters, 20 (time varying) signal-source clusters and 2 change points of signal-source clusters. The plot shows simulated observations (colored dots) and the gaussian densities from which they are sampled, according to the cluster of the signal source corresponding to the observation. The data is grouped by unit clusters and time intervals with no re-sampling of signal-source clusters. Here, for simplicity, $\sigma_r = 0.5$ for every signal source $r = 1, \dots, R$	106
4.3	Directed graphical representation of our dynamic biclustering model. The distribution of observed signals $y_{rt}^{(i)}$ (gray circles) at time $t = 1, \dots, T$ depends on the cluster s_i of unit $i = 1, \dots, N$ and on the cluster $c_{rt}^{(s_i)}$ of the signal source $r = 1, \dots, R$ among units in cluster s_i . When $s_i = z$, for $z \in \{1, \dots, Z\}$ and $c_{rt}^{(z)} = k \in \{1, \dots, K\}$, then $y_{rt}^{(i)}$ has a Normal distribution with mean μ_k and variance σ_r^2 . Unit clusters are sampled according to finite stick-breaking distribution π and are invariant over time. For units in cluster $s_i = z$, the cluster of signal source r at time t is the same as that at time $t-1$ when $\gamma_{rt}^{(z)} = 1$ or is re-sampled according to the finite stick-breaking distribution ω_z when $\gamma_{rt}^{(z)} = 0$. For every r and z , $\gamma_{rt}^{(z)} = 1$ with probability a_t . Yellow, diamond-shaped nodes denote hyperparameters with fixed, pre-specified values.	107

4.4	ROC for regional co-clustering across time in Simulation setting 2.	108
4.5	The posterior probabilities of subject level co-clustering. Each each i,j entry in this heatmap is given by $P(s_i = s_j Y)$. Subjects generally fall into either the larger subject cluster 4, or cluster with a single other subject.	109
4.6	The regional cluster assignments for subject cluster 4. The color denotes the cluster assignment corresponding to the means in Table 4.3. The vertical dotted line indicates the raising and lowering of the subject's arm respectively. We see the clusters are stable during the rest sections before reconfiguring during the active portions, aligning with inferences found in Section 2.3. . . .	110
4.7	The change points for subject cluster 4. Time is along the horizontal and the ROI are along the vertical axis. Each entry is $P(\gamma_{rt} = 0 Y)$. The vertical dotted lines are when the subjects are raising and lowering their arms. We see the model is able to capture the expected change points corresponding to the arm movements.	111
4.8	Clustering Similarity for each ROI during the rest and active portions of the time series. Each entry is the Rand Index comparing the row and column ROI cluster assignment agreement. We see during the rest portions there is tighter coupling among sections in the Default Mode Network. During the active sections the co-clustering are more prevalent among the Saliency and Frontal Parietal Control networks.	112

LIST OF TABLES

	Page	
2.1	Simulation Study 2: results over 30 repetitions. I report sensitivity and specificity metrics for the estimated graphs of the corresponding states, together with the overall accuracy of the estimated state sequences. Standard deviations across the 30 simulations are showed in brackets. The proposed method maintains the best balance between sensitivity and specificity as well as latent state estimation accuracy.	46
2.2	Simulation Study 3: Results over 30 repetitions. I show the entry wise true positive and true negative rates for the estimated graphs for the corresponding states. I also show estimated number of the estimated number of change points. PIBDFC performs comparably to CCPD in the setting where change-points are common among subjects despite no explicit assumption of this being the case.	48
3.1	The results of dPPI and gPPI under simulation setting 1. Each column entry is the mean and standard deviation of the estimates of the chosen effect modifications across all simulations. dPPI achieves less bias under a setting where physiological effects are dynamic. Overall	77
4.1	The clustering AUC and MSE from Simulation 1.	101
4.2	The clustering AUC and MSE from Simulation 2.	102
4.3	The cluster means for each region level cluster c_{rt}	104

LIST OF ALGORITHMS

	Page
1 The Baum-Welch Algorithm	23
2 The Viterbi Algorithm	24

ACKNOWLEDGMENTS

I would like to thank the Orange County chapter of the ARCS (Achievement Rewards for College Scientists) Foundation and Beckman-Coulter Inc for their generous financial support.

This work was also supported by the National Science Foundation Graduate Research Fellowship Grant No. DGE-1839285.

VITA

Jaylen Lee

EDUCATION

Doctor of Philosophy in Statistics University of California, Irvine	2022 <i>Irvine, California</i>
Master of Science in Statistics University of California, Irvine	2019 <i>Irvine, California</i>
Bachelor of Science in Mathematics James Madison University	2016 <i>Harrisonburg, VA</i>
Associate of Science in Math Education Thomas Nelson Community College	2014 <i>Hampton, VA</i>

RESEARCH EXPERIENCE

Graduate Research Assistant University of California, Irvine	2019–2022 <i>Irvine, California</i>
--	---

TEACHING EXPERIENCE

Teaching Assistant University of California, Irvine	2017–2019 <i>Irvine, CA</i>
High School Teacher Bruton High School	2016–2017 <i>Williamsburg, VA</i>

SOFTWARE

Predictor-Informed Bayesian Dynamic Functional Connectivity

<https://github.com/jayesrule/PIBDFC>

A Matlab toolbox for conducting Bayesian inference on Dynamic Functional Connectivity

ABSTRACT OF THE DISSERTATION

Bayesian State Space Models for Dynamic Functional Connectivity using fMRI Data

By

Jaylen Lee

Doctor of Philosophy in Statistics

University of California, Irvine, 2022

Professor Veronica Berrocal, Chair

Dynamic Functional Connectivity (DFC) investigates how the interactions among brain regions vary over the course of an fMRI experiment. A challenge in DFC inference is the between-subject heterogeneity in both brain region interactions and the evolution of these interactions over time. In this dissertation, I discuss three proposed state-space models to address these challenges. (1) I propose a multi-subject DFC model where transitions between functional connectivity (FC) states are modeled by a non-homogenous Hidden Markov Model: a state-space model where the transition distribution are informed by subject-specific physiological readings. I discuss the results of this model applied to fMRI data where pupil dilation is simultaneously tracked, suggesting FC stability is linked to attention and arousal. (2) I propose an extension to the single subject Psychophysiological Interaction (PPI) model, which seeks to estimate the effect modification of FC by a stimulus/task. The standard PPI model assumes static FC when the subject is at rest. The proposed model allows for DFC in absence of a stimuli, reducing bias in estimating PPI effects in DFC settings. This model is applied to a predictive learning task, where increased unpredictability of a sequence of images is estimated to be associated to tighter coupling in Anterior Cingulate and Caudate regions. (3) Lastly, I propose a multi-subject bi-clustering model that dynamically clusters brain regions based on their activation profile, then clusters subjects based on their regional clustering behavior over time. This model is applied to a *handgrip* task-fMRI dataset.

Chapter 1

Brain Imaging and Dynamic Functional Connectivity

The brain is extraordinarily complicated. It is estimated that there are 86 billion neurons in the average human brain, each functioning with the other in concert driving how we survive. The brain is responsible for processing our bodily functions, as well as how we perceive the world around us, through coordinating our cognitive functionality. (Herculano-Houzel, 2009). What drives much of the complexity of the brain is how these neurons work together by sending signals to and from each other in response to stimuli received from bodily receptors. Understanding the connections between neuronal cells is critical to understanding the brain and its functionality. Obtaining this understanding of the brain is important especially in the context of neurological diseases such as Parkinson's (Grafton et al., 1994), Schizophrenia (Lynall et al., 2010; Anderson and Cohen, 2013), and Multiple Sclerosis (Lowe et al., 2002). Through obtaining a greater understanding of the healthy brain and the neurological connections within, we can better understand the effects of these diseases and hopefully a path for therapies.

Broadly, the study of the connections within the brain can fall into 3 categories of connectivity. The first is *structural connectivity*, the mapping of the physical white matter tracts connecting cells within the brain. This is vital to understanding the physical mechanisms by which signals traverse the brain. While structural connectivity is important to understand the roads and highways that connect the brain, they do not explain the timing nor the frequency with which information flows between neurons. Such information is the domain of the second and third modes of connectivity: *functional connectivity* and *effective connectivity*. For a review of structural connectivity and its relationship to functional connectivity, I refer to Babaeeghazvini et al. (2021).

Functional connectivity is defined by Friston (2011) as the *undirected* statistical dependencies among remote neurophysiological events, while effective connectivity is the *directed* statistical dependence between remote neurophysiological events. Contrasting the two, effective connectivity takes on a more causal flavor than functional connectivity, seeking to make inferences about which neuronal firing causes others to fire. The objective of effective connectivity is to model the flow of information from stimulus to neural firing and speaks to a more intuitive idea of neural coupling. Functional connectivity does not seek to make these distinctions, instead seeking to form networks of neurons which tend to fire together (connected) or tend to fire independently (disconnected), with no sense of causal direction. While seemingly less informative on the surface, functional connectivity measures provide useful summaries of brain organization, and remains to be one of the more popular forms of analysis of brain imaging data (Friston, 2011). Differences in functional connectivity have been shown to be associated with clinical depression (Craddock et al., 2009), intensity of hallucinogenic symptoms in schizophrenia (Van Den Heuvel and Fornito, 2014; Anderson and Cohen, 2013), and aging (Zhang et al., 2019).

All three modes of connectivity are objects of active and come with their unique experimental and statistical challenges. The framework of interest for this thesis is solely within the realm

of functional connectivity. This thesis aims to present a compilation of statistical methods which seek to address three challenges presented in trying to summarize brain imaging data with functional connectivity:

- Capturing changes in functional connectivity over time and relating these dynamics to simultaneously recorded physiological data, such as heart rate or pupil dilation. Functional connectivity is established to vary over time, on the order of seconds (Chang and Glover, 2011). At this regard, it has become increasingly popular to talk of *Dynamic Functional Connectivity (DFC)* (Hutchison et al., 2013). Less studied is what drives these dynamics, whether they be fatigue, arousal, or even attentiveness. This is important to gain an understanding of how physiological information correlates with functional connectivity. In the second chapter, I introduce a model to quantify the effects of covariates on functional connectivity dynamics.
- Estimating DFC in task-based settings. The typical DFC analysis estimates the effects of tasks on activation and then subtracts them from the measured signals. DFC analysis is then performed on the residuals. This leads to the interpretation of DFC as statistical dependence between brain regions after the stimulus has been accounted for. Instead, in Chapter 3 I propose a model that estimates DFC as well as how it is modified by the task.
- Account for between subject heterogeneity in the DFC summaries. It is understood intuitively that different people’s brains behave in different ways. It has been shown empirically that this too applies to between-person differences in DFC (Tewarie et al., 2018; Lehmann et al., 2017). In Chapter 4 I discuss a multi-subject approach that clusters brain images on two levels, simultaneously grouping together brain regions based on the activation profiles as well as subjects based on their brain region clustering behavior.

In the remainder of this chapter, I will give a brief background of the brain imaging mode we consider in this thesis, i.e, functional Magnetic Resonance Imaging (fMRI), as well as other approaches to DFC.

1.1 Functional Brain Imaging

Friston’s definition of functional connectivity as the set of undirected statistical dependencies among remote neurophysiological events, while useful, must be contextualized on the basis of the available imaging techniques. What exactly are we measuring in the brain regions that we wish to characterize the statistical dependence of? Indeed, the ideal case would be to have neuronal action potentials for all 86 billion neurons across continuous time be our neurophysiological events of interests. Current technology limits the ability to map and measure the activity of this many neurons. Even further, the in-feasibility of both of the storage and processing of this data is quite clear. Thus the need for a proxy measure for neuronal firing that compromises on the following two properties: first, the spatial resolution, i.e, the degree of accuracy to which we can attribute the location in space of neural activity; secondly the temporal resolution, i.e, the sampling frequency with which we can capture neural events. There are a host of ways to measure brain activity and the difference lies in the imaging modality. Electroencephalogram recordings measure electrical signals in the brain; florescence imaging can track the spiking behavior of individual neuronal cells, and, the focus of this thesis, Functional Magnetic Resonance Imaging (fMRI), measures blood flow within the brain. All imaging modalities have their benefits and drawbacks concerning their spatial and temporal resolution or their ability to serve as an adequate proxy for neural activity. I refer to Markicevic et al. (2021) and the references within for a review of these methods and their properties.

1.1.1 Functional Magnetic Resonance Imaging (fMRI)

fMRI is a popular noninvasive imaging technique that detects brain activity in regions throughout the course of an experiment. Brain activity is quantified by changes in relative blood oxygenation. As brain activity increases within a region, localized oxygen consumption also increases. Blood rushes in to meet the demand of the increased oxygen need, but often over-compensates and delivers more oxygenated hemoglobin than the already present deoxygenated hemoglobin. The magnetic properties of oxygenated vs deoxygenated hemoglobin can be detected by the MR scanner and registered as Blood Oxygenation Level Dependent (BOLD) contrasts. The MR scanner scans a slice of the brain across a prespecified dimension and repeats that process until a volume that contains the brain regions of interest has been covered. Once a full tour of the brain has been measured, these slices are compiled into a scan. Each scan is often arranged in an array in which each elementary volume element, termed a voxel, contains a localized measure of BOLD signal (Bowman, 2014). Each scan takes approximately 2-3 seconds to collect, limiting the time resolution of our measurements. For each subject, these scans are collected repeatedly over the course of an experimental run (typically over 500 scans) and sometimes over multiple runs. As a data object, one can store the BOLD signal scans for a single subject as a four way array of dimension $V_x \times V_y \times V_z \times T$, where V is the number of voxels over the x, y, and z axis respectfully and T is the total number of scans collected across time. It is easy to see that a single subject generates a truly massive amount of data which can make inference challenging.

The typical fMRI experiment recruits subjects to lie in the scanner, having their brain repeatedly imaged in various conditions. The goal of these experiments is to relate observed brain activity, measured by the BOLD signal, to the presented conditions. Conditions can vary across experiments and depend on the research question of interest. Some require less active participation on behalf of the subject, such as laying at rest or passively observing pictures. Others are more active, asking the subject to tap their fingers, or make decisions.

The main mode of analysis in fMRI studies is to contrast the images obtained under each condition and quantify the differences within.

There are two main experimental paradigms for the presentation of the conditions. The first is the *block design*, which presents conditions in extended time intervals. For example, subjects can alternate between a rest condition and a finger tapping task every 30 seconds. In this way one can easily contrast the signals obtained in the finger-tapping images with the signals obtained during the at-rest images. The second is the *event-related design* the conditions are short discrete events, such as the presentation of a sound or image. In this paradigm, one can contrast multiple conditions without much concern about subject fatigue or boredom affecting the signals (Lindquist, 2008).

There are many challenges associated with fMRI data induced from data acquisition, preprocessing, modeling, and computation. Firstly, there are some controversies with the information the MR scanner collects itself. (Bowman, 2007) It is an important distinction to note that, while BOLD signal has been linked to neural activity, fMRI does not measure neural activity directly. Rather, it tracks the flow of blood within the brain. Thus, when studying brain response to stimuli using fMRI, it is essential to model the blood response to neural activity. This model is captured by the *hemodynamic response function* (hrf) (Lindquist, 2008). The hrf is known to vary greatly between brain regions and subjects, thus requiring assumptions about its functional form. Aspects of data acquisition such as hardware inconsistencies, head motion, and scanning time serve as sources of artifacts which serve as signal distortions that must be removed prior to statistical analysis. The series of steps where these artifacts are removed are known as the preprocessing pipeline; of which there are many variations. These steps include items like correcting for head motion; realigning slices in time so that each scan in total is associated with a single time point; and spatial smoothing for signal boosting. Each step within the preprocessing pipeline, whether it be motion correction or slice timing correction can have an effect on following analyses

and must be chosen carefully. For an overview of fMRI experiments and common statistical techniques, I refer to the summary work of Lindquist (2008).

Briefly revisiting the topic of the dimensionality of fMRI, considering our goal of quantifying statistical dependence between remote neurophysiological events, we must consider how to define sections of the brain. Different classification on *where* events are occurring obviously affect characterizations of the statistical dependence. fMRI measurements typically parcellate the brain into $3mm^3$ cubes. With such a fine partition of the brain, there is an expectation of spatial correlation among the measured signals: events happening close to each other probably being related. On the other hand, these cubes can be aggregated into larger functional regions, where immediate spatial dependence is of less interest than the dependence between the larger regions. Additionally, when data are aggregated, the dimensionality of the data are reduced, reducing the computational burden. For the work in this thesis, I will discuss methods applied to fMRI data aggregated into larger functional Regions of Interests (ROI). I refer to Tang et al. (2021) for a discussion about the dimensionality of brain imaging data.

1.1.2 The General Linear Model

Single subject statistical analysis of fMRI data is primarily based upon the *General Linear Model* (GLM). In a GLM, the BOLD signal measured for ROI r at time t , $Y_r(t)$ is expressed as a linear function of experimental conditions $X_r(t)$. The model can be written as follows:

$$Y_r(t) = \sum_{k=1}^K X_{rk}(t)\beta_{rk} + \epsilon_r(t) \quad t = 1, \dots, T; \quad r = 1, \dots, R; \quad k = 1, \dots, K, \quad (1.1)$$

where $X_{rk}(t)$ is represents the design matrix obtained as the convolution of the k^{th} stimulus of interest, here denoted by $x_k(t)$, with a pre-specified voxel-specific hrf, $h_r(t)$:

$$X_{rk}(t) = \int_{-\infty}^{\infty} x_k(\tau)h_r(a - t)d\tau. \quad (1.2)$$

Lastly the vector $\beta_r = [\beta_{r1}, \dots, \beta_{rK}]'$ can be seen as the unknown amplitude of bold response to the respective K stimuli and contains the main parameters of interest in activation-based fMRI experiments. Inference for activation is typically performed in terms of contrasts, where the typical form of the hypothesis test is $H_0 : c'\beta_r = 0$ and $H_a : c'\beta_r > 0$. c is the K -dimensional contrast vector which specifies the specific contrast of interest. For example, when testing the activation for a voxel r comparing a rest period, X_{r1} to a stimulus period, X_{r2} . By setting $c = [-1, 1, 0, \dots, 0]'$, we can test the alternative hypothesis $H_a : \beta_{r2} > \beta_{r1}$ or; whether the stimulus has a larger effect on BOLD response than during rest.

This model is important in the context of functional connectivity as the coefficients of the General Linear Model, the β_{rk} , present already some functional connectivity interpretation. Thinking under the previous scenario, if $\beta_{r2} > \beta_{r1}$ and $\beta_{r'2} > \beta_{r'1}$ for two voxels r and r' , we could say there is a marginal dependence between the two voxels in that they both respond to a stimulus. However, there could be additional statistical dependence between the two regions that is unaccounted for by the stimulus. This can be estimated by modeling the distribution of the errors, $\epsilon_r(t); r = 1, \dots, R$.

1.2 Functional Connectivity

Functional Connectivity (FC) has been defined by Friston (2011) as the undirected statistical dependence between distinct regions. Statistical dependence can be assessed in practice through various metrics by which FC can be quantified. In their review paper Li et al. (2009), separate methods for estimating FC in the three types: linear correlation based, decomposition based, and clustering based. I find these three broad categories encompass many of the methods presently used. Each type characterizes FC in different ways but each approach attempts to tackle the main goal in quantifying statistical dependence between regions. As FC is defined as undirected, let \mathcal{S} indicate the symmetric FC matrix for a

specific metric, where the $\mathcal{S}_{rr'}$ entry is the FC between regions r and r' . I give a brief, but not exhaustive listing of techniques under each type. For the sake of simplicity, I describe these methods in the context of a resting state paradigm, where fMRI signal is measured from subjects under no explicit task, thus requiring no modeling of task effects ($X_{rk}(t)$ in the context of Equation 1.1). This means that one can use these approaches on the BOLD signal $Y(t)$ itself. These approaches can usually be applied in the task-based setting simply by applying them to the residuals from Equation 1.1, $(Y(t) - \sum_k X_k(t)\hat{\beta}_k)$, for some estimate of the true stimulus response β_k .

1.2.1 Linear correlation Based Metrics

The following measures give an intuitive measure of statistical dependence between two regions $r, r' = 1, \dots, R$ with $r' \neq r$. These measures are symmetric and bounded between 0 and 1. Magnitudes of 0 imply no linear relationship between the activities of regions r and r' , while magnitudes of 1 imply activities in one region can be perfectly predicted by a linear function of the other. These measures are widely used as they have straight-forward interpretations and require minimal assumptions about the distribution of the data.

Pearson correlation: Let $Y_r(t)$ and $Y_{r'}(t)$ be the BOLD measurements at time t for regions r and r' respectively. The Pearson correlation between regions r and r' is $Corr(Y_r(t), Y_{r'}(t)) = \frac{E((Y_r(t) - E(Y_r(t)))(Y_{r'}(t) - E(Y_{r'}(t))))}{\sqrt{Var(Y_r(t))Var(Y_{r'}(t))}}$. This metric measures the strength of the linear relationship between two regions and is bounded between -1 and 1. A positive Pearson correlation between regions r and r' implies the increase in the BOLD signal in region r is associated with an increase in BOLD signal in region r' . A negative Pearson correlation implies that an increase in r is associated with a decrease in r' . Thinking of the distributions of $\epsilon_r(t)$ and $\epsilon_{r'}(t)$ from Equation 1.1, if we assume these error terms are jointly normally distributed, then $P(\epsilon_r(t)|\epsilon_{r'}(t)) = P(\epsilon_r(t))$ if and only if the Pearson correlation is 0 (Casella and Berger,

2002). That is, if the true Pearson correlation is 0, than the activity in regions r and r' are statistically independent.

Partial correlation: Let $\mathbf{Y}_{(rr')-}(t) = [Y_s(t)]_{s \neq r, r'}$. Consider the models $Y_r(t) = Y_{(rr')-}(t)\beta_r + \eta_r(t)$ and $Y_{r'}(t) = Y_{(rr')-}(t)\beta_{r'} + \eta_{r'}(t)$. The partial correlation between regions r and r' is $Corr(\eta_r(t), \eta_{r'}(t)) = \frac{Cov(\eta_r(t), \eta_{r'}(t))}{\sqrt{Var(\eta_r(t))Var(\eta_{r'}(t))}}$. This metric measures the strength of the linear relationship between 2 regions after all others have been accounted for. Similar to the relationship between pearson correlation and marginal independence under assumed joint normality, there is a relationship between partial correlation and conditional independence under joint normality assumptions.

Coherence: The above methods, while intuitive for measures of statistical dependence across regions, have noted issues when applied to BOLD fmri data. Since the shape of the hrf differs between regions, there could be correlation of the neural activities between regions but it would not be reflected in the correlation (or partial correlation) between the BOLD signals simply because the hemodynamic responses differ at two locations. Sun et al. (2004) propose using coherence, the spectral analog to correlation, instead. Coherence between two regions r and r' at frequency λ , $C_{rr'}(\lambda) = \frac{|f_{rr'}(\lambda)|^2}{f_{rr}(\lambda)f_{r'r'}(\lambda)}$, where f_{rr} is the power spectrum of r and $f_{r'r'}$ is the cross spectrum. This method has been shown in simulation to detect dependence between regions reliably in low frequencies (0.01-0.1 Hz generally).

1.2.2 Decomposition based Metrics

These model-free data-approaches are dimension reduction techniques where the R observed BOLD signals are decomposed into C latent signals, where $C < R$ typically. One can then analyze how the time series for each region is loaded on to the latent signals. If the two regions r and r' load heavily onto the same latent signal $c \in 1, \dots, C$, one can inform some form of statistical dependence between the two. There are a variety of methods to perform

this decomposition; here I focus on two popular ones below.

Principal Component Analysis: One can form the sample covariance $\hat{\Sigma} = \frac{1}{n-1} \sum_{t=1}^T Y_t Y_t'$ where $Y_t = [Y_{t1}, \dots, Y_{tR}]$. This sample covariance matrix can be decomposed into a product of three matrices using singular value decomposition. $\hat{\Sigma} = USV'$, where S is a diagonal $R \times R$ matrix with the elements being the square root of the eigenvalues of $\hat{\Sigma}$, called singular values, sorted in descending order; U and V are the $R \times R$ orthonormal matrices forming the principal components and the loading matrix respectively. Each principle component column can be seen as a linear combination of the regional time series, the coefficients of such can be found in the corresponding column of the loading matrix V . One can then threshold the number of principle components considered by the relative magnitude of the singular values, say S . This gives a lower dimensional representation of the BOLD data. Since the decomposition was on the sample covariance matrix, regions that tend to load highly on the same principal component can be seen as varying together in some sense. See (Baumgartner et al., 2000; Worsley et al., 2005) for more details on the use of Principal Component Analysis in fMRI.

Independent Component Analysis: Independent Component Analysis (ICA) seeks to represent the observed signals among R regions as noisy observations from S statistically independent sources. This assumes the model $Y_t = Ab_t + \eta_{rt}$ where b_t is an $S < R$ length vector representing the unobserved signal strength and $A \in \mathcal{R}^{R \times S}$ is known as the mixing matrix, indicating how the source signals sum together to form each of the observed signals. By finding an estimate for the "unmixing" matrix $W \in \mathcal{R}^{S \times R}$ such that $\hat{b}_t = WY_t$, we can find how the observed signals are loaded onto the latent source signals. Similar to principal component analysis one can look at this loading matrix to see how to observed signals are assigned to the S source signals and generate spatial maps. Regions that fall into the same area of the spatial map are assumed to be generated by the same latent source signal, b_s . In this sense, regions are said to be statistically dependent and functionally connected. For

more details on ICA and its extensions as it's applied to fMRI see (Beckmann et al., 2005; Lukemire et al., 2018; Li et al., 2007).

1.2.3 Clustering based metrics

Rather than approaches with a more explainable interpretations of connectivity tied to understood concepts of statistical dependence, clustering based approaches seek to partition brain regions into clusters which pass some threshold of similarity. Similarity is context specific and the resulting clusters can change depending on the measure of similarity used. There are three main types of clustering based analysis of FC. Distance-Based Clustering, and Model-Based clustering.

Distance-Based Clustering: In this approach, a certain number of clusters are pre-specified. Regions are then assigned to these clusters based on some measure of distance from the cluster centroid: a measure of average among regions assigned to the cluster. Algorithms generally proceed by initializing the clusters, computing the centroid, then reassign the regions to clusters that minimize the distance to the centroids. Venkataraman et al. (2009) proposed using K-Means clustering where the distance metric used was simple Euclidean distance between the BOLD values i.e $\sum_{t=1}^T (Y_{rt} - Y_{r't})^2$. This approach was found to consistently cluster together motor regions as well as visual regions of the brain. Ramezani et al. (2010) propose a spectral clustering approach, where the distance measure is based on the eigenvalues of the Laplacian of a similarity matrix. In brief detail, they calculate a symmetric similarity matrix $S \in \mathcal{R}^{R \times R}$ where $s_{rr'} = \exp(-\frac{Y_r - Y_{r'}}{2\sigma^2})$ for some kernel width $\sigma^2 > 0$. Then, they compute the Laplacian of this similarity matrix $L = I - DSD$ where D is a R sized diagonal matrix with entries $D_{rr} = \sum_{i=1}^R \sqrt{S_{ri}}$, followed by computing the eigenvalues of L and then performing k-means clustering on the rows of the matrix formed from the associated eigenvectors. These form clusters of regions that tend to covary over time while

avoiding limitations over the allowable shapes of the clusters as in Principle Component Analysis. Golay et al. (1998) propose a fuzzy clustering approach using a transformation of the Pearson Correlation coefficient as a distance metric: $\left(\frac{1-CC_{rr'}}{1+CC_{rr'}}\right)^\beta$ for some $\beta > 0$. This approach differs from k-means clustering in that all regions are not given a single cluster assignment but rather given a probability of membership to each cluster. This allows the possibility of more complex FC maps as regions may have high probabilities of belonging to more than one cluster.

Model Based Clustering: These approaches cluster on more distributional aspects of the data rather than a concrete measure of similarity or distance. Here, I briefly discuss useful applications of this clustering approach.

Amar et al. (2015) introduced a time series bi-clustering model to cluster brain regions along with subjects based on whether the region was activated or not. They specify a hierarchical model where regions are partitioned into 'modules' where every region within the modules is activated at the same time. In practice, this is enacted by letting each regional cluster be jointly distributed from an 'active' distribution F_1 , or a dormant distribution, F_0 . In their fMRI case study, they specify $F_0 = N(0, \sigma_0)$ and $F_1 = N(2, \sigma_1)$ for $0 < \sigma_0 \leq \sigma_1$. More specifically, let $Y_r^i(t)$ be the BOLD signal of subject i for ROI r at time t .

$$P(Y_r^i(t)|HS_r^i, C^i(t)) = HS_r^i C^i(t) F_1 + (1 - C^i(t)) F_0$$

$$P(H_r^i = 1|H_r) = H_r p_1 + (1 - H_r) p_0,$$

where HS_r^i is an indicator for region r belonging to the module set of subject i , or the set of voxels that are allowed to activate. They pool together information across subjects by allowing subjects to have shared modules $H = [H_1, \dots, H_R]$, deemed core modules. Time dependence is incorporated through module activations, $C^i(t)$, determining when the module of subject i is active, but not over assignments of clusters. A drawback of this is failing to allow the brain regions to dynamically reconfigure their spatial dependencies, an important

feature in brain imaging time-series (Chang and Glover, 2011; Calhoun et al., 2014).

Kim et al. (2010) propose a sort of second level analysis on the activation levels β_r from the GLM model 1.1. They assume a mixture of experts model, where the estimated activation values for each region r is modeled as a mixture of Gaussian distributions where the influence of each Gaussian is related to the distance of the region from the cluster center. The likelihood of this model is written as

$$P(\beta_r|x_r, \theta) = \sum_{z_i \in \mathcal{C}} p(\beta_r|z_i\theta)p(z_i|x_r, \theta),$$

where z_i denote the cluster membership labels, x_r is the spatial location of the center of region r in some coordinate system, and θ succinctly indicates the mean, covariance matrix, and weights of the associated Gaussian distributions for each cluster i.e $\mathbf{E}(\beta_r|z_i = c_m, \theta_m) = k_m \exp(-(x_r - b_m)' \Sigma_m^{-1} (x_r - b_m))$ for the likelihood of region r activation under cluster c_m . The weights on this mixture are spatially weighted, whereby clusters whose centers are closer to the region are more likely to be selected to make the prediction. The values of θ follow a Dirichlet process which induces the clustering mechanism (Ferguson, 1973). The combination of the Dirichlet process and the spatial influence weights allow for both close range and distant correlations in activation to be considered.

1.2.4 Graph-based FC Summaries

In the section we have described various measures of FC. It is important to note that all measure simply investigate the presence or strength of connection between only two regions of the brain over the course of measurement of a single subject. Given the possibly massive number of regions considered, it is important to describe these connections in a meaningful way. A popular approach is through the lens of undirected graph theory, where a graph is a collection of nodes and edges (Alderson-Day et al., 2016; Van Den Heuvel and Fornito,

2014; Fornito et al., 2013). In the case of fMRI data, each region can be considered a node and the presence of FC, as measured by the methods above, can be represented as an undirected edge between any two nodes. Once these graphs are constructed, one can further summarize them using network properties like efficiency, clustering behavior, and degree strength. Van Den Heuvel and Fornito (2014) found that in comparison to neurotypical subjects, schizophrenic subject displayed functional connectivity characteristics of decreased degree strength in the frontal lobe, lower overall efficiency (average path length) and increased segregation. Houck et al. (2017) found higher degree strength in auditory cortex among schizophrenic subjects compared to healthy controls. Those graph theoretical metrics are useful for describing FC, which can in turn help characterize population FC differences.

Inline with discovering differences between populations, the study of FC has shown to be a useful feature in classifying brain disease status. These classification approaches generally take the form of letting X be an $N \times (\frac{R(R-1)}{2} - 1)$ where each row is a subject and each column corresponds to the strength of the FC metric i.e an entry in the lower triangular of the FC matrix \mathcal{S} . y is a N length vector of each subject’s disease status. These classification approaches seek to construct a classifying function $f(X_i)$ that predicts subject i ’s disease status based on their FC. This approach has been applied to classifying mild cognitive impairment (Meszlényi et al., 2017), depression (Craddock et al., 2009), autism (Nielsen et al., 2013), and nicotine dependence (Pariyadath et al., 2014) using various machine learning approaches.

1.3 Dynamic Functional Connectivity

All approaches described in the previous section assume *static* connectivity patterns throughout the course of the fMRI experiment. Under such an assumption, functional brain connectivity is represented by spatially and temporally constant relationships among the regions of

the brain. However, in practice, the interactions among brain regions may vary during an experiment. For example, different tasks, or fatigue, may trigger varying brain interactions. Therefore, more recent work has pointed out that it is more appropriate to regard FC as *dynamic* over time (Chang and Glover, 2010; Hutchison et al., 2013; Calhoun et al., 2014). In particular, it has been shown that partial correlations between brain regions change on the the order of seconds, creating a need for a dynamic estimate of correlations (Chang and Glover, 2011; Vidaurre et al., 2017). This type of analysis is called time varying FC (TVFC). Within this paradigm, the patterns of brain connectivity across the whole brain are assumed to be time-varying and nonstationary, even within a single experiment. The term "chronnectome" has been introduced to describe the growing focus on identifying time-varying, but reoccurring, patterns of coupling among brain regions (Calhoun et al., 2014).

Methods for estimating the chronnectome can generally fall under three main categories. Sliding Window approaches apply static FC techniques to smaller overlapping sections (windows) of the BOLD time series, giving a sequence of connectivity summaries. Changepoint Detection approaches seek to partition the time series into non-overlapping windows by finding changepoints: an instance in the time series where the connectivity changes in some meaningful way. Finally, I recognize state-space models, the approach of this thesis, where the observed data are modeled as an emission of some dynamic latent process. In the next subsection, I will discuss the three approaches. As in the previous section, I discuss them under a single subject resting state paradigm for simplicity.

1.3.1 Sliding Window Analysis

In Sliding Window (SW) analyses, time windows are defined by their length, typically ranging from 30 seconds to 240 seconds. In terms of data points in the time series, let us call the length of the window w . A static metric of FC is then calculated on the first w time points,

say $f(Y(1), \dots, Y(w)) = \hat{\Sigma}_1$ for some function f that takes in a time series of size w and outputs a FC metric $\hat{\Sigma}$. The window is then shifted over by a single time points and, again, the FC metric of interest is computed, say $f(Y(2), \dots, Y(w + 1)) = \hat{\Sigma}_2$. This process is repeated until a collection of FC metrics are computed, $\hat{\Sigma}_1, \dots, \hat{\Sigma}_{T-w}$. Popular metrics of FC are generally correlation (Chang and Glover, 2010; Yu et al., 2015a; Damaraju et al., 2014) and partial correlation (Allen et al., 2014), noted especially for their fast computation time and interpretability.

It is usually assumed that this sequence of FC matrices $\hat{\Sigma}_1, \dots, \hat{\Sigma}_{T-w}$ exhibit possibly re-occurring patterns, especially so for temporally close windowed FC matrices due to the overlapping windows using many of the same time points for the calculation of $\hat{\Sigma}_t$ and $\hat{\Sigma}_{t+1}$. To uncover the predominant patterns, k -means can be run on the sequence $\hat{\Sigma}_1, \dots, \hat{\Sigma}_{T-w}$, typically using Frobenius distance as a distance (Hutchison et al., 2013). The resulting output consists of both cluster centers, termed FC states, $\mathcal{S}_1 \dots \mathcal{S}_K$ and cluster (state) assignments $c_1, \dots, c_{T-w}, c_t = 1, \dots, K$. Once these are obtained one can describe the cluster centers as in the static case of Section 1.2.4, as well as describe the dynamics of the cluster assignments.

There are benefits to using a SW approach, in that it requires few assumptions beyond whatever is required for the FC metric. The procedure for following a SW analysis is straightforward and easily adapted for different FC metrics. This simplicity produces interpretive results, leading to the popularity of SW as a method for assessing TVFC. Using SW, Damaraju et al. (2014) expanded on results found by Van Den Heuvel and Fornito (2014) to show that schizophrenic subjects do not solely exhibit decreased FC in the frontal lobe compared to healthy controls but rather switch between low and high FC states, spending a longer amount of time in the low FC state. Premi et al. (2020) showed increased fluidity between states of professional chess players compared to beginner players of similar age.

However, those SW approaches also present a number of limitations. Most notably with

choosing the window size w . Window sizes should be large enough to reliably estimate FC as well as capture slower low frequency correlations between regions. However, if w is too large, one runs the risk of averaging over the more rare transient states and missing key FC dynamics (Hutchison et al., 2013). Lindquist et al. (2014) showed in simulations that short window lengths may give rise to spurious connections even when the truth is the complete lack of statistical relationship between regions. Outlying observations are even more influential as the correlations are computed from a smaller set of time points. This is worsened due to correlations between subsequent FC matrices $\hat{\Sigma}_t$ and $\hat{\Sigma}_{t+1}$, so that if a spurious connection is found at time t it is also likely to be found at time $t + 1$. Shakil et al. (2016) show that the detection of state transitions and durations is most strongly influenced by the window length and offset, leading to poor estimates of correlation for each brain state. Additionally clustering based on the sliding window time course did not reliably reflect the underlying state transitions. There is no current consensus on how to choose this window size beyond the rule-of-thumb of choosing w larger than the longest BOLD wavelength (Leonardi and Van De Ville, 2015).

1.3.2 Change point Detection

To partially address the issues induced by arbitrary window lengths, Cribben et al. (2012) and Xu and Lindquist (2015) have investigated greedy algorithms, which automatically detect change points in the dynamics of the functional networks. This approach partitions the time series into contiguous blocks of time points where the FC is assumed to be constant at each block. Let B be the total number of blocks, these approaches seek to find the changepoints t_2, \dots, t_B where $f(Y(1), \dots, Y(t_2 - 1)) = \hat{\Sigma}_1, \dots, f(Y(t_B), \dots, Y(T)) = \hat{\Sigma}_B$, for some FC computing metric f . In an approach called Dynamic Connectivity Regression, Cribben et al. (2012) uses the GLASSO of Friedman et al. (2008) to estimate the inverse covariance matrix,

or precision matrix by minimizing a penalized loss function:

$$\hat{\Sigma}_b = f_\lambda(Y(t_b), \dots, Y(t_{b+1} - 1)) = \arg \min_{\Omega} \left(\text{tr} \left(\sum_{t=t_b}^{t_{b+1}-1} Y(t)Y(t)'\Omega \right) - \log|\Omega| + \lambda\|\Omega\|_1 \right),$$

where $\|\Omega\|_1$ is the L1 matrix norm of Ω . Xu and Lindquist (2015) instead estimate a sparse covariance matrix by using an adaptive thresholding technique and use that as their measure of FC. Both approaches use a greedy algorithm which recursively evaluates all possible change points and chooses the best split based on a likelihood function.

Kundu et al. (2018) extend this approach into the domain of a multi-subject analysis. They operate on the sequence of vectors ρ_1, \dots, ρ_T where each entry in ρ_t are the pairwise sample correlations at timepoint t . The sample correlation between regions r and r' at timepoint t can be estimated by aggregating across subjects. Let Y_{rt}^i be the BOLD signal of subject $i = 1, \dots, N$, of region r at timepoint t . Then if the j^{th} entry of ρ_t is the sample correlation between regions r and r' at timepoint t , $\rho_{jt} = \frac{\sum_{i=1}^N (Y_{rt}^i - \bar{Y}_{rt})(Y_{r't}^i - \bar{Y}_{r't})}{\sqrt{\sum_{i=1}^N (Y_{rt}^i - \bar{Y}_{rt})^2 \sum_{i=1}^N (Y_{r't}^i - \bar{Y}_{r't})^2}}$, where $\bar{Y}_{r't} = \frac{1}{N} \sum_{i=1}^N Y_{r't}^i$. They then use a multivariate extension of the Fused Lasso of Tibshirani and Wang (2008) to find the changepoints, t_b in the mean behavior of the sequence of ρ_t . Once the changepoints are found, the FC matrices are estimated using the GLASSO as in Cribben et al. (2012). Importantly, this approach makes a critical assumption that the FC is the same for all subjects at all timepoints, which may only be justifiable in a task-based setting.

The changepoint approach is powerful in that it avoids specifying a window size and automatically determines the FC states as they change through time by finding the change points. However, this very structure requires estimating the FC matrices Σ_b for each state separately whereas instead it may be desirable to borrow strength across similar connectivity states in order to increase the accuracy of the estimation.

1.3.3 The Hidden Markov Model

There are many approaches, including the novel ideas of this thesis, that use the modeling paradigm of the state space model. The state space model is useful for modeling time series dynamics. This technique relies on modeling two quantities in parallel: the observed data $Y(t)$, and an unobserved state $S(t)$, which serves as the underlying latent driver of the observed data. A discrete state space model, in which the latent state $S(t)$ may only take on one of Z values, is the Hidden Markov Model (HMM). In this set up, the observed data is modeled through an emission distribution $P(Y(t)|S(t), \theta_{S(t)})$ and the unobserved state is assumed to follow a discrete Markov Chain of order k : $P(S(t)|S(t-1), \dots, S(t-k)) = Q_{S(t-1)\dots S(t-k)}$ for some Z dimensional vector $Q_{S(t-1)\dots S(t-k)}$ where the entries sum to 1. Usual implementations of the HMM are of order $k = 1$, where the distribution of the of the state at timepoint t depends only on the state at timepoint $t - 1$. In such a case, the conditional distribution of the state transitions $P(S(t)|S(t-1))$ can be fully represented by a $Z \times Z$ matrix, Q where the (i, j) entry is $P(S(t) = j|(S(t-1) = i)$. The full likelihood for this model is in Equation 1.3:

$$P(\mathbf{Y}|\theta_{1:Z}) = \sum_{\mathbf{s} \in \mathcal{Z}^T} \pi_{S(1)} P(Y(1)|S(1), \theta_{S(1)}) \prod_{t=2}^T Q_{S(t-1), S(t)} P(Y(t)|S(t), \theta_{S(t)}), \quad (1.3)$$

where the sum is taken over \mathcal{Z}^T : the set of T -dimensional vectors where the entries can take values from $1 - Z$.

Of course, while physical phenomenon underpinning the states does not necessarily exist in truth, it is useful for creating more interpretable models so that the states represent some property of the the data. In the case of TVFC, the latent states of the HMM represent some measure of FC that, paired with a likelihood, can be inferred through some parameter θ_z for $z = 1, \dots, Z$. In the previous section, I discussed the sliding window technique and how k -means clustering is commonly used to illuminate the prominent and reoccurring patterns of FC. This process is automatic under the HMM as the FC at timepoint t is the equal to

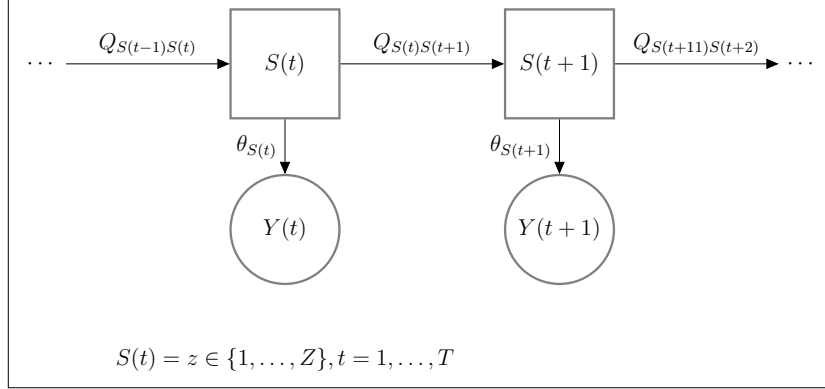


Figure 1.1: Graphical representation of the a simple HMM. The unobserved states $S(t)$ dynamically transition between a finite amount of states

the FC at timepoint t' if $S(t) = S(t')$. For specific HMM treatments of TVFC, we start with the order 1 HMM implementation by Eavani et al. (2013) on resting state subjects. The authors assume $Z = 4$ latent states, which parameterize the mean and covariance matrix of a Gaussian distribution. The log-likelihood for the mean and covariance conditional on the latent states can be written as

$$l(\{\mu_z\}_{z=1}^Z, \{\Sigma_z\}_{z=1}^Z) = \sum_{z=1}^Z N_z (\log|\Sigma_z^{-1}| - tr(S_z \Sigma_z^{-1})) \quad (1.4)$$

where N_z is the number of timepoints assigned to state z and S_z is the sample covariance computed from the same points. The covariance matrix is approximated with a state dependent mixture of K rank one matrices $\hat{\Sigma}_z = \sum_{k=1}^K c_{zk} \mathbf{b}_k \mathbf{b}_k^T$. All elements of the basis vectors \mathbf{b}_k are constrained to both be sparse and lie in the interval $(-1, 1)$, so that the outer product $\mathbf{b}_k \mathbf{b}_k^T$ produces a rank-one matrix that represents a measure of correlation between regions. The authors call these matrices basis networks. The mixture components $c_{rk} \geq 0$ change with the HMM states but importantly the basis components remain static. This set up reflects a belief that relatively few ROI change their co-variance pattern, from state to state, rather the prevalence of these basis networks as exhibited by the data fluctuate over time. The log likelihood in Equation 1.4 is modified by adding two terms. The first is a term for approximation error using the mixture of bases representation of the covariance matrix Σ_z : $\log|\hat{\Sigma}_z \Sigma_z^{-1}| - tr(\hat{\Sigma}_z \Sigma_z^{-1})$. The last encourages differences in the bases mixture across

states by encouraging the inner product $\langle c_i, c_j \rangle, i < j \leq Z$ to be small. The model is fit by maximizing the modified log likelihood through a nested EM algorithm.

One of the important challenges for making inference on HMM models is the complex likelihood of Equation 1.3. As the sum in Equation 1.3 is taken over Z^T possible configurations of the latent state sequence \mathbf{S} , methods that rely on computing this likelihood, i.e maximum likelihood estimation, quickly become infeasible for even moderately sized time series lengths T . There are two important dynamic programming algorithms used in fitting this model that help address this challenge: the forwards-backwards Baum-Welch algorithm (Baum et al., 1970) for computing the state transition distribution Q and the Viterbi algorithm (Viterbi, 1967) for computing the optimal state sequence.

The Baum-Welch Algorithm

The Baum-Welch algorithm is a special case of the EM algorithm that is useful for estimating the parameters of an HMM by way of maximum likelihood: namely the initial state distribution $\pi = P(S(1) = z)$ and the state transition distribution $Q_{zz'} = P(S(t) = z | S(t-1) = z')$. With a emission distribution $L_z(Y(t)) = P(Y(t) | S(t) = z, \theta_z)$, the Baum-Welch algorithm updates π and Q by traversing two recursive procedures. The first is the forward filtering procedure which computes $alpha_z(t) = P(Y(1), \dots, Y(t), S(t) = z)$. The second is the backwards smoothing procedure computing $\beta_z(t) = P(Y(t+1), \dots, Y(T) | S(t))$. Once these two quantities are computed, Bayes Rule can be used to obtain $\gamma_z(t) = P(S(t) = z | Y) = \frac{P(Y | S(t)) P(S(t))}{P(\mathbf{Y})} = \frac{\alpha_z(t) \beta_z(t)}{\sum_{j=1}^Z \alpha_j(t) \beta_j(t)}$, which is exactly the updated initial state distribution, π at $t = 1$. Additionally we can compute $x_{zz'} = \frac{P(S(t)=z, S(t+1)=z', \mathbf{Y})}{P(\mathbf{Y})} = \frac{Q_{zz'} \alpha_z(t) \beta_{z'}(t+1) L_{z'}(Y(t+1))}{\sum_{r=1}^Z \sum_{s=1}^Z Q_{rs} \alpha_r(t) \beta_s(t+1) L_s(Y(t+1))}$. Lastly the state transition distribution can updated as $P(S(t) = z' | S(t-1) = z) = \frac{\sum_{i=1}^{T-1} \xi_{zz'}(t)}{\sum_{t=1}^{T-1} \gamma_z(t)}$. The algorithm for this procedure is Algorithm 1.

Algorithm 1 The Baum-Welch Algorithm

Require: Y Initialize π, Q

for $z = 1 : Z$ **do** ▷ initialize forward procedure

$$\alpha_z(1) \leftarrow \pi_z L_z(Y(1))$$

end for

for $t = 2 : T$ **do** ▷ Begin forward recursion

for $z = 1 : Z$ **do**

$$\alpha_z(t) \leftarrow L_z(Y(t)) \sum_{j=1}^Z \alpha_j(t-1) * Q_{jz}$$

end for

end for

$$\beta_z(T) \leftarrow 1$$

for $t = (T-1) : 1$ **do** ▷ Begin backward recursion

for $z = 1 : Z$ **do**

$$\beta_z(t) \leftarrow \sum_{j=1}^Z Q_{zj} * \beta_j(t+1) * L_j(Y(t+1))$$

$$\gamma_z(t) \leftarrow \frac{\alpha_z(t) \beta_z(t)}{\sum_{j=1}^Z \alpha_j(t) \beta_j(t)}$$

for $z' = 1 : Z$ **do**

$$\xi_{zz'}(t) \leftarrow \frac{Q_{zz'} \alpha_z(t) \beta_{z'}(t+1) L_{z'}(Y(t+1))}{\sum_{r=1}^Z \sum_{s=1}^Z Q_{rs} \alpha_r(t) \beta_s(t+1) L_s(Y(t+1))}$$

end for

end for

end for

$\pi_z^* \leftarrow \gamma_z(t)$ ▷ update initial state distribution

$Q_{zz'}^* \leftarrow \frac{\sum_{t=1}^{T-1} \xi_{zz'}(t)}{\sum_{t=1}^{T-1} \gamma_z(t)}$ ▷ update state transition distribution

The Viterbi Algorithm

Another useful inference to obtain from an HMM is the most likely sequence of states $\hat{S}(t)$ given the observed data \mathbf{Y} and the HMM parameters π and Q . This can be obtained using the Viterbi algorithm, outputting a sequence of states called the Viterbi path. The first step of the algorithm is similar to the forward filtering procedure of the Baum-Welch algorithm. First, we compute $\alpha_z(t) = P(Y(1), \dots, Y(t), S(1), \dots, S(t-1), S(t) = z) = \max_j L_z(Y(t))\alpha_j(t-1) * Q_{jz}$. We additionally update the best path by keeping track of the maximizing state of $\alpha_z(t)$, i.e, we track a variable $V_z(t) = \arg \max_j L_z(Y(t))\alpha_j(t-1) * Q_{jz}$. Once these two quantities are computed, we set $\hat{S}(T) = \arg \max_z \alpha_z(T)$ and recursively update $\hat{S}(t) = V_{\hat{S}(t+1)}(t+1)$. An algorithm for this process is shown in Algorithm 2.

Algorithm 2 The Viterbi Algorithm

Require: $Y \pi Q$

```

for  $z = 1 : Z$  do                                     ▷ initialize forward procedure
     $\alpha_z(1) \leftarrow \pi_z L_z(Y(1))$ 
end for
for  $t = 2 : T$  do                                       ▷ Begin likelihood recursion
    for  $z = 1 : Z$  do
         $\alpha_z(t) \leftarrow \max_j L_z(Y(t))\alpha_j(t-1) * Q_{jz}$ 
         $V_z(t) \leftarrow \arg \max_j L_z(Y(t))\alpha_j(t-1) * Q_{jz}$ 
    end for
end for
 $\hat{S}(T) \leftarrow \arg \max_z \alpha_z(T)$                        ▷ Most likely ending state
for  $t = (T-1) : 1$  do
     $\hat{S}(t) \leftarrow V_{\hat{S}(t+1)}(t+1)$                        ▷ compute Viterbi path
end for

```

Other HMM Approaches to TVFC

The Baum-Welch and Viterbi algorithms are great tools for finding maximum likelihood solution to Equation 1.3. However there are Bayesian approaches the HMM model, where one defines a prior $P(\theta, Q, \pi, S(1 : T))$ and then makes inferences on the encoding and decoding tasks through the posterior distribution $P(\theta, Q, \pi, S(1 : T) | \mathbf{Y})$. An example of this

is implemented by Vidaurre et al. (2017) where they describe whole brain resting state FC using 12 state HMM. The authors assume a similar likelihood as in Equation 1.4 and place conjugate Wishart and Normal priors on the inverse covariance matrix, Σ_z^{-1} , and mean, μ_z , respectively. The authors used a Variational Inference approach to speed up computation and scale up analysis for 820 subjects. Interestingly, they perform clustering over the transition kernel to identify 2 metastates, where transitions between states that share a metastate is more probable. This model, as opposed to the one proposed by Eavani et al. (2013), parameterize changes in FC through the progression of the inverse covariance matrix, Σ_z^{-1} also known as the *precision matrix*, as opposed to the (uninverted) covariance matrix Σ_z . This reflects an interest in the progression of conditional independence between brain regions over unconditioned independence.

Another Bayesian HMM approach that models the precision matrix is the Bayesian Dynamic Functional Connectivity (BDFR) model by Warnick et al. (2018) which provides a basis for our future development in the Chapter 2. BDFR helps introduce the concept of *selection* when modeling TVFC through HMMs. They propose a principled, fully Bayesian approach for studying TVFC. More specifically, they cast the problem of inferring time-varying functional networks as a problem of dynamic model selection in a Bayesian framework. In their model formulation, they capture spontaneous brain activity via Gaussian graphical models, also known as covariance selection models (Lauritzen, 1996), as a way of estimating functional network connectivity. They do this by first adapting a proposal for inference on multiple related graphs put forward by Peterson et al. (2015). This model formulation further assumes that the connectivity states active at the individual time points may be related within a *super-graph* and imposes a sparsity inducing Markov Random field (MRF) prior on the presence of the edges in the super-graph. MRF priors have been used extensively in recent literature to capture network structures, particularly in genomics (Li and Zhang, 2010; Stingo et al., 2011, 2015) and in neuroimaging (Zhang et al., 2014; Smith and Fahrmeir, 2007; Lee et al., 2014). Finally, they embed a Hidden Markov Model on the space of the

precision matrices.

The basic idea lies in the specification of the Graphical Wishart distribution on the precision matrices:

$$P(\Omega_z | G_z, b, D) \propto I(\Omega_z \in \mathcal{P}_{G_z}) |\Omega|^{b/2} \exp\left(-\frac{\text{tr}(\Omega_z D)}{2}\right)$$

$$P(\mathbf{g}_{ij} | \nu_{ij}, \Theta) \propto \exp\left(-\nu_{ij} \sum_{z=1}^Z g_{ijz} + \mathbf{g}_{ij}' \Theta \mathbf{g}_{ij}\right),$$

where $b > 2$ is the degrees of freedom, D is a positive definite scaling matrix, and \mathcal{P}_{G_z} is the set of positive definite matrices with zero entries corresponding to zero entries in an adjacency matrix G_z . The adjacency matrix entries $\mathbf{g}_{ij} = [g_{ij1} \cdots g_{ijZ}]'$ are 0 when ROI i and ROI j are conditionally independent when FC state z is active. The prior on the super graph $P(\mathbf{g}_{ij} | \nu_{ij}, \Theta)$ can control how related the connectivity selection parameters g_{ijz} are between states, z , through the off-diagonal entries in the $Z \times Z$ similarity matrix Θ . $\Theta_{zz'} > 0$ implies stronger coupling between selection of connections in states z and z' . This approach is great in that it gives a hard-selection solution, where off-diagonal entries in the precision matrix Ω_z are either 0 or not according to the graph parameterizing FC state z : G_z . This is opposed to the models introduced above, which either performed no regularization of the FC matrices, in the case of the HMMs of Eavani et al. (2013) and Vidaurre et al. (2017); or performed shrinkage but not selection, in the case of the tapered sliding window (Allen et al., 2014) or the Dynamic Connectivity Detection model (Cribben et al., 2012), which both rely on the G-lasso (Tibshirani and Wang, 2008) to estimate Ω . However, one of the drawbacks of this graph-based selection approach is the computation burden of updating the super graph G . In chapter 2, I show taking a shrinkage approach can speed up inference without sacrificing selection capabilities.

Further, all the above HMM approach assume a stationary transition distribution Q . However, there are increasing studies that highlight how subjects are more likely to experience particular connectivity states when some underlying physiological conditions are present.

For example, Chandra and Bhattacharya (2019) have investigated the association between heart rate variations and functional connectivity. Similarly, in a sleep fMRI study, El-Baba et al. (2019) have shown that transitions between connectivity states slow as subjects fall into deeper sleep stages. As a further example, Kucyi et al. (2017) have described how connectivity dynamics are associated with attentiveness in a pencil tapping task. These studies, among others, have motivated the need for models that provide a better understanding of how the transitions between different brain connectivity states are modulated by internal or external conditions measured during the course of an experiment. In chapter 2, I propose a non-homogenous extension to the HMM, where Q is modeled as a function of time-varying physiological measurements. Working along the same trend, in Chapter 3 I investigate the modification of FC by an explicit task or stimulus. I propose an extension to the single subject Psychophysiological Interaction (PPI) model, which seeks to estimate the effect modification of FC by a stimulus/task. The standard PPI model assumes static FC when the subject is at rest, which recent work in TVFC has shown does not hold (Chang and Glover, 2010; Calhoun et al., 2014; Chiang et al., 2018). The proposed model allows for DFC in absence of a stimuli, reducing bias in estimating PPI effects in DFC settings. Lastly, in Chapter 4, I conclude this dissertation by investigating subject-level heterogeneity in FC evolution over time. Differences in TVFC have been associated with age (Chen et al., 2019), cognitive impairment (Yu et al., 2015b; Lombardo et al., 2020), and substance abuse severity (Prete et al., 2017). However, analysis of this heterogeneity has typically been ad-hoc, clustering FC matrices accross subjects after TVFC estimation. I propose a more efficient solution with a multi-subject bi-clustering model that dynamically clusters brain regions based on their activation profile, then clusters subjects based on their regional clustering behavior over time. Information can be shared among subjects in the same subject cluster by sharing regional cluster assignment distributions.

Chapter 2

Functional Connectivity Dynamics modulated by Physiological Data

In this chapter, I introduce a multi-subject non-homogeneous HMM modeling framework that allows the estimation of the effect of physiological factors on the transition probabilities between states which are shared between subjects. More specifically, for each subject, the state transition probabilities are assumed to vary over time and to depend on a set of time-varying, subject-specific, physiological covariates that are simultaneously recorded during the fMRI experiment. This setting allows the estimation of unique connectivity state transitions for each subject and relates the heterogeneity to the physiological data in an interpretable way. Furthermore, I also allow for recurring connectivity patterns and sharing of networks among the subjects, thus allowing for group-based inferences. With respect to multi-step inference strategies described above, in our approach both the dynamic connectivity states and their association with the physiological measurements are estimated in a single modeling framework, accounting for all uncertainties. Kundu et al. (2018) have recently proposed a two-step multi-subject fused-lasso approach for detecting change points in brain connectivity. Differently than their proposal, our method does not assume that the experimental

design and the timing of the change-points between connectivity states are shared among all subjects, and therefore it can be applied to more general experimental designs. Indeed, our approach allows for differing state transition behavior across multiple subjects by modeling the state transition parameters hierarchically. Our modeling approach further assumes sparsity in the network structures, attenuating spurious or weak connections using a shrinkage prior on the connectivity networks. Additionally, I propose a strategy for edge selection that combines the posterior shrinkage informed thresholding approach of Carvalho et al. (2010) with the Bayesian False Discovery Rate controlling method of Müller et al. (2006).

I apply our modeling framework on a resting-state experiment where fMRI data have been collected concurrently with pupillometry measurements, leading us to assess the heterogeneity of the effects of changes in pupil dilation on the subjects' propensity to change connectivity states. Changes in pupil diameter have been linked to attention and cognitive efforts modulated by the activity of norepinephrine-containing neurons in the locus coeruleus (LC). For example, Joshi et al. (2016) have shown that LC activation predicts changes in pupil diameter that either occur naturally or are caused by external events during near fixation, as in many experimental tasks. See also Joshi and Gold (2022). Therefore, pupil dilation has been used as a proxy for a metric of a person's willingness to exert more effort and involve greater mental resources to complete a task. In our application, I demonstrate how the model can recover expected change points in time-varying functional connectivity states, as those states align quite well with the experimental events regulated by the behavioral task. The rest of the paper is organized as follows. In section 2 I describe our proposed method and edge selection procedure. I also give a brief synopsis of our Markov Chain Monte Carlo (MCMC) approach to posterior inference. In section 3 I showcase our model performance on simulated data and provide comparisons to sliding window and homogeneous HMM approaches. Lastly, in Section 4, I apply our model to the LC handgrip data with accompanying results and analysis. Section 5 concludes the paper with a discussion.

2.1 Methods

In this section, I describe the proposed predictor-informed multi-subject model for dynamic connectivity. This is a single-step fully Bayesian approach that explicitly models the heterogeneity in the dynamics of connectivity patterns across all subjects and at the same time estimates the predictor effects on those dynamics. This is achieved by constructing a non-homogeneous Hidden Markov Model (HMM) where the transition probabilities are functions of time-varying covariates.

2.1.1 An HMM model for dynamic brain connectivity

Let $Y_t^i = (Y_{t1}^i, \dots, Y_{tR}^i)^T$ denote the vector of fMRI BOLD responses measured at time t in R regions of interests (ROIs), $t = 1, \dots, T$ on subject $i = 1, \dots, N$. I adopt a Gaussian graphical model framework, and assume multivariate normality of the bold signals, that is $Y_t^i \sim N_R(\mu_t^i, \Omega_t^{-1,i})$, where μ_t^i is a mean regression term and Ω_t^i indicates a time-varying precision matrix, i.e. the inverse covariance matrix at each time point. In graphical models, the zeros of the precision matrix correspond to conditional independence; that is, if an off-diagonal element $\omega_{jkt} = 0$, $j, k = 1, \dots, R, j \neq k$, then the signals Y_{tj}^i and $Y_{tk}^i (j \neq k)$ are conditionally independent. The mean term μ_t^i can be specified as a general linear model (Friston, 1994) to capture activation patterns over time, as done for example in Warnick et al. (2018). Here, however, since I am interested in capturing connectivity patterns through the modeling of the time-varying precision matrix, I assume without loss of generality that the BOLD signal has been mean-centered by removing any estimated trend and activation component. This “detrending” is not uncommon for studying functional connectivity, especially for task-based fMRI data, where the data are first mean-centered, to remove any systematic task-induced variance, and the analysis is then conducted on the time series of the residuals (see, e.g. Fair et al., 2007).

I propose to model the dynamics of functional connectivity using an HMM framework with S latent states characterizing functional connectivity and the brain transitions between during the fMRI experiment. This formulation captures the heterogeneity of individual-specific patterns of connectivity over time, since each subject’s fMRI data may be characterized by specific change-points and evolution of the brain functional organization. However, I assume that the connectivity patterns may also be re-occurring and they can possibly be shared among the subjects, thus allowing for group-based inferences. Let (s_1, \dots, s_T) be a T -dimensional vector of categorical indicators s_t , such that $s_t = s$ if state s is active at time t , $s = 1, \dots, S$. Then, I assume the data follow a Gaussian graphical model at time t ,

$$Y_t^i | s_t^i = s, \Omega_s \sim N_R(0, \Omega_s^{-1, i}), \quad s = 1, \dots, S \quad (2.1)$$

that is, the connectivity networks may vary by subjects and at each time they are characterized by the values of one among S precision matrices, identifying which state is active at that time.

2.1.2 Modeling connectivity transitions as a function of observed physiological factors

Many neuroscience experiments involve the simultaneous collection of fMRI data together with physiological, kinematics and behavioral data (Wilson et al., 2020). For example, this motivating application considers a handgrip task where pupillometry dilation data (i.e., measurements of pupil dilation sizes) are concurrently recorded. Pupillary dilation is regarded as a surrogate measure for activity in the locus coeruleus circuit, which plays a central role in many cognitive processes involving attention and effort, and it is considered the main source of the neurotransmitter noradrenaline, a chemical released in response to pain or stress. Neuronal loss in the locus coeruleus is known to occur in neurodegenerative disorders such as Alzheimer’s disease and related dementias as well as Parkinson’s disease dementia. It is

therefore important to understand how brain dynamics may be differentially modulated as a function of pupil dilation in different subjects. Recent methods for studying such association use a multi-step approach, first identifying switches in subjects’ state sequences and then calculating the difference between the normalized pupil size before and after the estimated switch (see, e.g. Hussain et al., 2022).

Here, I propose to model the dynamics of functional connectivity using a non-homogeneous HMM which assumes that switches between states are regulated by transition probabilities that also vary over time (non-homogeneity). More in detail, the transition probabilities are modeled as a function of time-varying covariates, e.g. physiological factors like the pupillary data in our motivating application,

$$Q_{rst}^i = P(s_{t+1} = s | s_t = r) = \frac{\exp(\xi_{rs}^i + \mathbf{x}_t^{i^T} \boldsymbol{\rho}_s^i)}{\sum_{k=1}^K \exp(\xi_{rk}^i + \mathbf{x}_t^{i^T} \boldsymbol{\rho}_k^i)}, \quad r, s = 1, \dots, S, \quad (2.2)$$

where \mathbf{x}_t^i denotes a $B \times 1$ vector of covariate values for subject i at time t , and $\boldsymbol{\rho}_s^i$ is the corresponding $B \times 1$ vector encoding the effect of each covariate on the probability of transitioning to state s for subject i . The parameter ξ_{rs}^i defines a baseline transition probability from state r to state s for subject i , that is the transition probability without any covariate effect. To ensure identifiability, I define a state as reference. Without loss of generality I set $s = 1$ as the reference state, and also set ρ_1^i and $\xi_{1k}^i = 0$, $i = 1, \dots, N$. Thus, the state transition coefficients are interpreted with respect to the reference state, and I can rewrite model (2.1) –(2.2) in terms of the log-relative odds of the transition from state r to state s compared to the transition from state r to the reference state 1,

$$\log\left(\frac{Q_{rst}^i}{Q_{r1t}^i}\right) = \xi_{rs}^i + x_t^i \rho_r^i, \quad r, s = 1, \dots, S \quad (2.3)$$

In this formulation, the transition coefficients $\exp(\xi_{rs}^i)$ and $\exp(\rho_{rb}^i)$ are more naturally interpreted, respectively, as the expected odds of transitioning from state r to s compared to transitioning from state r to 1, when the time varying covariates, x_t^i , are 0 (or at a baseline/average value), and the relative change in odds of transitioning to state s compared to

transitioning to state 1, after a one unit change in x_{tb}^i holding all other covariates constant.

I assume independent Gaussian priors for the transition parameters ρ and ξ . I further allow for sharing of information in estimating the state transition structure across subjects, by employing a hierarchical modeling formulation for the state transition parameters. More specifically, I assume that the individual coefficients ξ_{rs}^i and ρ_{rs}^i vary around population-level means, Z_{rs} and η_s , as follows:

$$\begin{aligned}
s_{t+1}^i | s_t^i = r &\sim \text{Multi}(Q_{r,\cdot,t}^i) \quad t = 1, \dots, T, \\
\xi_{rs}^i &\sim N(Z_{rs}, \sigma_\xi) \\
\rho_s^i &\sim N(\eta_s, \sigma_\rho) \\
Z_{rs} &\sim N(z_{rs}^0, \sigma_z) \\
\eta_s &\sim N(0, \sigma_\eta),
\end{aligned} \tag{2.4}$$

where $Q_{r,\cdot,t}^i = (Q_{r,1,t}^i, \dots, Q_{r,S,t}^i)^T$, and $r, s = 1, \dots, S$. By allowing each subject to have their own transition parameters the model allows for unique subject-level transition behavior but is also able to capture population-level estimates through the group level parameters. The interpretation of the group level parameters, η and Z , are similar as their single subject counter parts. The prior means Z_{sr} are usually set to 0 except for the diagonal elements z_{rr}^0 , $r \neq 1$, which is set to be positive to encourage state persistence over time (self-transitions) and thus more stable estimated state sequences. Keeping in mind that these state transition parameters operate on the log odds of transition relative to state 1, and that interpretation of the parameters require exponentiation, a small shift in value for the state transition parameters can result in rather large changes in state transition behavior. To this end, I recommend setting the variance parameters of the priors for ξ , ρ , Z_{sr} and η_r to some small positive value on the order of 0.1.

2.1.3 Modeling sparsity through a graphical horseshoe prior

Functional networks are thought to exhibit so-called small world behavior, indicating a high degree of clustering and high efficiency in the estimated networks (Alderson-Day et al., 2016; Essen and Tononi, 2016). This leads to an expectation of sparsity within functional networks and the associated precision matrices. In a Bayesian framework, sparsity can be enforced by postulating either selection- or shrinkage-inducing priors. Selection involves inferring which off-diagonal elements of the precision matrix should be suppressed to exact zeros. Warnick et al. (2018) achieve such a selection by using a G-Wishart prior to sample positive definite matrices according to estimated adjacency matrices that corresponds to the functional connectivity networks. This selection approach is intuitive and leads to straightforward inferences via the posterior probabilities of inclusion of the elements of the precision matrix. However, it is computationally challenging and doesn't scale well to relatively large dimensions of the networks. Here, instead, I take a shrinkage based approach and model the off-diagonal entries of the state-specific precision matrices Ω_s , $s = 1, \dots, S$ in (2.1) by employing a Graphical Horseshoe prior (Li et al., 2019). Thus, I set

$$p(\Omega_s | \tau, \Lambda) \propto \prod_{j < k} N(\omega_{jks} | \lambda_{jk}^2 \tau^2) \prod_{j < k} C_+(\lambda_{jk} | 0, 1) I(\Omega_s \in S_R), \quad s = 1, \dots, S \quad (2.5)$$

where $I(\Omega_s \in S_R)$ is an indicator function to ensure that samples of Ω_s belong to the space of positive definite $R \times R$ matrices and $C_+(\cdot; \mu, \sigma)$ denotes a half-cauchy distribution with location parameter μ and scale σ . In (2.5), I further assume a non-informative flat prior for the diagonal elements, i.e. $\omega_{jkt} \propto 1$. The shrinkage of the off-diagonal elements is obtained through the combined effect of the variance components λ_{jk}^2 and τ^2 in the normal priors for ω_{jkt} , $j = 1, \dots, k-1$, $k = 1, \dots, R$. The parameter τ is a *global shrinkage parameter*, that controls how sparse the precision matrix is in its entirety. The parameter $\lambda_{jk:j < k}$ defines instead a *local shrinkage parameter*, since it allows to shrink each individual off-diagonal entry ω_{jk} towards zero, whereas it maintains the magnitude of non-zero entries and thus reduces




Figure 2.1: For each value of τ_0 , I simulate 1000 undirected 100×100 graphs under the graphical horseshoe prior. Plotted above are the 2.5, 50, and 97.5 percentiles of the edge density as a function of τ_0 . A value of $\tau_0 = 1$ leads to approximately a 50% expected edge density, albeit with high variability, in the sampled graphs.

biases. Following Li et al. (2019), I assume a half-Cauchy prior on τ , $\tau \sim C_+(\cdot; 0, \tau_0)$, with τ_0 indicating an *a priori* belief about the global sparsity of the estimated graph. In order to specify τ_0 one can simulate graphs under the informal selection rule of Carvalho et al. (2010), where an edge j,k is selected if $E(\frac{1}{1+\lambda_{jk}\tau}) < 0.5$. I demonstrate such a process in Figure 2.1 by simulating 1,000 undirected graphs from the model, a larger τ_0 is associated with higher expected edge densities *a priori*. Additionally, I find that a $\tau_0 = 1$ gives an expected edge density of approximately 50% while having the largest spread. Figure 2.2 provides a graphical representation of the proposed predictor-informed Bayesian dynamic functional connectivity model (PIBDFC).

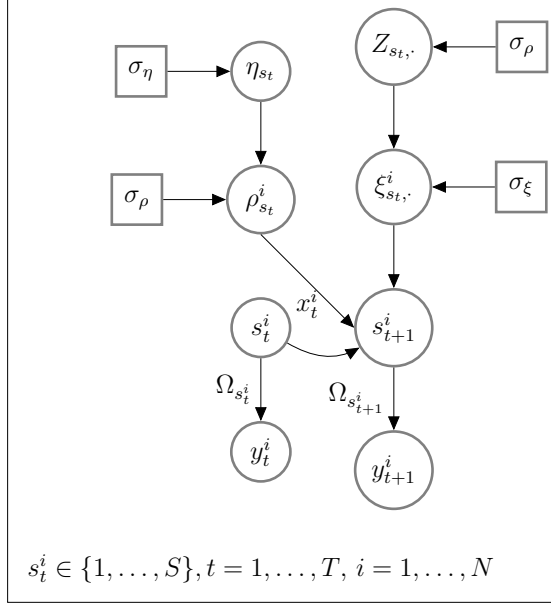


Figure 2.2: Graphical representation of the proposed PIBDFC. The data y_t^i are emissions from a distribution that is parameterized by a precision matrix $\Omega_{s_t^i}$, which encodes the functional connectivity and is determined by the state active at time t : $s_t^i \in \{1, \dots, S$, $t = 1, \dots, T$, $i = 1, \dots, N$. The probabilities of transitions from s_t^i to s_{t+1}^i are given by the (s_t^i, s_{t+1}^i) entry of the $S \times S$ matrix $Q_{\cdot, \cdot, t}^i$. This entry is modeled according to Equation 2.3.

2.1.4 Posterior Inference

The posterior distribution for the parameters in the proposed model is not available in closed form. Hence, I revert to Markov Chain Monte Carlo (MCMC) techniques for posterior inferences. In particular, I follow Holsclaw et al. (2017) and employ Polya Gamma auxiliary variables (Polson et al., 2013) to sample the state transition parameters. Based on the sampled $Q_{\cdot, \cdot, t}^i$, I can construct a sequence of transition matrices based on equation (2.3). After normalizing each row $Q_{s, \cdot, t}^i$ so that it sums to 1, I use a stochastic forward backward algorithm to sample the state sequence (Scott, 2002). Then, conditioned upon the estimated state sequence, it is possible to sample the precision matrix parameters using the blocked Gibbs algorithm presented in Li et al. (2019). Other parameters in the hierarchical model for the states' transitions (2.4) are sampled via simple Gibbs steps. By iterating through the steps above, I obtain samples from the posterior. I provide a brief summary below:

1. **Sample $\mathbf{Q}^i, \xi^i, \rho^i$:** I can rewrite the likelihood for ξ_{rs}^i according to Holmes and Held (2006) to be in the form of Equation 2.6.

$$L(\xi_{rs}^i) \propto \prod_{t:s_{t-1}^i=r} \frac{\exp(\xi_{rs}^i - c_{rst}^i)^{I(s_t^i=s)}}{1 + \exp(\xi_{rs}^i - c_{rst}^i)} \quad (2.6)$$

where $c_{rst}^i = \log \sum_{m \neq s} \exp(\xi_{rm}^i + x_t^i \rho_m^i - x_t^i \rho_s^i)$. Using the Polya-Gamma augmented logistic regression technique of Polson et al. (2013), we get the posterior of ξ_{rs}^i to be conditionally Gaussian.

$$\xi_{rs}^i | \cdot \sim N \left(\frac{Z_{rs}/\sigma_\xi + n_{rsi} - N_{ri} + 2 \sum_{t:s_{t-1}^i=r} \omega_{rst}^i c_{rst}^i}{\sum_{t:s_{t-1}^i=r} \omega_{rst}^i + 1/\sigma_\xi}, \left(\sum_{t:s_{t-1}^i=r} \omega_{rst}^i + 1/\sigma_\xi \right)^{-1} \right)$$

where n_{rsi} is the count of transitions from state r to state s during the timecourse of subject i and N_{ri} is the number of times subject i visited state r . ω_{rst}^i is a Polya-Gamma random variable distributed $PG(1, \xi_{rs}^i - c_{rst}^i)$. I use a similar strategy to update ρ_{rb}^i , the logistic component for subject i for state r and covariate b , achieving the posterior:

$$\xi_{rb}^i | \cdot \sim N \left(\frac{\eta_{rb}/\sigma_\rho + \sum_{t=1}^{T_i} x_{tb}^i (I(s_{t+1}=r) - 1/2 + \omega_{rbt}^i c_{s_t r b}^i)}{\sum_{t=1}^{T_i} (x_{tb}^i)^2 \omega_{rbt}^i + 1/\sigma_\rho}, \left(\sum_{t=1}^{T_i} (x_{tb}^i)^2 \omega_{rbt}^i + 1/\sigma_\rho \right)^{-1} \right)$$

where $c_{rst}^i = \log \sum_{m \neq s} \exp(\xi_{rm}^i + x_t^i \rho_m^i - x_t^i \rho_s^i)$.

2. **Sample \vec{s}_t :** I sample the sequence of states adapting the stochastic forward backwards algorithm presented by (Scott, 2002).
3. **Sample the matrices Ω_s^i , $s = 1, \dots, S$:** The conditional posterior for Ω_s is as follows:

$$P(\Omega_s | \mathbf{Y}, s, \lambda_{..s}, \tau_s) \propto \prod_{\{i, t: s_t^i=s\}} N_R(Y_t^i | 0, \Omega_s^{-1}) \prod_{j=2}^R \prod_{i=1}^j N(\omega_{ijs} | 0, \lambda_{ijs} \tau_s)$$

For MCMC inference purposes, Li et al adopt auxiliary variables ν_λ and ξ_τ , in order to modify the gibbs sampling procedure presented by Makalic and Schmidt (2016a). This procedure is performed for a column wise update in a fashion similar to Wang (2012). For each state, I update Ω_s by following the Graphical Horseshoe algorithm letting $S = n_s * \hat{\Sigma}_s$ where n_s and $\hat{\Sigma}_s$ are the sizes and sample covariance matrices of observations assigned to state s .

4. **Sample \mathbf{Z}, η :** These conditional posteriors follow the typical normal-normal update:

$$Z_{rs}|\cdot \sim N\left(\left(\frac{1}{\sigma_z} + \frac{n}{\sigma_\xi}\right)^{-1}\left(\frac{z_{rs}^0}{\sigma_z} + \frac{\sum_i \xi_{rs}^i}{\sigma_\xi}\right), \left(\frac{1}{\sigma_z} + \frac{n}{\sigma_\xi}\right)^{-1}\right),$$

$$\eta_b|\cdot \sim N\left(\left(\frac{1}{\sigma_\eta} + \frac{n}{\sigma_\rho}\right)^{-1}\left(\frac{\eta_b^0}{\sigma_\eta} + \frac{\sum_i \rho_b^i}{\sigma_\rho}\right), \left(\frac{1}{\sigma_\eta} + \frac{n}{\sigma_\rho}\right)^{-1}\right),$$

2.1.5 Graph Selection

Our model achieves sparsity of the estimated functional network thanks to the shrinkage of the off-diagonal elements of Ω provided by the graphical horseshoe prior. However, shrinkage priors do not lead to exact zeros. Hence, non-relevant connectivities need to be identified through post-MCMC inference. For example, Li et al. (2019) suggest using 50% posterior credible intervals of the inverse-covariance elements, and then thresholding the off-diagonal element to zero if the interval contains 0, reporting the posterior mean otherwise. However, the resulting selection may be too liberal, as it does not provide a multiplicity control for false discoveries.

I follow a decision theoretic approach and formulate the graph selection problem as a testing problem based on the posterior evidence of shrinkage for each off-diagonal element of the precision matrix Ω_s . Since I consider the posterior estimates of Ω_s for each state $s = 1, \dots, S$ separately, in the following, I drop the superscript s for notational simplicity, unless needed

for clarity. For any given state $s = 1, \dots, S$, the j, k off-diagonal element ω_{jk} ($j < k; k = 2, \dots, R$) provides a measure of the connectivity level, with $\omega_{jk} = 0$ indicating that the connectivity is truly zero, and $|\omega_{jk}| \neq 0$ otherwise. Let δ_{jk} indicate the decision (action) in the testing problem, that is $\delta_{jk} = 1$ corresponds to rejecting the null hypothesis of no-connectivity and $\delta_{jk} = 0$ failure to reject (acceptance). Let $D = \sum_{j < k} \delta_{jk}$ indicate the total number of positive (significant) decisions taken. Following Müller et al. (see 2007), for real numbers $c_1, c_2 > 0$, I can then determine the optimal set of decisions $\boldsymbol{\delta} = \{\delta_{12}, \delta_{13}, \dots, \delta_{R-1R}\}$ by minimizing the following loss function:

$$L_{\Omega_s}(\Omega_s, \boldsymbol{\delta}, \mathbf{Y}) = - \sum_{j < k} \delta_{jk} |\omega_{jk}| + c_1 \sum_{j < k} (1 - \delta_{jk}) |\omega_{jk}| + c_2 D.$$

The loss function compounds a reward for correct decisions (true positives), provided by the first addend, $-\sum_{j < k} \delta_{jk} |\omega_{jk}|$, where each correct decision is proportional to $|\omega_{jk}|$'s, and a penalty for false negative discoveries, represented by the second addend, $\sum (1 - \delta_{jk}) |\omega_{jk}|$. The last term encourages parsimony, by increasing the penalty as the number of significant elements increases. The optimal decision is obtained by minimizing the posterior expected loss,

$$E(L_{\Omega_s} | \mathbf{Y}, \tau) = - \sum_{j < k} \delta_{jk} E(|\omega_{jk}| | \mathbf{Y}, \tau) + c_1 \sum (1 - \delta_{jk}) E(|\omega_{jk}| | \mathbf{Y}, \tau) + c_2 D,$$

where $E(\omega_{jk} | \mathbf{Y}, \tau)$ is the posterior mean of the off-diagonal elements of the inverse matrix Ω . The minimizer corresponds to a threshold of the posterior means to identify the non-zero elements of the precision matrix,

$$\delta_{jk}^* = I \{E(|\omega_{jk}| | \mathbf{Y}) \geq c_2 / (1 + c_1)\}.$$

Li et al. (2019) show that the posterior mean is unbiased and it can be represented as a linear function of a shrinkage factor defined by the expected value of the random variable $\kappa_{jk} = \frac{1}{1 + \lambda_{jk}^2 \tau^2}$, which has a Compound Confluent Hypergeometric distribution (Gordy, 1998). More in detail, $E(\omega_{jk} | \mathbf{Y}, \tau) = (1 - E(\kappa_{jk} | \mathbf{Y}, \tau)) \hat{\omega}'_{jk}$ with $\hat{\omega}'_{jk}$ representing the least square

estimate of ω_{jk} , $j < k$. See Theorem 4.1 in Li et al. (2019), and related discussions in Bhadra et al. (2019). The result extends trivially to the folded normal distribution characterizing $|\omega_{jk}|$. Note that $\kappa_{jk} \in (0, 1)$, and larger values of $E(\kappa_{jk})$ indicate stronger shrinkage of the posterior estimates toward zero.

Therefore, graph selection can be conducted by thresholding an estimate $\hat{\kappa}_{jk}$ of the shrinkage factor κ_{jk} , i.e.

$$\hat{\delta}_i^* = I \{ \hat{\kappa}_{jk} \leq \eta \},$$

for some threshold $\eta \in (0, 1)$. For example, in the simple regression case, Carvalho et al. (2010) have previously recommended an informal decision rule thresholding ω_{jk} to 0 if $1 - \hat{\kappa}_{jk} < 0.5$ where $\hat{\kappa}_{jk}$ is the posterior median of κ_{jk} . However, such a rule does not take into account the multiplicity problem induced by the selection of the off-diagonal elements of the precision matrix. The posterior medians $\hat{\kappa}_{jk}$ provide a measure of the evidence in favor of the null hypothesis, $H_0 : \omega_{jk} = 0$.

Hence, a threshold η could be set by controlling a measure of the Bayesian False discovery rate (BFDR, Newton et al., 2004) at a certain level q^* , that is to satisfy the equation

$$BFDR(\eta) = \frac{\sum_{jk} \hat{\kappa}_{jk} I(\hat{\kappa}_{jk} \leq \eta)}{\sum_{jk} I(\hat{\kappa}_{jk} \leq \eta)} < q^*.$$

For a related but different solution to the problem of graph selection, see also Chandra et al. (2021), who consider inference on the partial correlation matrix derived from Ω .

2.2 Simulation Study

In this Section, I present three sets of simulation studies that aim at measuring the performance of our model with respect to the detection of non-zero connectivities and the estimation of the latent state trajectories over time. More specifically, in the first two simulation

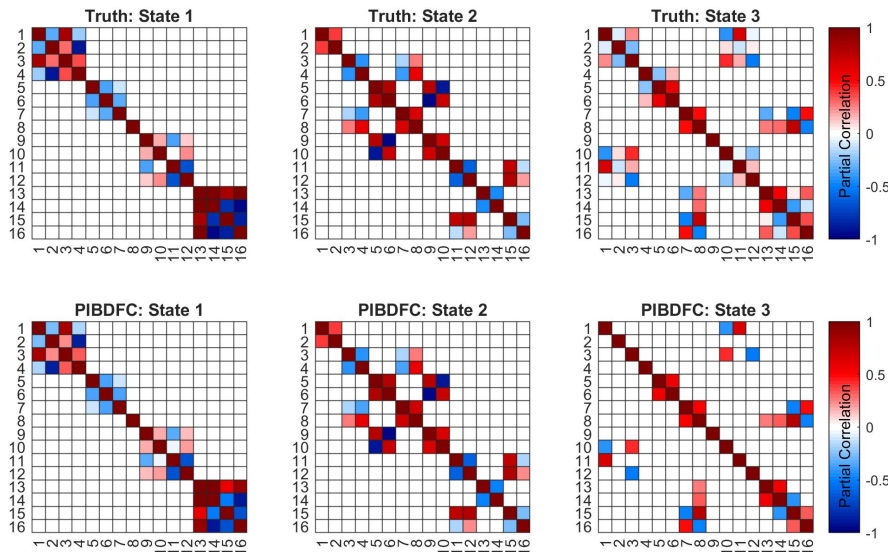


Figure 2.3: Simulation Study 1: *Top*: The true partial correlation matrices for each state responsible for generating the simulation data in the Simulation Study 1. *Bottom*: The estimated partial correlation matrix from the proposed PIBDFC from a single repetition of the simulation. Each estimated partial correlation is the posterior mean of their respective distributions. Cells are set to 0 in post-hoc MCMC by controlling the BFDR at the 0.2 level. See Sections 2.1.5 and 3.3 for details.

studies, I compare the proposed predictor-informed Bayesian dynamic functional connectivity (PIBDFC) model with two alternative models: a widely-used tapered sliding window (Tapered SW) approach, first outlined by Allen et al. (2014), and the Bayesian Dynamic Functional Connectivity (BDFC) model developed by Warnick et al. (2018). The Tapered SW represents a standard approach in the functional connectivity literature, whereas BDFC uses a homogeneous HMM to model latent connectivity state dynamics. However, BDFC provides a model-based estimation of exact zeros in the functional networks at the cost of computational scalability and speed, as opposed to our computationally faster soft-shrinkage based approach. Furthermore, the BDFC does not incorporate any predictor information in the latent state dynamics. Both competing approaches are developed for single-subject inference. I compare to our multi-subject model by concatenating the multi-subject data along the time axis for input into the respective algorithms. All models are run on a Linux computer with an Intel Xeon Gold processor (2x 3.10 GHz) and 4 GB of RAM. For both the

PIBDFC and BDFC, I simulated 5,000 posterior samples after 5,000 burn-in draws. When fitting PIBDFC, I set the hyperparameters $\tau_0 = 1$, $\sigma_\xi = \sigma_\rho = \sigma_z = \sigma_\eta = 0.1$, following the motivations of Section 2.1.2.

I assess the performance for our model as far as states' reconstruction by computing a set of metrics for each latent state separately. Let r_{jk} , $j < k; k = 2, \dots, R$, denote the binary indicator of a non-zero connection between two regions j and k . Following the discussion in Section 2.1.5, let δ_{jk} indicate the decision after model fit. Then I define the *edge true positive rate* (tpr) as $\sum r_{jk}\delta_{jk}/\sum r_{jk}$. Similarly, the *edge true negative rate* (tnr) is defined as $\sum (1 - r_{jk})(1 - \delta_{jk})/\sum (1 - r_{jk})$. The *Edge F1 score* (F1) is the product of the tnr and tpr, and serves as a measure of the overall performance in graph estimation, balancing between the tpr and tnr. Analogously, I define a metric to assess the performance of the model in the estimation of the states' sequences. Let s_t^i indicate the true latent state active at time t for subject i and let \hat{s}_t^i indicate its model estimate. Then, the *state sequence accuracy* for state s is defined as $\sum \{I(s_t^i = s)I(\hat{s}_t^i = s)\}/\sum I(s_t^i = s)$.

Simulation Study 1: In our first study, I investigate the performance of our model in an ideal setting where the data generation process matches the model closely. I set $T = 300$ time points, $R = 16$ ROIs, $N = 30$ subjects, and $S = 3$ connectivity states. In this setting, I simulate data $Y_t^i \sim N_{16}(0, \Omega_{s_t^i}^{-1})$ with $\Omega_{s_t^i}$ encoding the individual conditional independence structure at time t , identified by the value of the state indicator variables $s_t^i \in \{1, 2, 3\}$ and the prespecified graphs in the first row of Figure 2.3. In order to study the effect of a predictor information on the estimation of the transition probabilities and the functional connectivity dynamics, I introduce a single binary time-varying predictor variable, x_t , which transitions from 0 to 1 when $t = \frac{T}{2}$. For each value of the exogenous variable, I set the transition probabilities for the latent state trajectories as follows

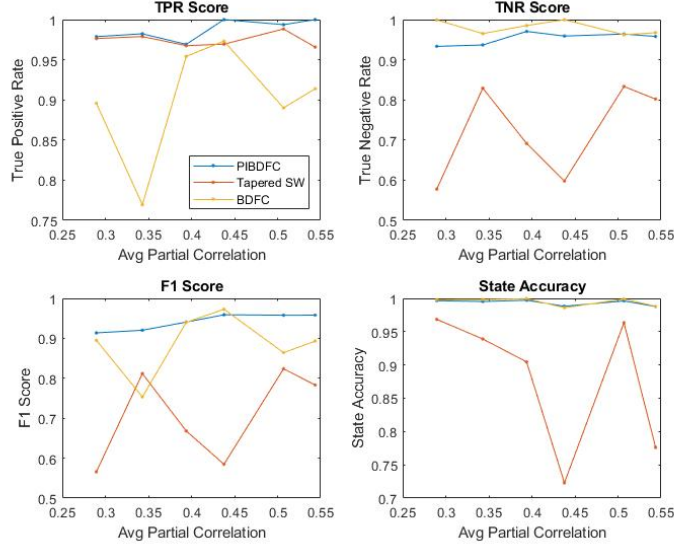


Figure 2.4: Simulation Study 1: True Positive Rate, True Negative Rate, F1 Score, and state accuracy metrics for the PIBDFC, BDFC, and Tapered SW approaches over different assumptions for the signal-to-noise ratio in the correlation structure. Along each horizontal axis is the average strength of the non-zero partial correlations for each state.

$$Q_t = \begin{bmatrix} 0.98 & 0.02 & 0 \\ 0.1 & 0.9 & 0 \\ 0 & 0.5 & 0.5 \end{bmatrix} \text{ when } x_t = 0 ; Q_t = \begin{bmatrix} 0 & 0.5 & 0.5 \\ 0 & 0.7 & 0.3 \\ 0 & 0.02 & 0.98 \end{bmatrix} \text{ when } x_t = 1.$$

Therefore, for each subject, the state sequence enforces transitions between states 1 and 2 for the first half of the time series, whereas it enforces transitions between states 2 and 3 for the second half. I then simulate different state sequences for each subject using equation (2.3), and replicated the process over 30 independent simulation data sets. In order to assess the performance of the methods for different levels of signal-to-noise ration, I repeated the simulation experiment using different precision matrices $\Omega_s, s = 1, 2, 3$ of the same structure of the top row of Figure 2.3 but allowing for different values of the non-zero entries. This is done by using the *sprandsym* function from the Mathematics toolbox of Matlab. This function takes in an adjacency matrix representation of a graph, $A_s \in \mathcal{R}^{R \times R}$ where $A_{ijs} = I(\omega_{ijs} \neq 0)$, and outputs a positive definite matrix with the same placement of 0's but




Figure 2.5: Simulation Study 1: *Top*: The true state transition path for each subject (vertical axis) across each time point (horizontal axis). The color in each cell identifies which precision matrix in Figure 2.3 generated the simulated the data for each subject-time point pairs. *Bottom*: The *maximum a posteriori* estimated state trajectories from PIBDFC.

random non-zero entries. This output matrix is then normalized to a partial correlation matrix. Thus, I obtained a total of six sets of precision matrices to learn the structure of. I show the aggregated results in Figure 2.4. The horizontal axis reports the average strength of the non-zero partial correlations for each of the six sets of precision matrices, indicating a level of signal strength. The PIBDFC consistently performs better in connectivity estimation with regards to F1 score, displaying the best balance of finding true non-zero partial correlations while controlling for false positives. A table representing the results with standard deviations computed over the replicated experiments is shown in the Appendix.

In the following, I illustrate the inferential analyses enabled by the proposed PIBDFC approach by showcasing a single replicate. In Figure 2.3 (bottom row) I show how the PIBDFC is able to recover the true conditional independence structure underlying the data generation process by estimating the partial correlations between regions and enforcing the true 0's through the BDFR approach devised in Section 2.1.5. The model is also able to recover the most likely state transition sequence for each subject, as determined by the *maximum a*




Figure 2.6: Simulation Study 1: Estimation of the connectivity change points in a representative subject. The horizontal axis indicates the time-points while the vertical axis reports the posterior probability $P(s_t^1 \neq s_{t-1}^1 | Y_{1:T}^i)$. The posterior probabilities of a change point are in red, whereas the black spikes represent the true change points for the subject. I also display a horizontal dotted line at 0.95 to reflect the informal rule of declaring a change-point if $P(s_t^1 \neq s_{t-1}^1 | Y_{1:T}^i) > 0.95$.

posteriori state estimate at each time point. See Figure 2.5. It is also important to assess the ability of the method to identify true change points in the connectivity structure. Figure 2.6 reports the estimated connectivity change points for a representative subject. PIBDFC is able to estimate the state sequence well while tying together the increased rate of appearance of state 3 when the stimulus changes from 0 to 1 halfway through the simulated experiment. All models were compared in terms of computation time as reported in Table A.2. PIBDFC is also able to draw as many posterior draws in a third of the computation time.

Simulation Study 2: In this second simulation study, I measure the performance of our approach with synthetic data that are similar to real fMRI data. More specifically, I use the simulation toolbox of Erhardt et al. (2012) and follow the simulation approach of Warnick et al. (2018). The SimTB toolbox implements a canonical hemodynamic response function (Lindquist et al., 2009), defined as a linear combination of two gamma functions, to simulate fMRI time series. This function is then convolved with a box stimulus function where gaussian noise with variance = 0.01 is added. Functional connectivity is then obtained by

predefining cliques, i.e. clusters of regions, that have signal (here, 0.5) added to or subtracted from all regions in the clique simultaneously at random time points within a connectivity state. This induces correlation while having non-gaussian errors. I then simulate the state sequence over $T = 150$ time points with x_t being 0 for the first 75 timepoints and 1 for the last 75 among all subjects. Similar to Simulation Study 1, I use the exact same Q_t among all subjects. I repeat this process for $N = 30$ subjects over 30 simulation replicates.

Metric	Method	State 1	State 2	State 3
Edge TPR	PIBDFC	1 (0)	0.8290 (0.032)	0.7652 (0.039)
	Tapered SW	1 (0)	1 (0)	1 (0)
	BDFC	0.9769 (0.070)	0.9014 (0.156)	0.7203 (0.189)
Edge TNR	PIBDFC	0.9278 (0.004)	0.8604 (0.041)	0.9250 (0.040)
	Tapered SW	0.3286 (0.109)	0.4583 (0.165)	0.2500 (0.157)
	BDFC	0.8294 (0.150)	0.8552 (0.148)	0.9531 (0.088)
Edge F1 Score	PIBDFC	0.9278 (0.004)	0.7134 (0.045)	0.7083 (0.055)
	Tapered SW	0.3286 (0.109)	0.4583 (0.165)	0.2500 (0.157)
	BDFC	0.8063 (0.138)	0.7717 (0.192)	0.6822 (0.176)
State Acc	PIBDFC	0.8526 (0.029)	0.7507 (0.022)	0.7727 (0.022)
	Tapered SW	0.7199 (0.175)	0.4133 (0.100)	0.6342 (0.108)
	BDFC	0.6110 (0.43)	0.7181 (0.11)	0.5541 (0.37)
Comp Time (min)	PIBDFC	161.23 (29.493)		
	Tapered SW	1.9241 (0.31)		
	BDFC	500.57 (18.11)		

Table 2.1: Simulation Study 2: results over 30 repetitions. I report sensitivity and specificity metrics for the estimated graphs of the corresponding states, together with the overall accuracy of the estimated state sequences. Standard deviations across the 30 simulations are showed in brackets. The proposed method maintains the best balance between sensitivity and specificity as well as latent state estimation accuracy.

In Table 2.1 I show the results to the application on the SimTB data. The PIBDFC still well detects the connectivities between the simulated regions despite a mis-specified likelihood. The performance in both graph estimation and state estimation slightly appears to decline slightly in comparison to the Gaussian Simulation 1 setting, which is expected. The Tapered SW approach suffers from low specificity. Compared to the standard HMM of BDFC, the proposed PIBDFC is slightly more capable at detecting changes in state transitions, thus improving graph estimation performance as a result. This result is likely due to the distortion of changes to the partial correlation by the convolution with the hemodynamic response function. In this setting, the covariate information becomes more relevant in helping the model identify changes in the state transition behavior. The computational time is also quite favorable compared to approach of Warnick et al. (2018), despite allowing for individual differences in state dynamics among the 30 subjects.

Simulation Study 3: In this simulation study, I compare the performances of our model and the Connectivity Change Point Detection (CCPD) model of Kundu et al. (2018), on the basis of edge- and change-point detection. Contrary to our model, the CCPD model employs a two-stage approach for estimating dynamic functional connectivity. In the first stage, the model learns the number and locations of the change points from all available subjects' data. In the second stage, a graphical lasso approach is applied independently to the time scans between two change points. Since the CCPD model assumes that every change point occurs at the same time for each subject, in order to fairly compare the two methods, I simulate data under the CCPD assumption of common change points. More specifically, I set $T = 300$ and generate $Y_t^i \sim N(0, \Omega_{s_t})$ where s_t varies across the following sequence of states: $\{1, 2, 3, 1\}$ switching at $t = 75, 150, 225$, for a total of 3 change-points overall. I use the same true partial correlation matrices to generate the data as in Simulation study 1. For the PIBDFC, a time point t for subject i was judged to be a change point if $P(s_t^i \neq s_{t-1}^i | Y_{1:T}^i) > 0.95$. PIBDFC does not assume common change points and as a result, will not output common

change points across individuals; therefore, I report the average the number of change points across all subjects.

Method	PIBDFC			CCPD		
	1	2	3	1	2	3
Edge TPR	0.9650 (0.02)	1.0000 (0)	0.9867 (0.01)	0.9333 (0.02)	1.0000	0.9800 (0.02)
Edge TNR	0.9674 (0.01)	0.9719 (0.01)	0.9615 (0.02)	0.9733 (0.09)	0.9978 (0.01)	0.7719 (0.06)
Edge F1 Score	0.9336 (0.02)	0.9719 (0.01)	0.9486 (0.02)	0.9078 (0.08)	0.9978 (0.01)	0.7564 (0.06)
Num ChgPts (3)	3.8 (0.97)			3.1 (0.38)		

Table 2.2: Simulation Study 3: Results over 30 repetitions. I show the entry wise true positive and true negative rates for the estimated graphs for the corresponding states. I also show estimated number of the estimated number of change points. PIBDFC performs comparably to CCPD in the setting where changepoints are common among subjects despite no explicit assumption of this being the case.

In Table 2.2, I show the results of the comparison between PIBDFC and CCPD under a shared change point model. CCPD is indeed able to accurately detect the number of change points and the resulting graph structure in each partition well. By thresholding the posterior probability of a change point, our model tends to over-estimate the number of change points on average, a result of the model sometimes estimating very sudden changes of state for a brief collection of time points in some subjects. As a contrast, in simulation studies 1 and 2, the change points are generated from a process that truly followed a hidden Markov model, leading to more accurate estimates of the occurred change points. By leveraging on the assumption of common change points, the two-stage CCPD model can achieve increased accuracy, while our model allows for the incorporation of individual transitions and covariates in the transition probabilities.

2.3 Case Study

I apply the proposed PIBDFC model to the motivating dataset, where fMRI data have been collected concurrently with pupillometry measurements. In our application, I demonstrate how the model can recover expected change points in time-varying functional connectivity states, as those states align quite well with the experimental events regulated by the behavioral task. I am also able to estimate the effect of pupil dilation on the subjects' propensity to change states.

2.3.1 Experimental design and data collection

In this experiment, subjects performed a handgrip task adapted from Mather et al. (2020). Thirty-one participants (18 females, mean age 25 years \pm 4 years) enrolled in this study at the University of California, Riverside Center for Advanced Neuroimaging, but one was excluded due to a history of attention deficit hyperactive disorder and consumption of related medication resulting in a total $N = 30$. All subjects provided written informed consent to participate, and received monetary compensation for their participation. The study protocol was approved by the University of California, Riverside Institutional Review Board (IRB). Magnetic resonance imaging (MRI) data were collected on a Siemens 3T Prisma MRI scanner (Prisma, Siemens Healthineers, Malvern, PA) with a 64 channel receive-only head coil. fMRI data were collected using a 2D echo planar imaging sequence (echo time (TE) = 32 ms, repetition time (TR) = 2000 ms, flip angle = 77° , and voxel size = $2 \times 2 \times 3 \text{ mm}^3$, slices=52) while pupillometry data were collected concurrently with a TrackPixx system (VPixx, Montreal, Canada).

All subjects underwent a 12.8-minute experiment in which they alternated between six resting state blocks and five squeeze blocks. In the squeeze blocks, they brought their dominant hand

to their chest while holding a squeeze-ball (Mather et al., 2020). The five squeeze blocks lasted 18 seconds while the interspersed six resting state blocks had durations of five-, two-, two-, five-, one-, and one-minute, respectively.

All subjects underwent two sessions: one where they executed the squeeze at maximum grip strength (active session), and one where they still brought their arm up to their chest but were instructed simply to touch the ball and not to squeeze it (sham session). The functional data underwent a standard preprocessing pipeline in the fMRI module of the brain software library (FSL). The pipeline consisted of slice time correction, motion correction, susceptibility distortion correction, and spatial smoothing using a kernel Gaussian smoothing factor set at a full width half maximum of 0.8475 (Smith et al., 2004; Woolrich et al., 2009). Finally, all data were transformed from the individual subject space to the Montreal Neurological Institute (MNI) standard space using standard procedure in FSL (Smith et al., 2004; Woolrich et al., 2009).

Pupillometry data were collected during the scans, using a sampling rate of 2kHz, pre-processed using the ET-remove artifacts toolbox (github.com/EmotionCognitionLab/ET-remove-artifacts), and downsampled to match the temporal resolution of the fMRI data (Mather et al., 2020). To measure pupil dilations relative to baseline, the dataset was normalized by dividing by subject-specific means of the first five-minute resting state block (prior to any squeeze or hand-raising), leading to percent signal changes. Three subjects' data were discarded due to technical difficulties during acquisition of pupil dilation measurements, resulting in $N = 27$ for all concurrent pupillometry and functional data related analyses.

Since I used a pseudo-resting state paradigm, our interest was focused on five networks of interest that have all been associated with resting state and have been related to attention in some manner. Default mode network (DMN; a resting state network) and dorsal attention network (DAN; an attention network) were selected because squeezing ought to invoke a

transition from the resting state into a task-positive state (Greicius et al., 2004). The fronto-parietal control network (FPCN) was chosen because it is linked to DAN and regulates perceptual attention (Dixon et al., 2018). Salience network (SN) was selected because it determines which stimuli in our environment are most deserving of attention (Mather et al., 2020; Menon and Uddin, 2010). Talariach coordinates for regions of interest (ROIs) within DMN, FPCN, and DAN were taken from Deshpande et al. (2011) and converted to MNI coordinates while SN MNI coordinates were taken directly from Raichle (2011) (Deshpande et al., 2011; Laird et al., 2005; Lancaster et al., 2007; Raichle, 2011). Two ROIs from FPCN (dorsal anterior cingulate cortex and left dorsolateral prefrontal cortex) were excluded due to their close location to other ROIs. The locus coeruleus (LC) was localized using the probabilistic atlas described in Langley et al. (2020). Blood oxygen level dependent (BOLD) signal from each voxel within an ROI were extracted and averaged to represent the overall signal for an ROI. I eventually considered 31 total ROIs: 9 from DMN, 7 from FPCN, 6 from DAN, 7 from SN, and 2 from LC. The MNI anatomical coordinates for the four attention networks and LC were used to center a 5 mm³ isotropic sphere (Deshpande et al., 2009; Stilla et al., 2007). See the Appendix for a list of the ROIs and corresponding MNI stereotaxic space coordinates and networks.

2.3.2 Model fitting

The 31 ROIs described above formed the vectors of BOLD responses $Y_t^i = (Y_{t1}^i, \dots, Y_{t31}^i)$ measured on subject $i = 1, \dots, 27$ at time t , for $t = 1, \dots, 1050$. I also included concurrently recorded pupillometry data as a proxy for the quantify the effect of LC engagement on the dynamics of functional connectivity (Joshi and Gold, 2022).

I fit our model considering several assumptions for the total number of states, $S = 3, 4, 5, 6$. However, when assuming more than 3 states, the fit simply degenerated to 3 states in the




Figure 2.7: Real Data Analysis: the estimated connectivity networks for the ROIs. Nodes represent ROIs and the edges denote the partial correlations between the connected nodes. The edge colors correspond to the directionality of the partial correlations and the width corresponds to the magnitude. Node colors identify clusters of regions into *a priori* defined networks. See Section 2.3.3 and Table A.1 in the Appendix

posterior inference, with no observations assigned to additional states. Thus, I present the model specification for 3 states with the following settings for the hyperparameters in (2.2). I set the group level baseline relative transition prior means $z_{rr}^0 = 2$ for $r = 2, 3$ while all other elements of z_{cdot}^0 are set to 0. I also set the prior spread of the baseline transitions and pupillary effects $\sigma_z, \sigma_\eta = 0.05$. This combination of setting is used to encourage self transitions, as they correspond to preferring smoother state sequences *a priori* among all subjects. I set the prior variability of the subject-level transition parameters around the group-level transition parameters, by choosing $\sigma_\xi, \sigma_\rho = 0.1$, therefore capturing individual differences between subjects on the log-odds of transitioning between states. Lastly, τ_0 , the hyperparameter informing prior knowledge of connectivity network sparsity, is set to 1, as this value corresponds to a prior distribution with a high spread over edge densities, as seen in Figure 2.1.

2.3.3 Inference results

Figure 2.7 plots the estimated connectivity networks for each of the three states. Nodes represent ROIs and the edges identify the non-zero partial correlations between the connected nodes. The edge colors correspond to the directionality of the partial correlations and the width corresponds to the magnitude. The dotted colors in the nodes identify clusters of regions within a *a priori*, knowledge-based, neuroscientific networks (from the top right section in counter-clockwise order): Default Mode Network (DMN), Frontal Parietal Control Network (FPCN), Dorsal Attention Network (DAN), Saliency Network (SN), and Locus Coeruleus (LC).

Figure 2.8 shows the *maximum a posteriori* (MAP) estimated state sequences from our model for all 27 subjects. The subjects' rows are ordered by similarity of the estimated state trajectories as captured by a hierarchical clustering using Euclidean distance.

By inspecting Figure 2.7, it is apparent that state 1 shows relatively sparser connectivity than the other two states. In state 1, I can see strong bilateral connectivity among homologous regions in the left and right hemispheres, as well as several nodes in FPCN (dark blue) showing strong connectivity with multiple nodes in SN (light red); likewise, several nodes in DMN (dark red) show connectivity with SN (light red) nodes. There is almost no presence of anti-correlation. The dominance of SN connectivities together with both DMN and FPCN suggests that arousal may be up-regulated in this state. Indeed, Figure 2.8 suggests that state 1 occurs predominantly during the 'squeeze' periods of the behavioral task, when subjects either squeezed the squeeze ball or held it to their chest. This observation suggests that our model was able to detect those objectively-definable events in the time-series of this experimental dataset.

In state 2, I see a quite different pattern: weaker average connectivity when compared to state 1, but also many more of these weaker connections both within-network and between

networks. In addition to relatively ubiquitous within-network connections within FPCN (dark blue) and DMN (dark red), state 2 is characterized by cross-network connectivity – and anti-connectivity – between DMN and FPCN. Interestingly, these parallel some of the strongest connectivities from state 1. The relatively occupancy in state 2 appears higher in the active condition (Figure 2.8, right half) than the sham condition (Figure 2.8, left half), suggesting subjects tended to occupy this relatively strong, broadly-connected state more often when periodically engaging in actively squeezing the squeeze ball.

The strongest connections in state 3 deviate from those identified in states 1 and 2. There is weaker overall connectivity than state 1, but the connections are stronger and sparser (fewer connections) than state 2. I do again see many within-network connections, as well as relatively strong connections between nodes in FPCN (dark blue) and SN (light red), and also again between DMN (dark red) and SN (light red). However, I also see many more connections with SN from DAN (light blue) than in either of the other two states. I can therefore characterize this state as more sparsely connected than state 2 but still with broad connectivity, which is also consistent with the differences visually apparent in this state between active and sham conditions (right and left halves of Figure 2.8): this state traded off with state 2 for relative percentage occupancy across the subjects.

Finally, I investigate how pupillary dilation modulates state transitions. Figure 2.9 provides the posterior distribution of the group (e^η , left) and individual (e^ρ) effects of pupil dilation on state dynamics. I start by assessing the relationship between between pupil dilation and state transitions for the group. Based on our findings, a 1% increase in pupil dilation relative to baseline is associated with a 31.4% (95%CI : 29.7% – 32.9%) decrease in the odds of transitioning to state 2 and a 34.9% (95%CI : 33.3% – 36.4%) decrease in the odds of transitioning to state 3, in comparison to remaining in the baseline state (state 1). This result is coherent with the findings outlined above, since increased pupil dilation (a proxy for increased arousal/effort) appears associated with transitioning toward the less densely




Figure 2.8: Real Data Analysis: Estimated states’ transition path for each subject. The horizontal axis indicates the TR with vertical dotted lines indicating portions where the subject raises their arm. Subject sequences are aligned so that the left half of the plot shows sequences from the sham condition and the right half shows sequences from the active condition. The vertical axis display the subject indices, ordered by similarity in state trajectory according to a hierarchical Euclidean distance based clustering of their MAP transition behavior.

connected connectivity structure of state 1, dominated by edges between SN and both DMN and FPCN. I should note the causal direction of the associations between arousal, pupil dilation, and the connectivity structure of state 1 can not be investigated by this model.

Further inspection of the right column of Figure 2.9 shows that the posterior distributions of the individual effects of pupil dilation e_i^p is decidedly below 1 for all subjects, i.e. the association between increased pupil dilation and state 1 holds for all subjects measured. Subjects are ordered along the horizontal axis according to their similarity in state trajectories obtained from a hierarchical Euclidean distance based clustering (similarly as in Figure 2.8. The horizontal dashed line represents the posterior mean from the group estimate in the right panel. It is interesting to note the differing clusters when comparing the posterior distributions of e_i^{p2} to e_i^{p2} : trending downwards and upwards respectively. Combining Figure 2.9 with Figure 2.8, it appears that differences in state trajectories between subjects lie

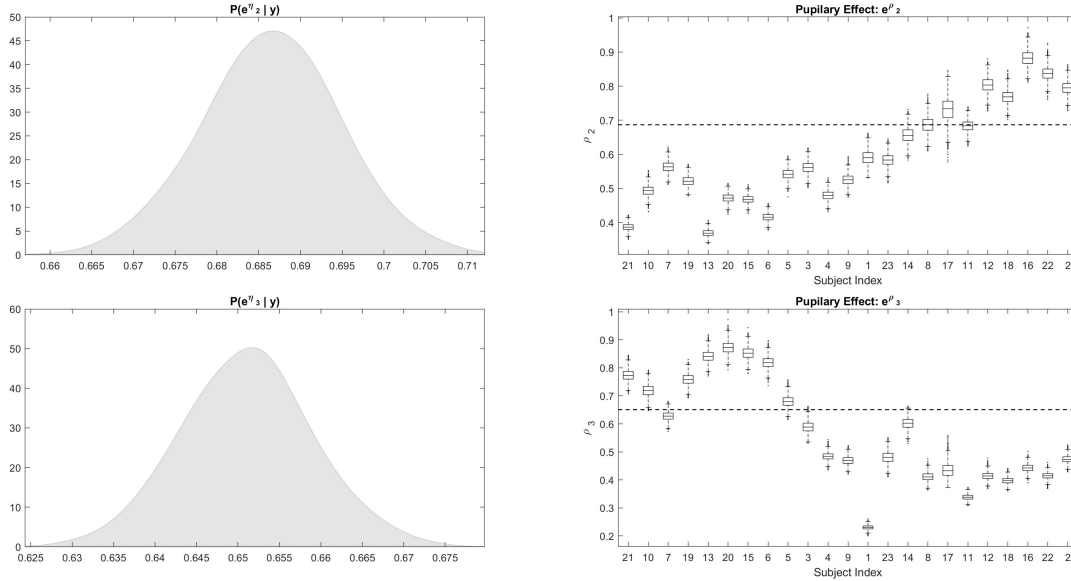


Figure 2.9: Real Data Analysis: The posterior distribution of the group effect of pupillary dilation e^η (left), and individual effects of pupillary dilation e^ρ . Rows indicate the propensity for transitioning into states 2 and 3 respectively. For the individual effects, subjects are identically ordered as in Figure 2.8. The horizontal dotted line is the posterior mean for the group level effects, $\eta_2 = 0.687$ and $\eta_3 = 0.651$ respectively.

in state occupancy when pupil dilation is not higher than the reference, despite all subjects tending to transition to state 1 during when raising their arm.

More specifically, subjects clustered in the first half of Figure 2.9 (right) tend to occupy state 3 during non-squeeze sections and so are even more likely to transition away from state 2 during periods of high pupil dilation. Similarly, subjects in the second half of the the Figure tend to occupy state 2 during non-squeeze sections, and are thus very likely to transition away from state 3. This heterogeneity is important as it provides a more thorough understanding of the relationship between increased pupil dilation and transitions toward different cognitive states.

2.4 Discussion

I have proposed a multi-subject Bayesian approach for estimating time varying functional connectivity where the brain network state transitions are dynamically informed by concurrently-recorded subject specific covariates. The proposed method allows for group level and subject level inferences on the heterogeneous effects on the connectivity dynamics of the recorded time-varying covariates. I have applied our model to multi-subject resting state fMRI data with pupillary physiological data and I have shown associations between pupil dilation and strengthened connectivity between the SN brain regions with both the FPCN and DMN. This association coinciding with subject arm-raising/squeezing suggests SN connections with both FPCN and DMN are associated with subject arousal.

While I focused here on covariates that were concurrently recorded on each subject, our model can also incorporate covariates that are subject-specific and not time-varying. For example, demographic information may be added to the regression terms in (2.2)–(2.3) and inform subject-specific transition probabilities to describe individual variability over the entire fMRI experiment.

Our model assumes a maximum number of states S to be pre-specified *a priori*. In our application, only a subset of the S available states was visited. However, more in general, the number of states could be learned by assuming a Bayesian-nonparametric specification where the number of functional connectivity states is learned directly from the data (see, for example, Beal et al., 2002; Fox et al., 2011). However, the computational complexity of the inferential algorithm would increase considerably. Variational Bayes approaches could be implemented to obtain approximate inference on the network connections.

Finally, the individual connectivity patterns could be associated with clinical or behavioral outcomes, e.g., to examine the individual heterogeneity of responses to treatments. A two-stage scalar-on-image approach can be devised where the posterior means of the precision

matrices are obtained from our model in the first stage and then used as predictors to investigate the association with the outcome in the second stage. These directions of research will be the object of future investigations.

Chapter 3

Time-Varying Psychophysiological Interactions

Psychophysiological interaction (PPI) models have been proposed by Friston et al. (1997) to study task modulated seed-based brain connectivity in fMRI studies. More in detail, in task based fMRI, it is of interest to investigate the effect modification of external stimuli (called psychological variables of the task design) onto the association between a seed region and other regions of interest (ROI; also called physiological variable). From a statistical perspective, the PPI model typically takes the form of a linear regression model, where the BOLD signal of the seed region is assumed to be an additive linear function of the BOLD time series of other ROIs plus an interaction term, which captures the effect modification of the task via a product of the ROI times series and the stimulus function. If the coefficient of an interaction term is significantly different from zero, after taking into account the main effects of all the psychological and physiological variables, then it implies a task-dependent connectivity between the two brain regions. The original model by Friston et al. (1997) considered a single task condition. More recently, (McLaren et al., 2012) have proposed a *generalized PPI* approach that includes interaction effects for more than two conditions.

Their framework directly models each task condition with respect to all other conditions in the regression framework. The PPI effects are then compared between the conditions of interest after the fit, rather than as in the original formulation by Friston et al. (1997), where the contrast information was used to define a psychological regressor.

In this manuscript, we propose a Bayesian characterization of the generalized PPI model in McLaren et al. (2012) and we further extend it in two important directions. On the one hand, popular implementations of the PPI model assume that the partial correlation between the seed region and the ROIs is static in the absence of a stimulus. However, current developments in neuroimaging suggest that functional connectivity is dynamic, that is the coupling between brain regions may be better characterized by possibly time-varying levels of correlated or mutually informed activity (Iraji et al., 2019; Calhoun et al., 2014). Therefore, our first contribution is to extend the generalized PPI model to include a *dynamic process* which allows for estimating task-modulated time-varying background functional connectivity during the course of an fMRI experiment. Since in our model the proposed time-varying background effect is present during either the resting-state and task conditions, our framework effectively allows to characterize time-varying functional connectivity due to both task and spontaneous fluctuations of neural activity over time. To model the dynamics of the background regression coefficients, we leverage a time-varying scale-mixture shrinkage prior recently proposed for variable selection in sparse Bayesian state-space models (Cadonna et al., 2020).

Seed-based FC analysis investigates the association of the fMRI signal from specific regions and the fMRI signals from any other region in the brain. The seed region is often selected based on the available prior literature or a testable research hypothesis. The primary benefit of seed analysis is that typically both the calculation and the interpretation of the neuroscientific findings are straightforward. However, if the seed region changes, the results will likely vary; hence, the dependence on the choice of seed regions can impact the reproducibil-

ity and generalizability of the findings (Lv et al., 2018). A possible solution is to rotate the seed regions across the brain and implement separate PPI models for each seed. However, testing for non-zero PPI effects in this circumstance is guaranteed to find spurious non-zero PPI effects as the false positive rate is greatly inflated. Also, a Bonferroni correction will be much too conservative, decreasing the ability to detect true correlations and PPI effects.

The second contribution of our manuscript is to propose a viable multicomparison procedure for identifying significant associations between brain regions independently of the choice of the seed region. First, separate PPI models are run, with each ROI serving as the seed region. Since it can be parallelized among numerous computer nodes, this task is computationally efficient. Then, we aggregate findings from each model and identify significant partial correlations, by adapting a non-marginal decision-theory-based multicomparison framework recently developed by Chandra and Bhattacharya (2019). This multiple testing procedure allows to control the rate of false positive decisions while maintaining the ability of the model to detect true non-zero correlations between regions.

The rest of the manuscript proceeds as follows. In Section 3.1, we introduce the proposed dynamic PPI model, with the shrinkage-inducing priors for the dynamic coefficients described in Section 3.1.1. Section 3.2 describes how to obtain posterior inference from the Bayesian model, with Section 3.2.2 detailing the proposed multicomparison procedure to combine inference across different seed-based analyses. Sections 3.3 and 3.4 illustrate the properties of the model in a simulation analysis and an application to data from a serial reaction time experiment. Finally, Section 3.5 concludes with some discussion of the advantages and limitations of our proposal, together with plans for future investigation.

3.1 A Time-Varying PPI Model.

In this Section, we introduce the proposed Bayesian dynamic PPI model, which considers as outcome the fMRI signal of a specific time-series for a region of interest, assumed as the seed region. In our application, we run the model varying the choice of seed region across all the ROIs in the whole brain, and then we synthesize the results of the inference via the multicomparison procedure described in Section 3.2.2. Here, we illustrate the model for the BOLD signal of a generic seed ROI. More specifically, we assume that we have available fMRI time-series from R ROIs across the whole brain, and we let $Y_r(t)$ be the BOLD signal for a seed ROI, $r \in \{1, \dots, R\}$, observed at time points $t = 1, \dots, T$. Following the general PPI framework, we collect the BOLD signals of every other ROI, *except* for the seed region r , into a vector $\mathbf{Z}_t^r = \{Y_d(t)\}_{d \neq r}$, of dimension $R - 1$. Let s_{tk} indicate a stimulus function (more generally, a psychological variable of the task design) at time t , with $k = 1, \dots, K$, $K \geq 1$. The observed BOLD signal Y_{tr} can be written as a noisy realization of a convolution of the underlying neural signal y_{tr} and its hemodynamic response function (hrf), that is,

$$Y_r(t) = \int_{-\infty}^t h_r(t) y_r(t - \tau) d\tau, \quad (3.1)$$

where $h_r(t)$ represents the value of the hrf in the ROI r at time t . It is well established in the neurosciences that the shape of the hrf can vary between regions and subjects, and that inaccurate estimation of the shape of the hrf can lead to spurious non-zero correlations between ROIs (Gitelman et al., 2003). Consistent and efficient estimation of the hrf is an on-going challenge within the context of PPI analysis. For the purpose of this manuscript, we assume a canonical hemodynamic response function, i.e. the hrf is defined as the difference of two gamma functions (Lindquist et al., 2009). The latent neural signals y can be obtained from the deconvolution of the observed BOLD signal with the assumed hrf, using either the Weiner deconvolution (Wu et al., 2013) or latent feature encoding (Bush and Cislser, 2013). The proposed dynamic psychophysiological interaction (dPPI) model posits that the

observed BOLD signal in region r is a function of dynamically changing associations with all the other ROIs, which can in turn be modified by the observed stimulus,

$$Y_r(t) = \mathbf{Z}^r(t)^T \boldsymbol{\beta}^r(t) + \sum_{k=1}^K \{S_k(t) \alpha_k^r + \mathbf{X}_k(t)^T \boldsymbol{\gamma}_k^r\} + \epsilon_r(t), \quad (3.2)$$

where $\mathbf{Z}^r(t)$ denotes the $(R - 1)$ -dimensional vector of the BOLD signals recorded at time t in all the ROIs, excluding ROI r ; the scalar valued time-series $S_k(t) = \int_{-\infty}^t h(\tau) s_k(t - \tau) d\tau$ is the convolution of the stimulus k with the canonical hrf, $k = 1, \dots, K$; and each entry of the $(R - 1)$ -dimensional vector of time-series $\mathbf{X}_k(t)$ is obtained as $X_{kd}(t) = \int_{-\infty}^t h(\tau) s_k(t - \tau) y_d(t - \tau) d\tau$, i.e. as the product of the neural signal (deconvolved BOLD signal) of each ROI d , $d = 1, \dots, d \neq r$, with the stimulus k convolved with the canonical hrf (Di et al., 2021). The scalar $\alpha_k^r \in \mathbb{R}$ and the vector $\boldsymbol{\gamma}_k^r \in \mathbb{R}^{R-1}$, respectively, denote the psychological and psychophysiological interaction effects for stimulus k , $k = 1, \dots, K$. Lastly, the $(R - 1)$ -dimensional vector $\boldsymbol{\beta}^r(t)$ indicates the *time-varying* physiological effect capturing the functional connectivity with the seed region independently of any stimulus effect. Using the terminology employed in the Introduction, this effect can be regarded as capturing time-varying *background* functional connectivity. For simplicity, in the following, we will refer to (3.2) as the time-varying Psychophysiological interaction (tvPPI) model.

The proposed model effectively extends the Generalized PPI model of McLaren et al. (2012) by allowing the physiological effects, i.e the correlation between ROIs, to change over time. Indeed, if $\boldsymbol{\beta}^r(t) = \boldsymbol{\beta}$, i.e. if the elements of $\boldsymbol{\beta}^r(t)$ are constant over time, then the proposed tvPPI model reduces to the Generalized PPI model. Figure 3.1 provides an illustrative example of the type of dynamic relationship between two time series that the proposed model can help capturing. More specifically, we generate data from a centered bi-variate gaussian process with marginal variance 1, and we assume that the correlation between the realizations at each time t varies over time, according to the function $\rho(t) = 0.2 * \sin(t/10) + s_t \gamma$, where the task stimulus indicator s_t switches between 0 and 1 every 25 seconds and $\gamma = 0.6$ is the coefficient capturing the stimulus effect.

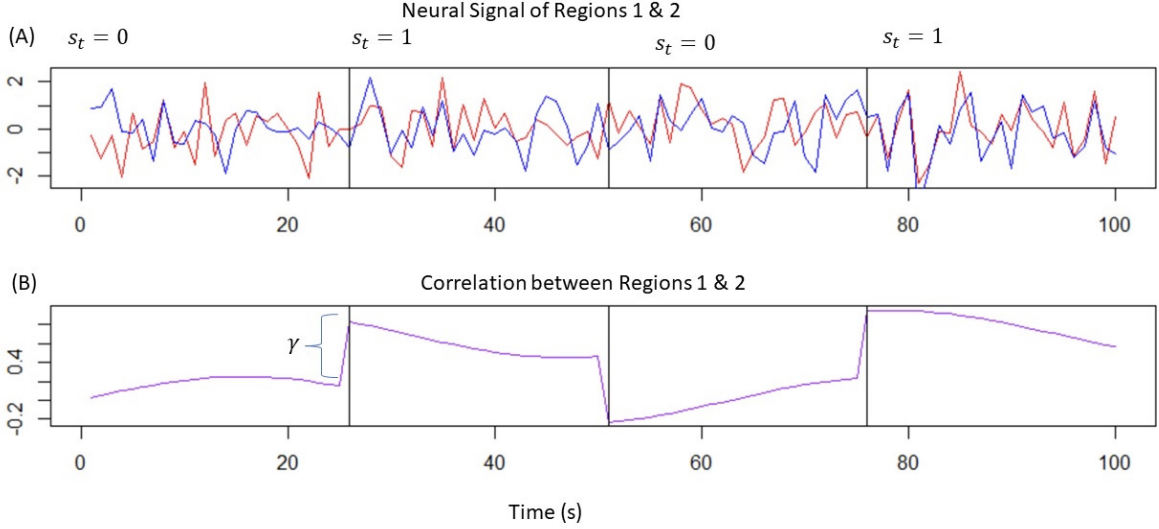


Figure 3.1: *Illustration of the proposed time-varying PPI paradigm*: similarly to the Generalized PPI model, the proposed model allows to capture stronger dependence between two realized time-series due to stimulus effects (top). In addition, it allows to capture smooth changes in the correlation over time, independent of the stimulus effect (bottom). For more details, see Section 3.1.

The two simulated time series in the displayed realization of Figure 3.1 (top) are only weakly dependent when s_t , the task indicator, is equal to 0. The dependence increases and the time series are more strongly coupled when $s_t = 1$. The PPI γ effect of the simulated stimulus s_t is shown by the sudden jump when the stimulus becomes active in Figure 3.1 (bottom). However, the correlation between the two time series changes over time and slowly weakens the longer the block remains active. The smooth changes of the correlation over time are modeled by β_t . The existing non-dynamic versions of PPI models are incapable of capturing the subtle fluctuations in the strength of connectivity between the two time-series.

For a given seed-region r , we can employ the partial correlation between region r and any other ROI $d = 1, \dots, R, d \neq r$, as a measure of functional connectivity. Based on model (3.2), the partial correlation is proportional to $\rho(r, d, t) \propto \beta_d^r(t) + \sum_{k=1}^K S_k(t) \gamma_{kd}^r$. By allowing the seed region r to vary, $r = 1, \dots, R$, we can fit R versions of model (3.2), and obtain partial correlations for any pair (r, d) , $r, d = 1, \dots, R$, at each time point $t = 1, \dots, T$.

Correspondingly, with the seed regions indicating the row number, at each time point we can form functional connectivity matrices, say $P_t \in \mathcal{R}^{R \times R}, t = 1, \dots, T$, where the diagonal elements are ones and each off-diagonal term is the partial correlation for the associated seed region-ROI pair. The resulting functional connectivity matrices are not necessarily symmetric. More importantly, a significant partial correlation $\rho(r, d, t)$ when ROI r is assumed as seed region does not guarantee that the partial correlation $\rho(d, r, t)$ is also significant, when ROI d is assumed as seed region. In Section 3.2.2, we describe a multi-comparison procedure to control the rate of false positive connectivities and regularize the inference across model fits.

3.1.1 Shrinkage Prior for the Dynamic Physiological Effects

The time-varying coefficients in the vector $\beta^r(\mathbf{t})$ capture the relevant psychological effects between the seed region r and any other ROI $d, d = 1, \dots, R, d \neq r$. It is reasonable to expect that only a subset of the ROIs will be associated with the seed region at any time $t = 1, \dots, T$. Hence, we employ variable selection priors to enforce regularization of the inference problem and identify relevant associations between the BOLD signals. Recently, there have been a few proposals for dynamic variable selection priors in sparse state-space models and time-varying models in the Bayesian literature (Rockova and McAlinn, 2020; Kowal et al., 2019). Here, we use the double gamma prior recently proposed by Bitto and Frühwirth-Schnatter (2019) and Cadonna et al. (2020). This prior generalizes the popular Horseshoe prior (Carvalho et al., 2010) to the time-varying modeling framework. In our practical experience, this choice of prior is computationally efficient, avoids overfitting by effectively shrinking the process variances of the dynamic coefficients to zero, and can flexibly adapt to varying degrees of sparsity in the data.

More in detail, the prior assumes that the psychological effect between a seed region r and

an ROI d , $\beta_d^r(t)$, $d = 1, \dots, R$, $d \neq r$, changes over time according to a discrete random walk,

$$\beta_d(t) = \beta_d(t-1) + w_d(t), \quad (3.3)$$

where we have dropped the superscript of the seed region, r , on the parameters for ease of notation. The step size $w_d(t)$ is a random variable following a scale-mixture normal-gamma-gamma shrinkage prior on the variance components,

$$\begin{aligned} w_d(t) &\sim N(0, \theta_d), \\ \theta_d | \zeta_d &\sim G\left(\frac{1}{2}, \frac{1}{2\zeta_d}\right), \\ \zeta_d | a_\zeta \kappa_B &\sim G\left(a_\zeta, \frac{a_\zeta \kappa_B}{2}\right). \end{aligned} \quad (3.4)$$

The model specification requires an initial condition, which assumes that at time $t = 0$, $\beta(0)_d \sim N(0, \theta_d)$. This construction emphasizes the role of θ_d in both the identification of significant connections and whether or not the connections vary over time. A very small value of θ_d encourages smooth transitions over time and shrinking of the coefficients to static (i.e., not time-varying) values. In particular, if $\beta(0)_d$ is close to 0, the coefficient should remain in close proximity of 0 over the entire time domain due to the small step size of the random walk process. The process prior (3.4) is an example of a *global-local* shrinkage prior: the parameter κ_B is a global shrinkage parameter as it controls the overall level of sparsity among all the θ_d 's, $d = 1, \dots, R$, $d \neq r$. On the other hand, the parameter ζ_d is a local shrinkage parameter, since it allows shrinking the corresponding θ_d to 0 as necessary, with the result that non-zero θ_d 's do not correspondingly decrease toward 0. The prior specification (3.3)–(3.4) depends on two hyperparameters: a_ζ and κ_B . We follow Cadonna et al. (2020) and use a default of $a_\zeta = 0.1$, while we let κ_B be user-defined. A simulation from the distribution of the double gamma prior for specific values of κ_B can be informative toward specifying an *a priori* level of global shrinkage. Figure 3.2 shows such a simulation. The conditional distribution of $\theta|\zeta$ is heteroskedastic with both increasing mean and variance

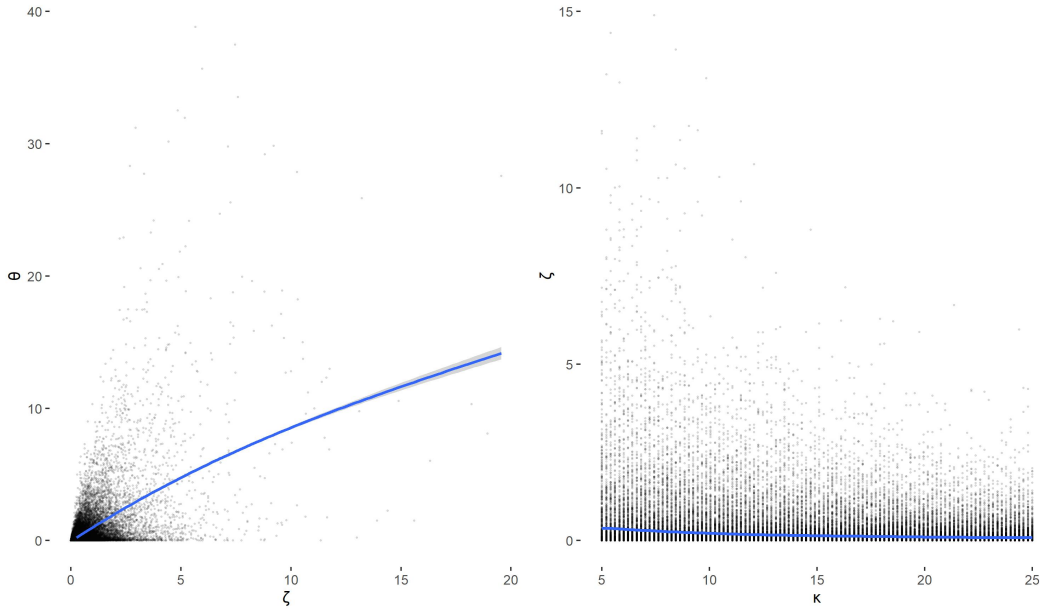


Figure 3.2: *Choice of the global shrinkage parameter κ_B* : Simulations from $P(\theta|\zeta)$ where the blue line is a smoother estimating $E(\theta|\zeta)$ (left) and from $P(\zeta|\kappa_B, a_\zeta = 0.1)$ (right) in (3.4). The conditional distribution of $\theta|\zeta$ is heteroskedastic with increasing mean and variance as a function of ζ . Larger and more varying values of ζ are associated with smaller values of κ_B .

as a function of ζ . Larger and more varying ζ are associated with smaller values of κ_B . Thus setting κ_B to be large (> 10) leads to more sparse solutions, characterized by less dynamic changes in the coefficients. Instead, setting κ_B to be small would correspond to prefer less sparse solutions *a priori*. Recently, Knaus et al. (2021) have suggested $\kappa_B = 20$ as a default setting, which has been shown to exhibit good simulation performance.

3.1.2 Shrinkage Prior for the PPI Effects.

Model (3.2) includes $K \times (R-1)$ static psychophysiological interaction vectors, $\gamma_k^r = \{\gamma_{dk}\}_{d \neq r}$, and K psychological effects, α_k . We expect that only a sparse subset of the parameters will identify significant associations between the fMRI BOLD signal in the seed region and either the psychological or the PPI factors. Thus, we adopt the horseshoe prior of Carvalho et al. (2010) to shrink the irrelevant PPI and psychological effects to zero. More specifically, for

notational simplicity, we introduce the vector $\boldsymbol{\eta}_k = [\gamma_{1k}, \dots, \gamma_{R-1,k}, \alpha_k]$ to indicate the R dimensional vector concatenating the PPI and psychological effects of each stimulus k , and we let η_{kh} denote the generic element of such vector, $k = 1, \dots, K$, $h = 1, \dots, R$. Then, we assume the *global-local* shrinkage prior,

$$\begin{aligned} \eta_{kh} | \tau^2, \lambda_{kh}^2 &\sim N(0, \tau^2 \lambda_{kh}^2) \quad h = 1, \dots, R \\ \tau^2 &\sim C_+(0, 1) \\ \lambda_{kh}^2 &\overset{ind}{\sim} C_+(0, 1) \quad h = 1, \dots, R. \end{aligned} \tag{3.5}$$

where the parameters τ and λ induce global and local shrinkage, respectively, of the coefficients toward zero. Here, the interest is in identifying the main effects and the interaction effects of psychophysiological factor on brain activity. These effects are assumed as static, i.e. not-time varying to ensure identifiability of all the model parameters. Indeed, if the stimuli are k are blocked indicators – as is often the case – it would be challenging to learn the dependence structure between time-varying interaction parameters, say γ_{kdt} and $\gamma_{kd(t+1)}$, when $s_{kt} = s_{k(t+1)} = 0$. In the lack of substantive prior information about how the modification of the connection between seed region r and region d by stimulus k were modified by successive exposures to stimulus k , the choice of a static specification of these parameters allows clearly interpretable inferences and reduced computational complexity.

3.2 Posterior Inference

In this Section, we outline how to obtain inference on the parameters of model (3.2). First, we discuss the computational algorithm for drawing posterior samples of our parameters of interest, namely α_k , $\beta_{dk}(t)$, and γ_{dk} . The use of continuous global-local shrinkage priors for the model parameters does not lead to a positive probability mass to *exact zero*'s in the posterior distribution of irrelevant parameters. Therefore, we propose a selection procedure inspired by the work of Chandra and Bhattacharya (2019) on testing the related hypotheses. Quite

notably, by casting the selection problem in a *non-marginal* multi-comparison framework, our approach is able to combine inferences across models that assume different seed-regions, and obtain unique connection graphs across models. Lastly, we discuss constructing point estimates for the time varying partial correlation matrices $P_t \in \mathcal{R}^{R \times R}, t = 1, \dots, T$.

3.2.1 Markov Chain Monte Carlo Algorithm

For every choice of the seed region, posterior inference is conducted for each seed-region in parallel through Markov Chain Monte Carlo (MCMC). In the following, we discuss the implementation for a single region, say $r \in \{1, \dots, R$, as in equation (3.2). The prior setup allows simulation from the joint posterior of the parameters by combining the collapsed sampler employed by Cadonna et al. (2020) for the time-varying regression coefficients $\beta_d(t)$ and the auxiliary variable augmented Gibbs sampler of Makalic and Schmidt (2016a) for the static regression coefficients η_k . The general procedure for sampling the parameters for each regional model is straightforward and computationally efficient. Once all parameters are initialized, one can simply iterate between the following two steps:

1. **Update** $\beta, \theta, \zeta, \sigma$: Let $Y_{t\beta} = Y_t - \sum_{k=1}^K s_{tk} \gamma'_k X_t, \forall t$. Then follow steps A-D of Algorithm 1 in Cadonna et al. (2020).
2. **Update** η, τ, λ : Let $Y_{t\gamma} = Y_t - \beta'_t X_t, \forall t$. Use the conditional distributions in Makalic and Schmidt (2016a) equations 7 and 8, by substituting η for β in their notation.

The first step is efficiently implemented in R using the *shrinkTVP* package by Knaus et al. (2021). Once again, it is worth noting that since the model specified in Equations 3.2–3.5 is independent across each seed region, we can fit each of the R models in parallel leading to a sizeable computational gain.

3.2.2 Selection of Non-Zero Components

In the Bayesian literature, variable selection is often implemented either by using continuous shrinkage priors or discrete selection models, e.g. through a spike-and-slab prior specification (George and McCulloch, 1993, 1997). The latter approach explicitly allows computing the posterior probability that a coefficient is exactly zero, but imposes additional computational complexity. Continuous shrinkage priors result in increased speed of implementation; however, they lack generally accepted straight-forward selection criteria and require careful posterior assessment of the inference. For example, Carvalho et al. (2010) as well as Cadonna et al. (2020) prescribe using a thresholding approach on the estimated values of a latent shrinkage factor to identify relevant coefficients.

Our problem presents two sets of challenges to the direct applicability of the conventional thresholding approaches described above. On the one hand, at each time point $t = 1, \dots, T$, lack of connectivity between the seed region r and any other region d should be assessed by testing the statement: $H_{0rdt} : \beta_{td}^r + \sum_{k=1}^K s_{tk} \gamma_{kd}^r = 0$, which encodes conditional independence between the two regions by setting the corresponding partial correlations to zero. Therefore, any formal post-MCMC selection procedure should take into account the combination of the inference on the static and dynamic coefficients of the model. Indeed, Cadonna et al. (2020) apply a thresholding rule on a local shrinkage factor which tests the relevance of the coefficient $\sqrt{\theta_d}$ in (3.4). However, this solution does not apply to our hypothesis of interest due to the dual role of θ_d informing not just variable inclusion but also in determining whether β_{td} is time-varying or static.

A second important consideration is that decisions based on a model studying the association between seed region d and all the other brain regions should take into account also the inference for models that assume any other brain region as seed region. In other words, for any pairs of regions (r, d) , inference assessing evidence for the null hypothesis $H_{0rdt} :$

$\beta_{td}^r + \sum_{k=1}^K s_{tk} \gamma_{kd}^r = 0$ when r is assumed as seed-region is inherently related to inference on the corresponding hypothesis H_{0drt} when d is assumed as seed region, in that it would be desirable for inference on H_{0rdt} to agree with inference on H_{0drt} , $t = 1, \dots, T$. This consistency would maintain the non-directed property of functional connectivity. In other words, once we consider the combined results of all models obtained by rotating the seed-regions, we are facing a multi-comparison problem with dependent hypotheses.

To identify relevant functional brain connectivity associations, we adapt the non-marginal decision theoretic approach recently developed by Chandra and Bhattacharya (2019). Here, the word non-marginal means that the decision theoretical framework does not rely solely upon marginal decision rules, e.g., marginal distributions of the test statistics, marginal p-values or marginal posterior probabilities, as it is often the case for standard multi-comparison procedures. On the contrary, the non-marginal decision theory framework proposed by Chandra and Bhattacharya (2019) allows to incorporate the joint structure inherent in the hypotheses and the dependence between parameters directly in the error measures, leading to a modified false discovery rate (mFDR) criterion.

More in detail, let $G_{rdt} = \{H_{0rdt}, H_{1rdt}, H_{0drt}, H_{1drt}\}$ denote the collection of null and alternative hypotheses statements for identifying the association between regions r and d , where $H_{1rdt} : \beta_{td}^r + \sum_{k=1}^K s_{tk} \gamma_{kd}^r \neq 0$ (similarly for H_{1drt}) and $d < r$, $d, r = 1, \dots, R$. Let $D_{rdt} = I(H_{1rdt} \text{ is accepted})$ indicate the decision in favor of conditional dependence (non-null functional connectivity association) between regions r and d at time t , when region r is assumed as the seed ROI. Correspondingly, let $z_{rdt} = I(H_{1drt} \text{ is accepted})$ indicate the decision made in support of conditional independence between d and r at time t when region d is assumed as the seed ROI. Lastly, let $h_{rdt} = I(H_{1rdt} \text{ is true})$ indicate the a true association (true state of nature). Following Chandra and Bhattacharya (2019), at each time $t = 1, \dots, T$, a non-marginal decision rule will aim at maximizing the posterior expectation

of true positive decisions (TP):

$$TP = \sum_{i=1}^{rdt} D_{rdt} h_{rdt} z_i$$

while controlling for possible decision errors. In the context of non-marginal decision rules, errors also derive from the fact that related hypotheses are wrongly accepted, wrongly rejected, or are not coherent with each other. Chandra and Bhattacharya (2019) show that these errors add up to the following cumulative sum,

$$E = \sum_{rdt} D_{rdt} (1 - h_{rdt} z_{rdt}). \quad (3.6)$$

The resulting loss function can be seen as a constrained maximization of the posterior expectation of true decisions penalized by both making false positive decisions as well as non-symmetric decisions. The objective function can be also rewritten as a function of the admissible decisions in the set of all possible decision configurations, under some penalization constant $\eta \in \mathbb{R}^+$:

$$f_\eta(\mathbf{D}) = \sum_{rdt} D_{rdt} (w_{rdt}(\mathbf{D}) - \eta), \quad (3.7)$$

where $w_{rdt}(\mathbf{D}) = P(H_{1rdt} \cap H_{1drt} | Y)$ denote the posterior joint probability of the alternative hypotheses related to regions r and d . Since all the models are fit separately across all seed regions, we compute $w_{rdt}(\mathbf{D})$ as the product of the two probabilities,

$$w_{rdt}(\mathbf{D}) = P\left(|\beta_{td}^r + \sum_{k=1}^K s_{tk} \gamma_{kd}^r| \geq c\sigma^r\right) P\left(|\beta_{tr}^d + \sum_{k=1}^K s_{tk} \gamma_{kd}^d| \geq c\sigma^d\right)$$

for some small $c \in \mathcal{R}$. In the following implementations, we have used $c = 0.1$.

Following Chandra and Bhattacharya (2019), the objective function (3.7) can be initially maximized by assuming $\beta = 1 - mFDR_x$, where $mFDR_x \in (0, 1)$ indicates the *posterior*

modified false discovery rate, defined as

$$mFDR_x = \sum_{\mathbf{D} \in \mathcal{D}} \frac{\sum_{r,d,t} D_{rdt} (1 - w_{rdt}(\mathbf{D}))}{\sum_{r,d,t} D_{rdt}} \delta(\mathbf{D} | \mathbf{Y}), \quad (3.8)$$

where \mathcal{D} denotes the space of all possible decision configurations, and the indicator variable $\delta(\mathbf{d} | \mathbf{X}_n)$ is equal to one when we consider the decision configuration chosen to be the final decision rule according to the associated multiple testing procedure.

For each given η , the maximization of $f_\eta(\mathbf{D})$ employs the simulated annealing method in Algorithm 2 of Chandra and Bhattacharya (2019). However, Chandra and Bhattacharya (2019) show how setting $\beta = mFDR_x$ may be overly conservative due to the additional penalty for incorrect decisions encoded in the posterior $mFDR_x$ criterion, particularly when the weights w_{rdt} are small. As a way around the problem, they suggest to run the simulated annealing algorithm multiple times, decreasing η each time until the realized $mFDR_x$ is at the desired level – a strategy that is computationally intensive for large datasets. Here, we follow an alternate strategy, motivated by the following observation. In our experiments, the achieved $mFDR_x$ is often below the nominal false discovery rate when the initial run of the simulated annealing algorithm leads to optimal decision pairs of decision that are discordant, that is of the type $D_{rdt} \neq D_{drt}$, $r, d = 1, \dots, D$. Subsequent runs of the annealing algorithms appear to relax these inconsistencies. Therefore, in order to gain computational speed, after the run of the simulated annealing algorithm, all the decisions are inspected we conclude that two regions are not associated only if the optimal solution identified by the simulated annealing requires both $D_{rdt} \& D_{drt} = 0$. This simple modification appears to perform well in our investigations.

3.2.3 Inference on Partial Correlations

Based on the associations estimated from the non-marginal decision rules and our modified algorithm, we can then proceed to estimate the degrees of association between pairs of regions. More specifically, we employ the posterior samples from the MCMC algorithm to estimate the physiological and PPI effects, and construct the time varying precision matrices: $\mathbf{\Omega}_{rt} \in \mathcal{R}^{R \times R}$, $t = 1, \dots, T$, with the subscript r indicating that region r is assumed as seed, $r = 1, \dots, R$

Since model (3.2)) assumes gaussianity, conditional independence can be assessed by considering the distribution of $Y_r(t)$ conditioned on all other regions $\mathbf{Z}^r(t)$,

$$Y_r(t) | \mathbf{Z}^r(t) \sim N \left(-\frac{\mathbf{\Omega}_{rz^rt}}{\mathbf{\Omega}'_{rrt}} \mathbf{Z}^r(t), \mathbf{\Omega}_{rrt}^{-1} \right), \quad (3.9)$$

where $\mathbf{\Omega}_{rz^rt}$ denotes the r^{th} row of $\mathbf{\Omega}_t$, excluding the entry on the diagonal, and $\mathbf{\Omega}_{rrt}$ is the r^{th} diagonal entry of $\mathbf{\Omega}_t$. Comparing with (3.2), it follows that $\mathbf{\Omega}_{rrt} = Var(\epsilon_t)^{-1}$, whereas the vector of partial correlations $\mathbf{\Omega}_{rz^rt}$ can be approximated by $\mathbf{\Omega}_{rz^rt} \approx \boldsymbol{\beta}^r(t) + \sum_{k=1}^K S_k(t) \boldsymbol{\gamma}_k^r$, after subtracting the psychological effects. Here, we assume that although the elements $X_{kd}(t) = \int_{-\infty}^t h(t) s_k(t-\tau) y_d(t-\tau) d\tau$ of $\mathbf{X}_k(t)$ do not coincide with $S_k(t) Z_d^r(t)$, $d = 1, \dots, R$, $d \neq r$, their values should be similar in a small neighborhood of each time point. Thus, using Equation (3.9) and the non-marginal decision theoretic approach outlined in Section 3.2.2, we can estimate the partial correlations from each fitted seed-based model, say $\tilde{\mathbf{\Omega}}_{rt}$, from the posterior samples of β_t , γ_k and σ_t^r . Then, we can combine the results across runs and obtain a final estimate of the time-varying precision matrix $\hat{\mathbf{\Omega}}_t$ by setting each off-diagonal entry as follows:

$$\hat{\mathbf{\Omega}}_{rdt} = \begin{cases} 0 & \text{if } D_{rdt} \ \& \ D_{drt} = 0 \\ \frac{\tilde{\mathbf{\Omega}}_{rdt} + \tilde{\mathbf{\Omega}}_{drt}}{2} & \text{otherwise} \end{cases} \quad r, d = 1, \dots, R \quad (3.10)$$

where – for any two regions r and d – the decisions D_{rdt} and D_{drt} are the optimal decisions

obtained by the simulated annealing algorithm of Chandra and Bhattacharya (2019) in the models fitted assuming seed region r and d , respectively. See Section 3.2.2.

3.3 Simulation

In this section we compare our proposed dPPI model with the more standard generalized PPI (gPPI) of McLaren et al. (2012). In this setting we show how the model perform estimating PPI effects under an assumption of resting state dynamic functional connectivity. We simulate from a gaussian distribution with a dynamic covariance matrix where PPI effects are known explicitly. In this way we can compare the bias and standard errors for the PPI effects of interest.

We show the appropriateness of our model in estimating PPI effects under time-varying functional connectivity data generating mechanism. We simulate a response time course for $R = 12$ over $T = 600$ in an event related design with $K = 2$ stimuli. We simulate the dynamic partial correlations through a mixture of two covariance processes.

The first is a stimulus linked process $\Omega_{s(t)}^1$ where $s(t) = 0, \dots, K$ is the stimulus sequence. Ω_0^1 is the precision matrix during rest and is block diagonal consisting of 4 3×3 matrices. Ω_1^1 , the precision matrix when $s(t) = 1$, is equal to Ω_0^1 for the first and last 3×3 matrix entries, while the middle 6×6 entries are changed to create a PPI effect with stimulus 1. Similar is done for Ω_2^1 , but now the middle 6×6 are equal to Ω_0^1 and the first and last 3×3 block matrices are changed. In this way, there are no PPI effects for the first 3 regions, while other regions have their connectivity modified by at least one stimulus.

The second covariance process Ω_t^2 is dynamic precision matrix equal to the identity matrix except for the (1,2)-(2,1) entries and the (14,15)-(15,14) entries. We set $\Omega_t^2(1, 2) = \Omega_t^2(2, 1) = 0.43\sin(\frac{0.01\pi t}{2} + 2.2)$ and $\Omega_t^2(10, 11) = \Omega_t^2(11, 10) = 0.47\sin(\frac{0.017\pi t}{2} - 1.3)$. This process

simulates the partial correlation between the above pairs fluctuating regardless of the active stimulus. With these two process defined we simulate $Y(t) \sim N_{12}(0, (\Omega_{S(t)}^1)^{-1} + (\Omega_t^2)^{-1})$. We use this process to create 60 simulation datasets to fit the dPPI and generalized PPI models. We then construct partial correlation matrices Ω_t^{dppi} and Ω_t^{gpqi} using the method outlined in Section 3.2.3.

We compare the bias and standard errors of three PPI effects of interest, which we note using the notation of Equation 3.2. γ_{21}^1 , which is the PPI between regions 1 and 2 under stimulus 1. Since $\Omega_0^1(1, 2) = \Omega_1^1(1, 2)$, $\gamma_{21}^1 = 0$ even though there is a dynamic correlation between the two regions. Secondly we look at γ_{51}^4 as the (4,5) pair has a constant resting state correlation and a non-zero PPI with stimulus 1. Lastly we look at $\gamma_{11,2}^{10}$ as the (10,11) pair exhibits both a dynamic resting state correlation and a non-zero PPI with stimulus 2. Lastly we compute the mean squared error defined as $MSE = \sum_{t=1}^T \sum_{r=2}^R \sum_{d=1}^r (\Omega_{rdt} - \hat{\Omega}_{rdt})^2$ for the true partial correlation matrix at time t Ω_t . The results are compiled in Table 3.1. The dPPI and gPPI rows are simply the mean and standard deviations across simulation datasets for the PPI effects, γ , of interest. The dPPI w/ selection row computes the PPI effects after selection by computing the difference in partial correlation when the simulated subject transitions from the rest condition to the stimulus condition of interest after the dPPI selection procedure has been applied. dPPI exhibits less bias in estimating partial correlation under a resting state dynamic functional connectivity model. This is especially true when considering the region pair (10,11) which has both a dynamic resting state partial correlation and a non-zero PPI with stimulus 2. dPPI also achieves lower MSE when estimating the partial correlation trajectories between all regions, since gPPI fails to capture changes in partial correlation that are independent of the stimuli considered. Selection is shown to have additional benefits by reducing both bias for the PPI effects and MSE for the partial correlation path among all regions.

The rows in Figure 3.4 are the paths of the estimated partial correlations for the (1,2), (4,5),

Method	$\gamma_{2,1}^1 = 0$	$\gamma_{5,1}^4 = -0.055$	$\gamma_{11,2}^{10} = 0.269$	MSE
dPPI	0.090 (0.097)	-0.062 (0.088)	0.210 (0.100)	188.546 (27.169)
dPPI w/ selection	0.073 (0.099)	-0.041 (0.125)	0.238 (0.122)	147.557 (26.033)
gPPI	0.170 (0.088)	-0.076 (0.106)	0.101 (0.106)	240.231 (29.471)

Table 3.1: The results of dPPI and gPPI under simulation setting 1. Each column entry is the mean and standard deviation of the estimates of the chosen effect modifications across all simulations. dPPI achieves less bias under a setting where physiological effects are dynamic. Overall

and (10,11) pairs from a single simulation dataset. dPPI is able to recover the underlying dynamic partial correlations while also accommodating the effects of stimulus 2 modifying the partial correlations in the (1,2) and (10,11) pairs. With the (4,5) row, we show dPPI is also able to perform similarly to gPPI in the absence of resting state dynamic partial correlations, showing the benefits of the shrinkage prior on $\beta_d(t)$. Figure 3.3 averages the partial correlation over each rest/stimulus block for a single simulation dataset, outputting a form of static functional connectivity matrix. We see that our proposed selection procedure is able to recover the changes in the structure of the partial correlation matrix under each stimulus block.

3.4 Case Study

A topic of continuing interest in cognitive neuroscience is the manner in which individuals learn predictive associative relationships that can be used to support decision-making and planning for rewards Daw and Shohamy (2008). It has been shown that the brain learns multiple sets of such relationships, simultaneously, and that these transition probabilities are reflected in the activity of distinct neural structures. Our interest lies in the modification of functional connectivity by learning in a predictive learning task. We consider a task where where a probabilistic sequence of 4 images were shown to 20 participants. Data from eight of the participants are analyzed here Bornstein and Daw (2012). Participants were

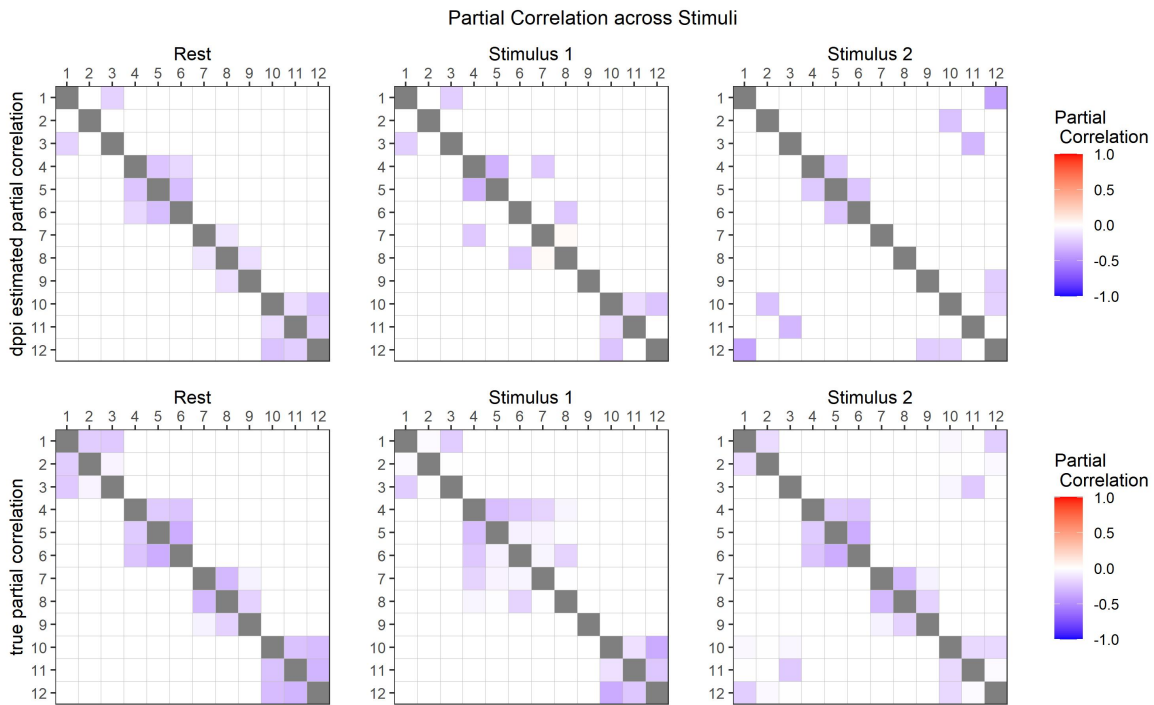


Figure 3.3: The average partial correlation matrices across each condition for a single simulation dataset. (Top row) The average estimated partial correlation matrices across the rest, stimulus 1, and stimulus 2 conditions respectively. (Bottom row) The average data-generating partial correlation matrices across the rest, stimulus 1, and stimulus 2 conditions respectively. We see there are some false positive connections among the 4,5,6 block where

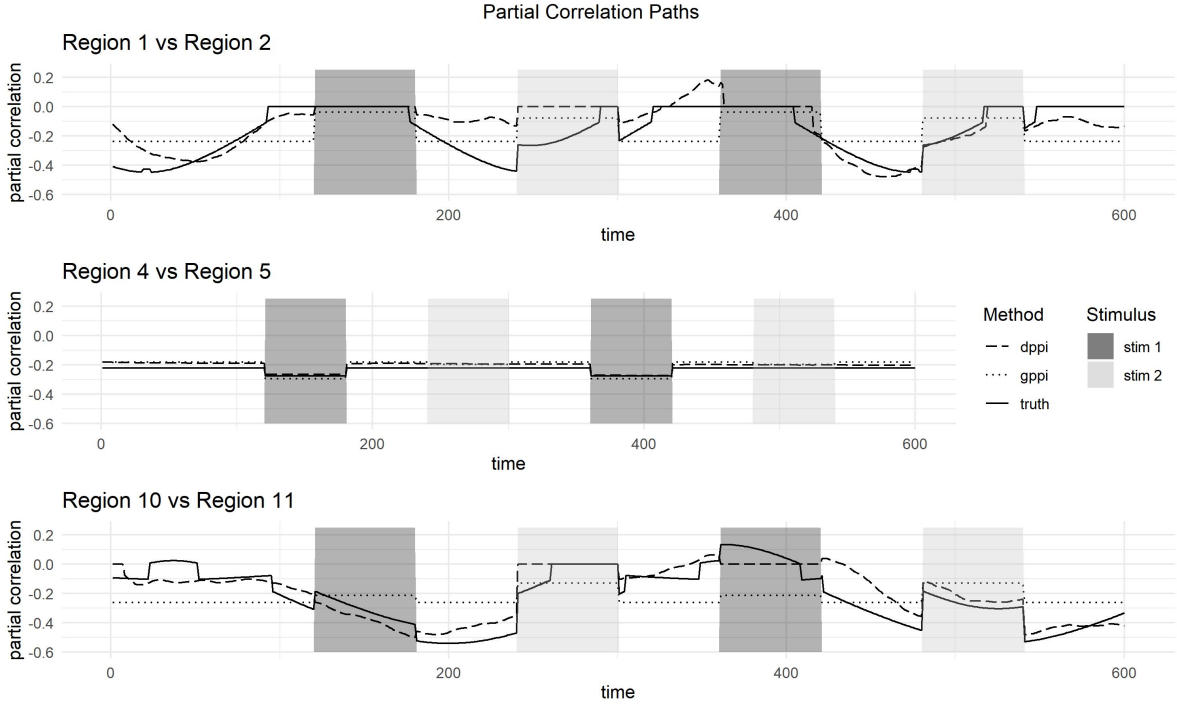


Figure 3.4: The paths of the estimated partial correlations for the (1,2), (4,5), and (10,11) pairs from a single simulation dataset. The solid black line shows the true partial correlations for the respective region pairs over time. The dashed and dotted lines are the estimates from dppi and gppi respectively, constructed using the method outline in Section 3.2.3. dppi is able to recover the underlying dynamic correlations while also accommodating the effects of stimulus 2 modifying the partial correlations in the (1,2) and (10,11) pairs. With the (4,5) row, we show dppi is also able to perform similarly to gppi in the absence of resting state dynamic partial correlations, showing the benefits of the shrinkage prior on $\theta_d(t)$.

tasked with responding to the shown image with a matched pre-trained button press. The sequence of images were randomly generated (unknown to the participants) according to a first order Markov process, where the probability a picture is shown in trial $e = 1, \dots, 999$ depends solely on the picture shown in trial $e - 1$. The full conditional distribution can be fully represented by a 4×4 matrix, where the (i, j) specifying the probability the picture represented by the j^{th} column entry is shown on trial e given that the picture represented by the i^{th} row is shown on trial $e - 1$. To encourage continual learning, the Markov transition kernels were changed twice at trials $e = 334$ and $e = 667$. Participant reaction times (RTs) were measured as a proximity for predictive learning, where faster RTs were indicative of learning the conditional distribution of the image sequence. The transition matrices were generated anew for each subject, and selected to minimize mixing time so that first-order dependencies were the sole consistent source of information about the identity of the next stimulus and thus could be used as the primary predictor of behavior.

Consistent with the experiment design and previous findings using a similar task, response times matched the first-order conditional distribution, suggesting that participants implicitly learned the probabilities incrementally, via trial-by-trial experience. Further, the relationship between RTs and the underlying stimulus probabilities evolved over learning, with recent experiences providing exponentially more influence on subsequent predictions. This phenomenon is measured in delta rule learning by the learning rate: the weight a system places on new information relative to previous experience. Following previous work on the division of learning into multiple systems Poldrack et al. (2001); Doll et al. (2015), we fit two Rescorla-Wagner learning rules under both a slow and fast learning rate: $P^s, P^f \in \mathcal{R}^{4 \times 4 \times 999}$. These represent the participants' learned conditional distribution of the images according to each learning rate. Let $I(e) \in \{1, 2, 3, 4\}$ be the label of the image shown at trial e . We

compute the learning rules as in Equation 3.11.

$$P_{i,j,e}^z = \begin{cases} P_{i,j,e-1}^z + \alpha_z(1 - P_{i,j,e-1}^z) & j = I(e) \\ P_{i,j,e-1}^z + \alpha_z(0 - P_{i,j,e-1}^z) & j \neq I(e) \end{cases} \quad z = s, f; \quad e = 1, \dots, 999, \quad (3.11)$$

with $\alpha_s = 0.0138$ and $\alpha_f = 0.5499$ as median parameter values across the population estimated in the previous analysis (Bornstein and Daw, 2012). We then use these learning rules to compute the forward entropy of the system under the different learning rates: $H_s(e)$ and $H_f(e)$ for the slow and fast rates respectively. The forward entropy is computed in Equation 3.12 and represents a measure of expected surprise for the subsequent image given the participant’s experience under the respective learning rate and the currently shown image.

$$H_z(e + 1) = - \sum_{j=1}^4 \log(P_{I(e),j,e}^z) P_{I(e),j,e}^z \quad z = s, f \quad (3.12)$$

This quantity is designed to index the amount of “lookahead” activity to be expected in anticipation of the upcoming stimulus Johnson and Redish (2007); Bornstein and Daw (2013). Lookahead activity has been identified in multiple regions associated with learning and memory Johnson and Redish (2007); Johnson et al. (2007), thus implicating these regions in *prediction* as well. Given that each region learns a distinct representation of sequential structure, and that behavior reflects a mixture of the predictions generated by each of these representations, a critical open question in the field is where and how these predictions are integrated Daw and Shohamy (2008). A key candidate structure is the Anterior Cingulate Cortex (ACC), a region that has the necessary anatomical connectivity to combine inputs from the hippocampal formation with those from the striatum, and which has itself been strongly associated with conflict between, and control of, competing motivational drives Shenhav et al. (2013). Indeed, ACC is a key overlapping region correlated with forward entropy from both processes in the original analysis Bornstein and Daw (2012). Intuitively, if the brain represents sets of two or more competing predictions that can drive behavior,

these predictions should be weighted according to their relative precision Wang et al. (2022) – here, the inverse of their entropy. Therefore, by this model, as the two transition matrix representations change with learning, the more predictive representation should be more tightly coupled with ACC. Here, we test a question that was unapproachable using previous methods: Namely, does the relationship between ACC and each ‘source’ region (here, hippocampus and caudate) change as a function of the difficulty of lookahead (entropy) in each process? In other words, the goal of this analysis is to estimate the PPI effect of the two forward entropy timeseries across the brain, with particular interest in the relationship between the two learning structures and ACC.

Imaging was performed on the 3T Siemens Allegra head only scanner with time resolution of 2.0 seconds per acquisition, across four sessions of 300 acquisitions each. Images were normalized into a template and resampled into $2 \times 2 \times 2$ -mm voxels in the normalized template space (MNI). These voxels were then further combined into $R = 18$ ROI by taking the mean BOLD of all voxel time series attributed to each ROI. ROI were defined anatomically according to a standard atlas (AAL). The regions selected were structures previously associated with learning, memory, and decision-making, along with several “control” regions, and separated bilaterally.

There are $K = 4$ covariates of interest included the model:

- An indicator for the participant seeing an image
- An indicator for the participant pressing a button
- The forward entropy under the slow learning rate: H_s
- The forward entropy under the fast learning rate: H_f

Each covariate was convolved with the canonical hrf (the default in SPM) to create the psychological regressors. To lessen the effect of false positive PPI effects, PPI regressors

were formed by first deconvolving the ROI time series into latent neural signals using the method of Bush and Cisler (2013). We then take the product of the deconvolved time series with the above covariates, before scaling and convolving once again with the canonical hrf (Di et al., 2017).

Consistent with the above-described neural circuit model, we observed a consistent pattern of entropy-modulated time-varying PPI between each source region and ACC. Specifically, we found that functional connectivity between ACC and Caudate increases when the opposing (slow, hippocampally-linked) process becomes more uncertain (Fig. 3.5). This suggests that when hippocampus has greater difficulty predicting the next stimulus in the sequence, Caudate is more tightly coupled with activity in the putative output structure, ACC. Notably, the observation that the predominance of Caudate depends on the predictiveness of the hippocampal representation, independent of the predictiveness of the Caudate representation itself, is consistent with an inductive bias towards the former structure, perhaps due to its faster neural dynamics and reflecting its more recency-weighted predictions Wang et al. (2022).

Further, we observed time-varying PPI that supported the role of each region in performing “lookahead” to support behavior. Specifically, as prediction becomes more difficult in the fast LR process, we find greater coupling between Caudate and Putamen (Fig. 3.6), a structure widely observed to represent the action sequences upon which the Caudate transition matrix may draw during lookahead Smith and Graybiel (2013). Paralleling the fast entropy-mediated Caudate-Putamen relationship, greater entropy in the slow (hippocampal) process was associated with stronger connectivity between hippocampus and parahippocampal cortex (Fig. 3.7). The latter region is widely referred to as the “parahippocampal place area” (PPA; Epstein and Kanwisher 1998) due to its sensitivity to the presentation and prospective imagination of complex “scene” images. Given that the stimuli in this experiment were images of natural scenes, this finding is consistent with the idea that, during lookahead,

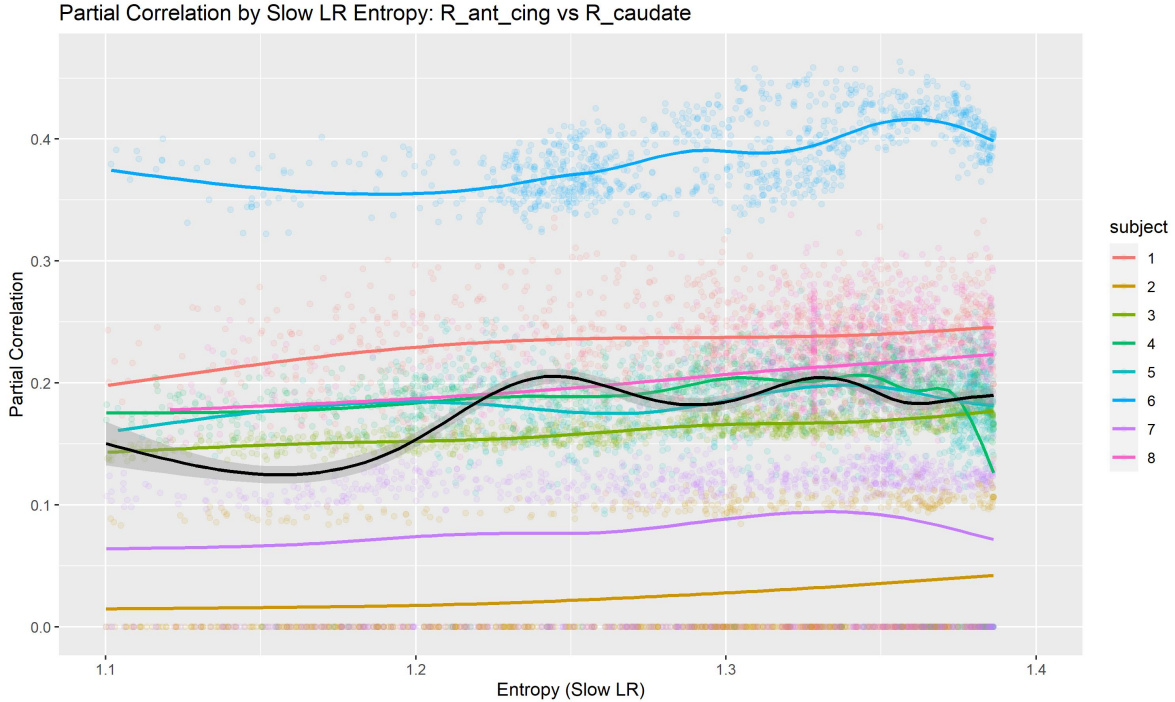


Figure 3.5: The estimated partial correlation between right side Caudate and ACC by the entropy of the slow learning rate process.

hippocampus is orchestrating reinstatements of scene images – as entropy increases, so does the number of potential next-step scene images that may be reinstated (see Bornstein and Daw (2013); Bornstein and Norman (2017) for corresponding findings in sequential decision-making tasks).

3.5 Discussion

We introduced model able to estimate PPI effects in the presence of background possibly non-linear dynamic functional connectivity. In simulation, dPPI exhibited less bias in estimating PPI effects when there is dynamic functional connectivity not predicted by stimuli, and performs similar to gPPI in estimating PPI effects of non-dynamically correlating regional pairs. We introduced a method to combine inference across regions through combining the results of multiple related linear models into a precision matrix and performing selection

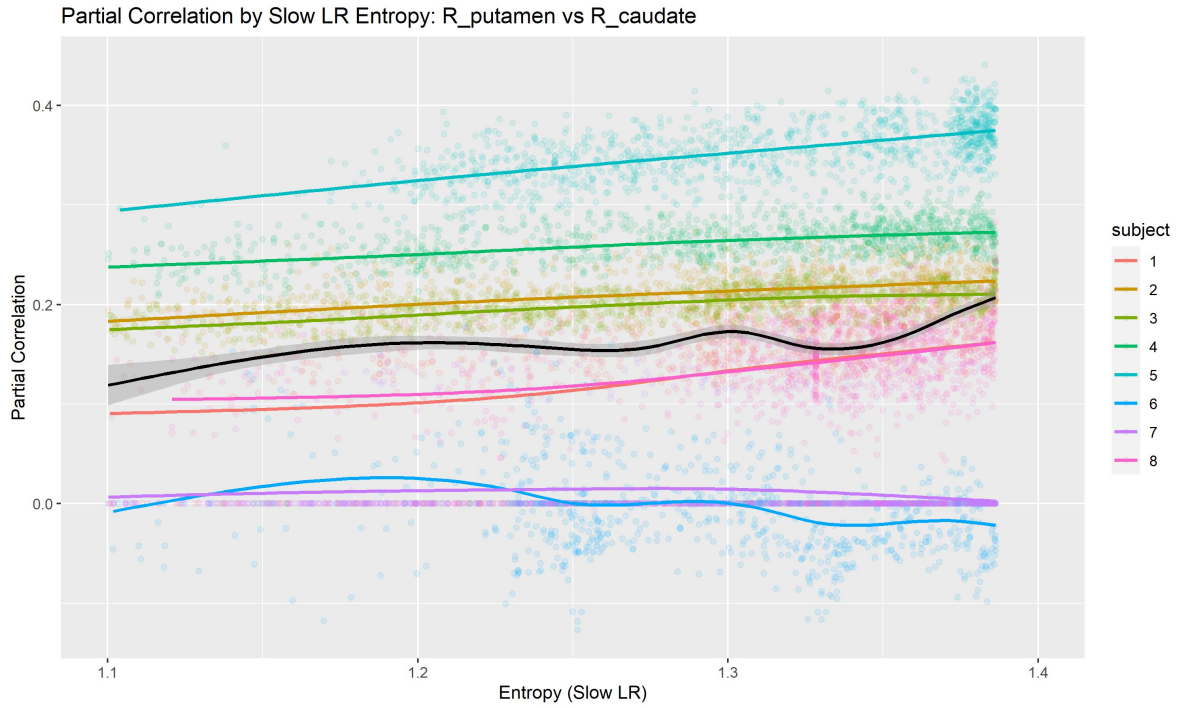


Figure 3.6: The estimated partial correlation between right side Caudate and Putamen by the entropy of the slow learning rate process.




Figure 3.7: The estimated partial correlation between right side Parahippocampus and the Hippocampus by the entropy of the slow learning rate process.

on the precision matrix entries. There is still work to be done in overcoming the assumptions inherent in combining posterior distributions from fitting the dynamic linear model in 3.1 across all regions. Firstly, there is no sharing of information between the models when proceeding through the MCMC scheme to sample from the posterior distributions of β_{td} , α_k , and γ_{dk} . While this allows the sampling for each region's parameters in parallel, a computational benefit, we then must pay by assume $w_{rdt} = P(H_{1rdt} \cap H_{1drt}|Y) = P(H_{1rdt}|Y)P(H_{1drt}|Y)$, implying independence in the posterior distributions between related hypotheses. While this assumption does not hurt in simulation, there is clear further work to be done to allow computation of w_{rdt} without this restriction.

Chapter 4

Dynamic Bi-Clustering of fMRI Data

Data are increasingly large and complex, requiring sophisticated methods to reveal patterns while also respecting heterogeneity across samples as more data are collected. Bi-clustering is a well established unsupervised technique for grouping observations across two axes of information, which can be collected within rows and columns within a matrix. In gene-expression data, observations are collected on thousands of genes in possibly hundreds of conditions. Bi-clustering can be used to group together conditions which exhibit similar gene clusters while simultaneously grouping together genes with similar expression profiles (Madeira and Oliveira, 2004). In text mining, documents are collected each with an abundance of different words. Bi-clustering is used to group together documents based on their patterns of word usage (Dhillon, 2001). Further examples of this are in brain imaging, where subjects can be grouped-together based on the clustering behavior of their respective brain regions (Ambrosen et al., 2014)

However, the patterns revealed by bi-clustering may not be stable over time, changing as successive samples are collected. This situation calls for a more flexible technique that allows for clusters to change over time. There are some approaches within this realm that simply

add time as a third mode to cluster with: so called tri-clustering approaches. Guigourès et al. (2018) analyzed short term bike rental data in London, where information was collected on bike pick-up locations and bike drop-off locations. They used a graph-based tri-clustering approach to simultaneously cluster pick-up and drop-off locations as well as the pick-up time stamps over a 9 month period. They found the bi-clustering between pick-up and drop locations changed heavily depending on the type of day. Wu et al. (2018) used a deterministic tri-clustering algorithm to cluster temperature data across three modes: station location in the Netherlands, days of the year, and across years itself. Through this they can describe reoccurring spatial temperature patterns over the course of a year while grouping similar years together.

While the above approaches indeed describe bi-clustering behavior changing over time, there is little temporal dependence assumed across successive cluster assignments. This is important when collecting data across time scales where one would apriori expect clustering behavior to be similar for time points that are closer together. Within fMRI brain imaging data, Amar et al. (2015) used a time series bi-clustering model to cluster brain regions along with subjects based on whether the region was activated or not. They specify a hierarchical model where regions are partitioned into 'modules' where every region within the modules is activated at the same time. They then pool together information across subjects by allowing subjects to have shared modules, deemed core modules, and then allowing further subject specific modules. Time dependence is incorporated through module activations but not over assignments of clusters. A drawback of this is failing to allow the brain regions to dynamically reconfigure their spatial dependencies, an important feature in brain imaging time-series (Chang and Glover, 2011; Calhoun et al., 2014).

We propose the dynamic bi-clustering (dbc) model for separate exchangeable matrix valued data collected over time. Rows (signal sources) are clustered within column (unit) clusters such that, all columns within a column cluster share the exact same partitioning of the rows

across all time points. Each row is then clustered together based on their similarity to other rows within the same column cluster. We incorporate time dependence using the approach of Page et al. (2021) where a row specific latent variable is introduced determining whether row cluster assignments of successive timepoints are the same. If this latent variable determines the assignment is to be different across successive timepoints, the model will assign the row a new cluster assignment. As not all rows are reassigned simultaneously, this allows cluster configurations between successive timepoints to be relatively similar. We also do not require the number of clusters to be specified apriori. Through the sparse specification of the dirichlet process, we allow for a user-specified upper limit on the number of clusters. The model will then automatically determine the number of unit clusters and signal source clusters visited over time (Frühwirth-Schnatter and Malsiner-Walli, 2019).

Additionally, we discuss an extension of the dbc model where units assigned to the same unit cluster are not forced to share the exact same sequence of signal source clusterings, but share distributions over potential signal sources clusters. Further, we allow for signal source cluster assignment changepoints to be informed by simultaneously collected covariates. This is motivated by task-based fMRI experiments where subjects are introduced to a stimulus, say looking at a picture, with possibly differing timings. In this setting, there is a reasonable expectation of cluster assignments changing as brain regions respond to the stimulus in differing ways.

In Section 4.3 we show the results of a simulation study, indicating the importance of accounting for the unit-level heterogeneity in signal source clustering. Finally, in Section 4.4, we apply the dbc model to an behavioral task fMRI experiment conducted on 24 subjects. We demonstrate the ability of the model to detect expected differences in FC configurations when subjects undergo differing conditions in a hand-grip task.

4.1 DBC Model

4.1.1 Motivation

Let us begin by considering the type of data and the objective that this work is motivated by. Suppose that we observe realizations of some random signal

$$Y_{rt}^{(i)}, \quad \text{for } \begin{cases} i = 1, \dots, N \text{ units} \\ r = 1, \dots, R \text{ signal sources} \\ t = 1, \dots, T \text{ time points} \end{cases}$$

Such data may be organized in a length- T time series of $N \times R$ dimensional matrices. In practice, $Y_{rt}^{(i)}$ could represent the blood oxygenation level dependant (BOLD) signal - that is a measure of neuron activation - for subject i in brain region r at time t as illustrated in Figure 4.1a. In microbiome analysis, $Y_{rt}^{(i)}$ may represent the abundance of microbe taxa r for subject i at time t , in music analysis $Y_{rt}^{(i)}$ may be the pitch of the sound played in song i by musical instrument r at time t , and so on.

From the data, we would like to learn what signal sources produce similar signals at given time points and what units tend to behave similarly over time. Specifically, our objectives are (i) clustering *signal sources* based on observed signals, allowing for such clusters to change over time; and (ii) simultaneously clustering *units* whose signal sources exhibit a homogeneous clustering behavior across time (see Figure 4.1b for an illustration). More formally, let $\tilde{c}_{rt}^{(i)} \in \{1, \dots, K\}$ be the cluster to which signal source r belongs at time t for unit i and let $s_i \in \{1, \dots, Z\}$ be the *static* cluster to which unit i belongs. We are interested in the case in which clusters of signal sources differ *across* unit clusters but are homogeneous *within* unit clusters, that is $\tilde{c}_{rt}^{(i)} = c_{rt}^{(s_i)}$, where $c_{rt}^{(s_i)}$ is the cluster to which signal source r belongs at time t for *all* units in unit-cluster s_i .




Figure 4.1: Toy example of data and objectives considered by our model, with 4 units, 4 signal sources and 3 time points (i.e. $N = 4$, $R = 4$, $T = 3$). (a) BOLD signals (represented by blue bars) recorded in multiple units (in this case subjects, represented by rows) at multiple sources (in this case brain regions, represented by the location of a bar) and time points (represented by columns). (b) The objectives of the model are (i) to identify time-specific clusters of signal sources that generate similar signals in a given time point (here 3 clusters, represented by the color orange, green or magenta of a bar, i.e. $c_{rt}^{(s_i)} \in \{1, 2, 3\}$); and (ii) to learn time-invariant clusters of units that exhibit the same clustering of signal sources across time points (here 2 clusters, represented by the color red or blue of a row label, i.e. $s_i \in \{1, 2\}$).

4.1.2 Dynamic bi-clustering

We specify a normal likelihood where the mean of the distribution depends on the unit and signal source clusters and where the variance is specific to the signal source:

$$Y_{rt}^{(i)} | s_i = z, c_{rt}^{(z)} = k \sim N(\mu_k, \sigma_r^2) \quad (4.1)$$

$$\mu_k \sim N(\mu_0, \sigma_0^2) \quad \sigma_r^2 \sim IG(\alpha_\sigma, \beta_\sigma), \quad (4.2)$$

for $\mu_0 \in \mathcal{R}$ and $\sigma_0^2, \alpha_\sigma, \beta_\sigma > 0$. Equation 4.1 illuminates the meaning of unit and signal source clusters. If two units i and i' are clustered together, i.e. $s_i = s_{i'}$, then the corresponding signal sources will share the exact same clustering behavior over the entire time series, as in Figure 4.1b. That is, for any fixed $r \in \{1, \dots, R\}$, $s_i = s_{i'}$ implies $Y_{rt}^{(i)}$ has the same distribution as

$Y_{rt}^{(i')}$ for every $t = 1, \dots, T$.

Let $\pi_z \in [0, 1]$ be the probability that a unit belongs to cluster $z = 1, \dots, Z$ and let $\omega_{zk} \in [0, 1]$ be the probability that at time 1 a signal source of a unit in unit-cluster z belongs to signal-source cluster $k = 1, \dots, K$. Then, let $\pi = (\pi_1, \dots, \pi_Z)$ with $\sum_{z=1}^Z \pi_z = 1$ and similarly for every $z = 1, \dots, Z$ let $\omega_z = (\omega_{z1}, \dots, \omega_{zK})$ with $\sum_{k=1}^K \omega_{zk} = 1$. We seek to place prior distributions on π and on each of $\omega_1, \dots, \omega_Z$ that both allow us to learn the number of necessary (i.e. nonempty) clusters and are amenable to tractable posterior inference. One such choice is the finite stick breaking prior of Frühwirth-Schnatter and Wagner (2010). The probability vector π follows a finite stick breaking prior of dimension Z with concentration parameter $\alpha > 0$, in formulas $\pi \sim \text{SB}_Z(\alpha)$, if it can be constructed in the following manner:

$$\nu_1, \dots, \nu_{Z-1} \stackrel{\text{iid}}{\sim} \text{Beta}(1, \alpha), \quad \nu_Z = 1 \quad (4.3)$$

$$\pi_1 = \nu_1, \quad \pi_h = \nu_h \prod_{j=1}^{h-1} (1 - \nu_j), \quad h = 1, \dots, Z - 1. \quad (4.4)$$

Similarly, we assume $\omega_1, \dots, \omega_Z \stackrel{\text{iid}}{\sim} \text{SB}_Z(\beta)$ with $\beta > 0$. By constraining the concentration parameters α and β to a small value, e.g. 0.01, the model from Equations 4.3-4.4 favors sparse probability vectors (Malsiner-Walli et al., 2016), i.e. those where some entries are approximately zero, representing empty clusters. Thus, if we set α and β small and the dimensions Z and K to be an upper bound to the number of unit and signal-source cluster respectively, this choice of priors allows learning the number of inhabited unit clusters $Z^+ < Z$ and signal source clusters $K^+ < K$. In practice, we can let Z be equal to the total number of units N and K equal the number of signal sources R times the number of time points T . This modeling choice on unit and signal source clusters can be summarized as follows:

$$\pi \sim \text{SB}_Z(\alpha), \quad (4.5)$$

$$s_i \mid \pi \stackrel{\text{iid}}{\sim} \text{Categorical}(\pi), \quad i = 1, \dots, N \quad (4.6)$$

$$\omega_z \stackrel{\text{iid}}{\sim} \text{SB}_K(\beta), \quad z = 1, \dots, Z \quad (4.7)$$

$$c_{r1}^{(s_i)} \mid \omega_{s_i} \stackrel{\text{iid}}{\sim} \text{Categorical}(\omega_{s_i}) \quad r = 1, \dots, R, \quad i = 1, \dots, N \quad (4.8)$$

There are additional considerations concerning setting the concentration parameters α and β . Frühwirth-Schnatter and Malsiner-Walli (2019) discuss placing on the concentration parameters a Gamma hyperprior with expected value equal to 0.5 to allow for the appropriate level of sparsity in the number of inhabited clusters. In our experience, however, setting $\alpha = \beta = 0.01$ worked well in simulation studies for recovering the true number of both unit and signal source clusters.

Equation 4.8 describes the distribution of only the initial signal-source cluster assignment. To describe the evolution of signal-source cluster assignments as time progresses, we would like a model complying with the a-priori expectation that sources that are clustered together (co-clustered) at time t are more likely to also be co-clustered at time t' when t is closer to t' and less so when t and t' are further apart. For units in cluster s_i , let the set of indices of signal sources in cluster k at time t be $S_{kt}^{(s_i)} = \{r \in \{1, \dots, R\} : c_{rt}^{(s_i)} = k\}$, so that the clusters $(c_{1t}^{(s_i)}, \dots, c_{Rt}^{(s_i)})$ define a partition $\rho_t^{(s_i)} = \{S_{1t}^{(s_i)}, \dots, S_{Rt}^{(s_i)}\}$ of signal sources of units in cluster s_i at time t . The a-priori expected similarity between the partition $\rho_t^{(s_i)}$ of signal sources of units in cluster s_i at time t and the partition $\rho_{t+\delta}^{(s_i)}$ of the same signal sources at time $t + \delta$, $\delta > 0$, should be decreasing in δ . Page et al. (2021) argued that for partitions to have this property, one should not rely on partitions *induced* by Bayesian nonparametric models but rather model the partition evolution directly, as we detail in Section 4.1.2.

Temporal random partitions

We describe the evolution of $c_{rt}^{(s_i)}$ through the temporal random partition model proposed by Page et al. (2021). This is done by introducing an indicator $\gamma_{rt}^{(s_i)}$ for the signal sources of units in cluster s_i being reassigned from time $t - 1$ to t . When $\gamma_{rt}^{(s_i)} = 1$, $c_{rt}^{(s_i)}$ equals the previous cluster assignment $c_{r(t-1)}^{(s_i)}$. Otherwise, when $\gamma_{rt}^{(s_i)} = 0$, a cluster c^* is re-sampled for signal source r among units in cluster s_i using the same distribution ω_{s_i} used to draw the initial cluster assignment (Equations 4.7 and 4.8). The following equations outline the evolution of cluster assignments of signal sources over time:

$$c_{rt}^{(s_i)} \stackrel{d}{=} \gamma_{rt}^{(s_i)} c_{r(t-1)}^{(s_i)} + (1 - \gamma_{rt}^{(s_i)}) c^* \quad (4.9)$$

$$c^* | \omega_{s_i} \sim \text{Categorical}(\omega_{s_i}) \quad (4.10)$$

$$\gamma_{rt}^{(s_i)} \sim \text{Bernoulli}(a_t^{(s_i)}) \quad (4.11)$$

$$a_t^{(s_i)} \stackrel{\text{iid}}{\sim} \text{Beta}(\alpha_a, \beta_a), \quad (4.12)$$

where $\alpha_a, \beta_a > 0$. The time-varying change-point parameter $a_t^{(s_i)}$ controls the dependence among successive partitions within the same unit cluster: $\rho_t^{(s_i)}$ and $\rho_{t+1}^{(s_i)}$. As $a_t^{(s_i)}$ approaches 1, cluster assignments are more likely to remain unchanged between successive timepoints, increasing the dependence between $\rho_t^{(s_i)}$ and $\rho_{t+1}^{(s_i)}$. It is important to note that $\gamma_{rt}^{(s_i)} = 0$ does not imply $c_{r(t-1)}^{(s_i)} \neq c_{r(t)}^{(s_i)}$. The new cluster assignment drawn, c^* , may in fact match the previous cluster assignment: $c_{r(t-1)}^{(s_i)}$. Figure 4.2 shows simulated data with $T = 100$ time points and 2 change points common to all of $R = 20$ signal sources and all of $Z = 3$ clusters of $N = 10$ units. Specifically, for all $i = 1, \dots, 10$ and all $r = 1, \dots, 20$, $\gamma_{r,33}^{(s_i)} = \gamma_{r,66}^{(s_i)} = 0$ while, for $t = \{2, \dots, 100\} \setminus \{33, 66\}$, $\gamma_{r,t}^{(s_i)} = 1$. Figure 4.2 illustrates that observations from signal sources that are in the same cluster at consecutive time points before the change points $t = 33$ and $t = 66$ are assumed to be generated from the same gaussian distributions.

The model described in Equations 4.1-4.12 accounts for the three dimensions of the data (multiple units, signal sources and time points) and achieves our objective of simultaneously clustering units and signal sources. Figure 4.3 summarizes the model visually by means of a directed graphical model and illustrates the (conditional) dependencies between its variables.

Covariate-Informed Changepoints

A property of the dynamic bi-clustering model is that units that are clustered together share the exact same clustering pattern among all signal sources at every timepoint. This is plausible in the case where all units are likely to experience the same phenomenon over the measured timecourse. An example could be in a block-designed experiments in fMRI, where subjects are presented with a stimulus, say a photo in a visual setting, in a synchronous fashion. In this way, there could be a reasonable expectation of similar timing of activations, and resulting clustering behavior, of brain regions (signal sources) within unit (subject) clusters. However, there may be circumstances where there is no synchronization of events across units. Subjects may be presented photos at different timepoints across each subject's respective scanning sessions. Further still, there may differing number of fMRI scans per subject, say T_i . In these task-based scenarios, we can leverage the timing of these stimuli as covariates to infer the changes in the clustering behavior. Consider the indicator of a changepoint at timepoint $t_i : \gamma_{rt_i}^i$. We model the log odds of this indicator via a logistic regression model using a set of p covariates for subject i : $X_t^i = [X_{t1}^i, \dots, X_{tp}^i]$. The covariate-informed changepoint dynamic bi-clustering (cic-dbc) variant of the model is outlined in

4.13.

$$\begin{aligned}
Y_{rt}^{(i)} | c_{rt}^i = k &\sim N(\mu_k, \sigma_r^2) & \gamma_{rt}^i &\sim \text{Bern}(a_t^i) \\
\mu_k &\sim N(\mu_0, \sigma_0^2) \quad \sigma_r^2 &\sim IG(\alpha_\sigma, \beta_\sigma) & \log\left(\frac{a_t^i}{1-a_t^i}\right) = X_t^i \beta_{s_i} \\
P(c_{r1}^i = k | s_i = z) &= \omega_{zk} & \beta_z &\sim N(\mu_\beta, \Sigma_\beta) \\
c_{rt}^{(s_i)} &= \gamma_{rt}^{(s_i)} c_{r(t-1)}^{(s_i)} + (1 - \gamma_{rt}^{(s_i)}) c^* & \omega_z &\sim SB(\beta_\omega) \\
c^* | \omega_{s_i} &\sim \text{Categorical}(\omega_{s_i,1}, \dots, \omega_{s_i,K}) & \pi &\sim SB(\alpha_\pi)
\end{aligned} \tag{4.13}$$

We draw attention to the differing properties of the subject-level clustering. In the cicdbc model outlined in Equations 4.9-4.12, all subjects assigned to the same subject cluster shared the exact same regional clusters at all timepoints. In this version of the model, subjects assigned to the same cluster share only the distribution of the proposed clusters, ω_z , and the effects of the covariates on the changepoints, β_z .

When considering the model for $\log\left(\frac{a_t^i}{1-a_t^i}\right) = X_t^i \beta_{s_i}$, it is important to note the covariates are modifying the probability of a changepoint and the not clustering behavior itself. It may be useful to use the absolute first difference of a covariate, $D_{tj}^i = |X_{tj}^i - X_{(t-1)j}^i|$, rather than the raw value itself, as there is more of an expectation of changes in clustering behavior when conditions change in the covariates.

4.2 Posterior Inference

We present a Markov Chain Monte Carlo (MCMC) algorithm for sampling from the posterior distribution: $P(s_{1:N}, c_{1:R,1:T}^{1:Z}, \gamma_{1:R,1:T}^{1:Z}, \pi, \omega_{1:Z}, \mu_{1:K}, \sigma_{1:R}^2 | Y)$. Many of the priors outlined in our model are conjugate, allowing very simple Gibbs sampling updates. The non-standard update lies in sampling from the conditional distribution of γ_{rt}^z . We rely on the independence between γ_{rt}^z and $(c_{ru}^z, \gamma_{rv}^z)$ for $u < v < t$ conditioned the current and previous regional cluster

assignments: $c_{rt}^z, c_{r,t-1}^z$. This allows us to compute the conditional distribution up to a proportionality constant by iteration through $t = 1, \dots, T$. Equation 4.14 gives us the form of the conditional distribution.

$$\begin{aligned}
P(\gamma_{rt}^z | c_{rt}^z, c_{r,t-1}^z, a_t^z, \omega_z) &= P(\gamma_{rt}^z | c_{r,1:t}^z, \gamma_{r,1:(t-1)}^z, a_t^z, \omega_z) \propto \\
&P(\gamma_{rt}^z | a_t^z) P(c_{r,1:t}^z | \gamma_{r,1:t}^z, a_t^z, \omega_z) \propto \\
&P(\gamma_{rt}^z | a_t^z) \prod_{u=1}^t P(c_{r,u}^z | \gamma_{r,u}^z, c_{r,u-1}^z, \omega_z), \tag{4.14}
\end{aligned}$$

where $P(c_{r,u}^z | \gamma_{r,u}^z, \omega_z) = \gamma_{r,u}^z I(c_{r,u}^z = c_{r,u-1}^z) + (1 - \gamma_{r,u}^z) \omega_{z, c_{r,u}^z}$. Notice that for successive t , the product can be recurrently computed by setting $L_{t-1} = \prod_{u=1}^{t-1} P(c_{r,u}^z | \gamma_{r,u}^z, c_{r,u-1}^z, \omega_z)$. Equation 4.14 can then be simplified to

$$P(\gamma_{rt}^z | c_{rt}^z, c_{r,t-1}^z, a_t^z, \omega_z) \propto P(\gamma_{rt}^z | a_t^z) P(c_{r,t}^z | \gamma_{r,t}^z, c_{r,t-1}^z, \omega_z) L_{t-1},$$

with $\gamma_{r1}^z = 0$ by construction, giving us the beginning of our procedure. The $\gamma_{r,2:T}^z$ are then updated through the following steps:

1. Sample γ_{r2}^z from a Bernoulli with success probability: $\frac{a_t^z 2\omega_{z, c_{r2}^z}}{a_t^z 2\omega_{z, c_{r2}^z} + (1 - a_t^z) I(c_{r2}^z = c_{r1}^z)}$
 2. Set $L_2 = P(c_{r2}^z | \gamma_{r2}^z, c_{r1}^z) \omega_{z1}$
 3. Sample γ_{r3}^z by computing $P(\gamma_{r3}^z | -) \propto P(\gamma_{r3}^z | a_3^z) P(c_{r3}^z | c_{r2}, \gamma_{r3}^z) * L_2$ for $\gamma_{r3}^z = 1$ and $\gamma_{r3}^z = 0$, then normalizing.
 4. Set $L_3 = P(c_{r3}^z | c_{r2}, \gamma_{r3}^z) * L_2$
- ⋮ Proceed by iterating between steps t1 and t2.
- t1 Sample γ_{rt}^z by computing $P(\gamma_{rt}^z | -) \propto P(\gamma_{rt}^z | a_t^z) P(c_{rt}^z | c_{r,t-1}, \gamma_{rt}^z) * L_{t-1}$ for $\gamma_{rt}^z = 1$ and $\gamma_{rt}^z = 0$, then normalizing.
 - t2 Set $L_t = P(c_{rt}^z | c_{r,t-1}, \gamma_{rt}^z) * L_{t-1}$

This procedure would be repeated over differing regions r and subject clusters z .

All other parameters are updated through standard Gibbs updates using the conditionally conjugate priors. We outline a brief summary of the steps below:

- **Update μ_k :** We sample from $P(\mu_k|-) = N(\mu_k^*, \sigma_k^{*2})$ where $\mu_k^* = \sigma_k^{*2} \left(\sum_{\{i,t:c_{rt}^i=k\}} \frac{Y_{rt}^i}{\sigma_r^2} + \frac{\mu_0}{\sigma_0^2} \right)$ and

$$\sigma_k^{*2} = \left(\sum_{\{i,t:c_{rt}^i=k\}} \frac{1}{\sigma_r^2} + \frac{1}{\sigma_0^2} \right)^{-1}$$
. Repeat for $k = 1, \dots, K$
- **Update σ_r^2 :** Sample from $P(\sigma_r^2|-) = IG \left(a_\sigma + \frac{T \times N}{2}, b_\sigma + \sum_{t=1}^T \sum_{i=1}^N (Y_{rt}^i - \mu_{c_{rt}^i})^2 \right)$. Repeat for $r = 1, \dots, R$
- **Update π :** Let $M_z = \sum_i I(s^i = z)$. Sample ν_z from $B(1 + M_z, N - \sum_{j=z+1}^Z M_j)$ for $z = 1, \dots, Z - 1$. Set $\pi_1 = \nu_1$ and $\pi_z = \pi_z \prod_{j=1}^{z-1} \nu_j$.
- **Update s_i :** Sample from the discrete distribution $P(s_i = z|-) \propto \prod_r \prod_t \pi_z N(Y_{rt}^i | \mu_{c_{rt}^z}, \sigma_r^2)$. Repeat for $z = 1, \dots, Z$
- **Update c_{rt}^z, γ_{rt}^z :** Similar to sampling s_i , sample c_{rt}^z from the discrete distribution $P(c_{rt}^z = k|-) \propto \prod_{i:s_i=z} \omega_{zk} N(Y_{rt}^i | \mu_k, \sigma_r^2)$. Then update γ_{rt}^z by sampling from the discrete distribution given in Equation 4.14. Repeat for $t = 1, \dots, T$, then for $r = 1, \dots, R$ and $z = 1, \dots, Z$.
- **Update ω_z :** Let $N_{zk} = \sum_r \sum_t I(\gamma_{rt}^z = 0) I(c_{rt}^z = 0)$. Sample ν_{zk} from $B(1 + N_{zk}, \sum_{k=1}^K N_{zk} - \sum_{j=k+1}^K N_{zj})$ for $k = 1, \dots, K - 1$. Set $\pi_1 = \nu_1$ and $\omega_{zk} = \nu_{zk} \prod_{j=1}^{k-1} \nu_{zj}$. Repeat for $k = 1, \dots, K$.

4.2.1 Updating the Covariate-Informed Model

For the Covariate Informed Change point model outlined in 4.13, all Gibbs updates are unchanged except for the conditional distribution of s_i , which is now conditionally independent

of Y , and the newly introduced changepoint regression coefficients β_z . We sample β_z using the Polya-Gamma augmentation introduced by Polson et al. (2013). For the sake of notational simplicity, we assume all unit time series are the same lengths, i.e $T_i = T$ for all units i . Let N_z be the number of units assigned to unit cluster z . Let $X^z = [X_{2:T}^1, \dots, X_{2:T}^{N_z}]'$ be the $N_z(T-1) \times p$ vertical concatenation of the covariates for units belonging to unit cluster z . We construct a working response $g_{it}^z = \sum_{r=1}^R \gamma_{rt}^i$ for $i \in \{i : s_i = z\}$ and $t = 2, \dots, T$. We ignore the first time point since γ_{r1}^i is assumed constant at 0, and is not modeled as a result. Under this construction $g_{it}^z \sim \text{Bin}(R, \frac{1}{1+\exp(-X_t^i \beta_z)})$, where $\text{Bin}(R, p)$ is distribution of a Binomial experiment with R trials and probability of success p . We then sample from the conditional distributions:

$$\begin{aligned} w_{it} | \beta_z &\sim PG(R, X_{it}^z \beta_z) & \beta_z | \vec{g}, w &\sim N(m_w, V_w) \\ V_w &= (X^{z'} \Omega X^z + \Sigma_\beta)^{-1} & m_w &= V_w (X' (\vec{g} - \frac{R}{2})), \end{aligned}$$

where Ω is a $N_z(T-1) \times N_z(T-1)$ diagonal matrix with diagonal entries $\vec{w} = [w_{12}, \dots, w_{N_z T}]$, and $PG(R, X_{it}^z \beta_z)$ is the Polya-Gamma distribution which can be efficiently sampled using the `pgdraw` package in R (Makalic and Schmidt, 2016b).

The other changed component is the discrete full conditional distribution for the unit cluster indicators: s_i .

$$\begin{aligned} P(s_i | c_{1:R,1:T}^i, \gamma_{1:R,1:T}^i, \beta_z) &\propto P(s_i) P(c_{1:R,1:T}^i | \gamma_{1:R,1:T}^i, s_i) P(\gamma_{1:R,1:T}^i | \beta_z, s_i) \\ &\propto \pi_{s_i} \prod_{r,t:\gamma_{rt}^i=0} \omega_{s_i, c_{rt}^i} \prod_{t=1}^T \prod_{r=1}^R (a_t^i)^{\gamma_{rt}^i} (1 - a_t^i)^{1-\gamma_{rt}^i}, \end{aligned}$$

which can be evaluated for $s_i = 1, \dots, Z$ then normalized and sampled from.

4.3 Simulation

We compare the clustering capabilities of our model to that of the Dependent Random Partition Model (drpm) of Page et al. (2021) and the partitions induced by the standard Dirchlet Process Mixture Model (dp). In our first setting we simulate data that exhibits dynamic bi-clustering characteristics. $N = 10$ subjects are assigned to $Z = 3$ unit-clusters. For each unit-cluster $z = 1, 2, 3$, we draw with replacement 8 out of 15 a set of inhabitable signal source clusters, inducing a unit-cluster distribution ω_z over the signal source clusters. Each signal source cluster is paired with a mean μ_k sampled from a centered gaussian with standard deviation 3. We then assign the initial signal-source cluster assignments c_{r1}^z according to ω_z . We then allow the regional clusters to change over time according to the temporal random partition model in 4.9. Once the cluster assignments are completed, our simulation dataset is generated by $Y_{rt}^i \sim N(\mu_{c_{rt}^i}, 1)$.

We evaluate the clustering performance by casting the clustering task as a classification problem. We create a co-clustering indicator $E_{rjit} = I(c_{rt}^i = c^i jt)$ for $j < r = 1, \dots, R; t = 1, \dots, T; i = 1, \dots, N$. We then consider the classifiers $P(c_{rt}^i = c^i jt | Y)$ for each method mentioned above. Then, like evaluating any binary classifier, we can compute the Receiver Operator Curve (ROC) and accompanying Area-Under-the-Curve (AUC), to evaluate the model's ability to cluster regions together that were generated from the same distribution. Secondly we measure the Mean Squared Error in estimating the means of the time series as $MSE = \frac{1}{RNT} \sum_{r,i,t} (\mu_{c_{rt}^i} - \hat{\mu}_{\hat{c}_{rt}^i})^2$.

In Table 4.1 we show the results of the models as applied to 60 simulation datasets under the first setting. Each entry shows the average and standard error for each metric across the simulation repetitions. This highlights the progressive importance of incorporating dependence across axes when such structure is present in the data. A large jump in recovering the true partition structure occurs when assuming dependence in the partitions across time, as shown

by the improvement from dp to drpm. Further improvement is gained by incorporating unit-level dependence, evidenced by the improvement by dbc and dbc-cic. An interesting point to note is that dbc-cic performs just as well as dbc despite the lack of covariates informing the changepoints. This could suggest the assumption of common changepoints across units in the same unit-cluster that is explicit in dbc may be stronger than needed, even when the data are generated under such an assumption.

	dbc	dbc-cic	drpm	dp
AUC	0.99 (0.01)	0.99 (0.01)	0.93 (0.04)	0.85 (0.02)
MSE	0.03 (0.04)	0.18 (0.42)	1.29 (0.88)	0.51 (1.575)

Table 4.1: The clustering AUC and MSE from Simulation 1.

4.3.1 Simulation Setting 2

We simulate from a simplified model of neural connectivity. Our goal is to simulate BOLD timecourses for $N = 9$ subjects and $R = 10$ regions over $T = 200$ timepoints in a block design setting with $K = 2$ stimuli. Broadly, we specify subject cluster assignments and then we group the regions $r = 1, \dots, R$ into clusters. We then simulate the neural activity for each region where regions assigned to the same cluster will tend to fire together. We then convolve the neural activity with the canonical hemodynamic response function $h(t)$, the default in SPM, and add Gaussian noise to create our simulated BOLD dataset for input into our procedure. We outline the data generating model in Model 4.15.

$$\begin{aligned}
 Y_{rt}^i &= \int_{-\infty}^t h(t) y_r^{si}(t - \tau) d\tau + \epsilon_{rt}^i \\
 y_r^{si}(t) | q_C(t), C_{rk}^i = C, x_t = k &\sim q_C(t) \text{Bern}(0.95) + (1 - q_C(t)) \text{Bern}(0.05) \\
 q_C(t) | x_t &\sim (I(x_t = C) \delta(1) + (1 - I(x_t = C)) \text{Bern}(0.3)) \\
 \epsilon_{rt}^i &\sim N(0, 0.5),
 \end{aligned} \tag{4.15}$$

Our first step is to simulate underlying neural firing under three conditions for $x(t) = 0, 1, 2$, where $s(t) = 0$ is our rest condition being active at time t and $x(t) = k$ implies stimulus $k = 1, 2$ is active at time t . We do so by assigning each region to one of five clusters, where this cluster assignment is allowed to change according to the current active stimulus. We denote such an assignment by $C_{rk} = 1, 2, 3$, being the currently assigned cluster of region r when stimulus k is active. We then generate a binary cluster firing sequence $q_C(t) \sim (I(x(t) = C)\delta(1) + (1 - I(x_t = C))Bern(0.3))$. This essentially generates a cluster-stimulus connection that induces a positive psychological effect α_{rk} in the regions r where $C_{rk} = k$. Additionally, the regional changing of the clusters, i.e $C_{r0} \neq C_{rk}, k = 1, 2$ implies a PPI with stimulus k . Next we simulate the regional stimulus sequence, $y_r(t)$ conditional on the cluster firing sequence $q_C(t)$. If the cluster firing indicator is 1, each regional neuron will fire with 95% probability, otherwise firing with 0.05% probability. As this applies to all regions clustered together, this induces a dependence across all regions. Once these neural signals are generated, they are convolved with the canonical hrf before adding Gaussian noise with a standard deviation of 0.1. All subjects that are clustered together share the same neural firing patterns: $y_r^{s_i}(t)$, differing only in their white noise component ϵ_{rt}^i .

	dbc	dbc-cic	drpm	dp
AUC	0.86 (0.04)	0.82 (0.04)	0.78 (0.03)	0.74 (0.02)
MSE	1.14 (0.22)	1.19 (0.22)	1.51 (0.83)	1.02 (0.2)

Table 4.2: The clustering AUC and MSE from Simulation 2.

4.4 Case Study

We apply the DBC model to the handgrip task data from Chapter 2. As a reminder, there are 31 apriori specified ROI measured across 26 participants. Each participant lies in rest in the scanner before being presented with a prompt to raise their arm to the chest. A prompt appears shortly after instructing them to lower their arms. As prompts appear with similar

timings to each subject, there is an expectation of similar change points across subjects, encouraging more grouping. For more detail concerning the experimental design, I refer back to Section 2.3.

4.4.1 Model Specification

We fit the dbc model setting a maximum of $Z = 26$ unit (subject) clusters and $K = 16$ signal source (ROI) clusters. Additionally we set $\alpha_\pi = \beta_\omega = 10^{-3}$ to encourage sparsity in the number of visited subject and ROI clusters. We set the priors on the cluster means $\mu_k \sim N(0, 3)$, regional variances $\sigma_r^2 \sim IG(15, 2)$, and the unit-cluster intercepts $\beta_{0z} \sim N(2, 0.5)$ to encourage strong correlations among successive partitions in the absence of a stimulus. We use the dbc model as opposed to the dbc-cic as subjects perform similar tasks at similar times

4.4.2 Results

The subject clustering is shown in Figure 4.5. There seems to be a strong cluster filled with most subjects with 3 more smaller clusters. Interestingly the clusters seem to be rather insular, exhibiting near 0 posterior probabilities of co-clustering outside of the established cliques. Within the largest cluster, subject cluster 4, ROI clustering results fit within expectations. The ROI clustering is shown in Figure 4.6. Assignments are stable during the rest sections before re-configuring when the raise-arm prompt appears. This is further evidenced by the posterior distribution of γ_{rt} shown in Figure 4.7. The strong vertical black lines, indicating likely cluster assignment re-configurations align with the appearance of the prompts.

Considering the partitions of ROIs during the rest and active portions, we investigate pat-

terns of co-clustering using the Rand Index, a measure of cluster similarity. Let $P_r u = Rand(c_r, c_u)$, where c_r is the collection of cluster assignments for the time points corresponding to either the 'rest' or 'active' portions of the experiment. P^{rest} and P^{active} are shown in Figure 4.8. We see during the rest portions there is tighter coupling among sections in the Default Mode Network. During the active sections the co-clustering are more prevalent among the Salience and Frontal Parietal Control networks.

ROI Cluster	μ_k
1	0.36 (0.33,0.4)
2	-0.38 (-0.42,-0.34)
3	-0.35 (-0.38,-0.31)
4	1.45 (1.33,1.56)
5	-2.14 (-2.32,-1.98)
6	-4.99 (-5.26,-4.74)
7	-1.18 (-1.25,-1.1)
8	0.98 (0.9,1.06)
9	0.1 (0.09,0.12)
10	-2.57 (-3.53,1.16)
11	2.8 (-1.95,4.7)
12	-0.09 (-3.56,4.44)
13	0.28 (-3.64,4.51)
14	-10.38 (-11.05,-9.73)

Table 4.3: The cluster means for each region level cluster c_{rt} .

4.5 Discussion

We have introduced a dynamic bi-clustering approach to account for subject-level heterogeneity in the patterns of brain activation and functional connectivity. We further account for dynamic re-configurations of the clustering partitions through a latent time-varying binary process, capturing change points in the persistence of partitions over time within the same unit cluster. An important property of the proposed model is that subjects are clustered together based on the distribution of latent regional clusters over time. In the case study discussed in Section 4.4, we have obtained a clustering of the ROIs informed by the

average activation in the BOLD signal at any given time. Under this paradigm, two regions who may be functionally connected according to traditional measures of co-variation like correlation and conditional independence patterns may be not grouped together under this model. However, as the time-varying regional clusters are grouped based on a first order basis, the mean, this approach can capture sudden and transient configurations of partitions. The investigation of time-varying clustering methods for FC, sh

Lastly, the scalability of the provided MCMC inference approach is a good opportunity for further improvements. The benefit of clustering approaches to define partitions of the brain is attractive. However, the scale of the thousands of voxels measured in fMRI served as a computational barrier, requiring dimension reduction into ROIs. Allowing the model to operate on the voxel space, rather than the ROI space will allow dbc to form more precise spatial clusters in the brain. This requires a much more scalable inference algorithm than MCMC. A future direction for research could involve using Variational Inference to speed up computation (Blei and Jordan, 2006).




Figure 4.2: Illustration of the dynamic biclustering model with data simulated with 3 (time invariant) unit clusters, 20 (time varying) signal-source clusters and 2 change points of signal-source clusters. The plot shows simulated observations (colored dots) and the gaussian densities from which they are sampled, according to the cluster of the signal source corresponding to the observation. The data is grouped by unit clusters and time intervals with no re-sampling of signal-source clusters. Here, for simplicity, $\sigma_r = 0.5$ for every signal source $r = 1, \dots, R$.




Figure 4.3: Directed graphical representation of our dynamic biclustering model. The distribution of observed signals $y_{rt}^{(i)}$ (gray circles) at time $t = 1, \dots, T$ depends on the cluster s_i of unit $i = 1, \dots, N$ and on the cluster $c_{rt}^{(s_i)}$ of the signal source $r = 1, \dots, R$ among units in cluster s_i . When $s_i = z$, for $z \in \{1, \dots, Z\}$ and $c_{rt}^{(z)} = k \in \{1, \dots, K\}$, then $y_{rt}^{(i)}$ has a Normal distribution with mean μ_k and variance σ_r^2 . Unit clusters are sampled according to finite stick-breaking distribution π and are invariant over time. For units in cluster $s_i = z$, the cluster of signal source r at time t is the same as that at time $t - 1$ when $\gamma_{rt}^{(z)} = 1$ or is re-sampled according to the finite stick-breaking distribution ω_z when $\gamma_{rt}^{(z)} = 0$. For every r and z , $\gamma_{rt}^{(z)} = 1$ with probability a_t . Yellow, diamond-shaped nodes denote hyperparameters with fixed, pre-specified values.




Figure 4.4: ROC for regional co-clustering across time in Simulation setting 2.

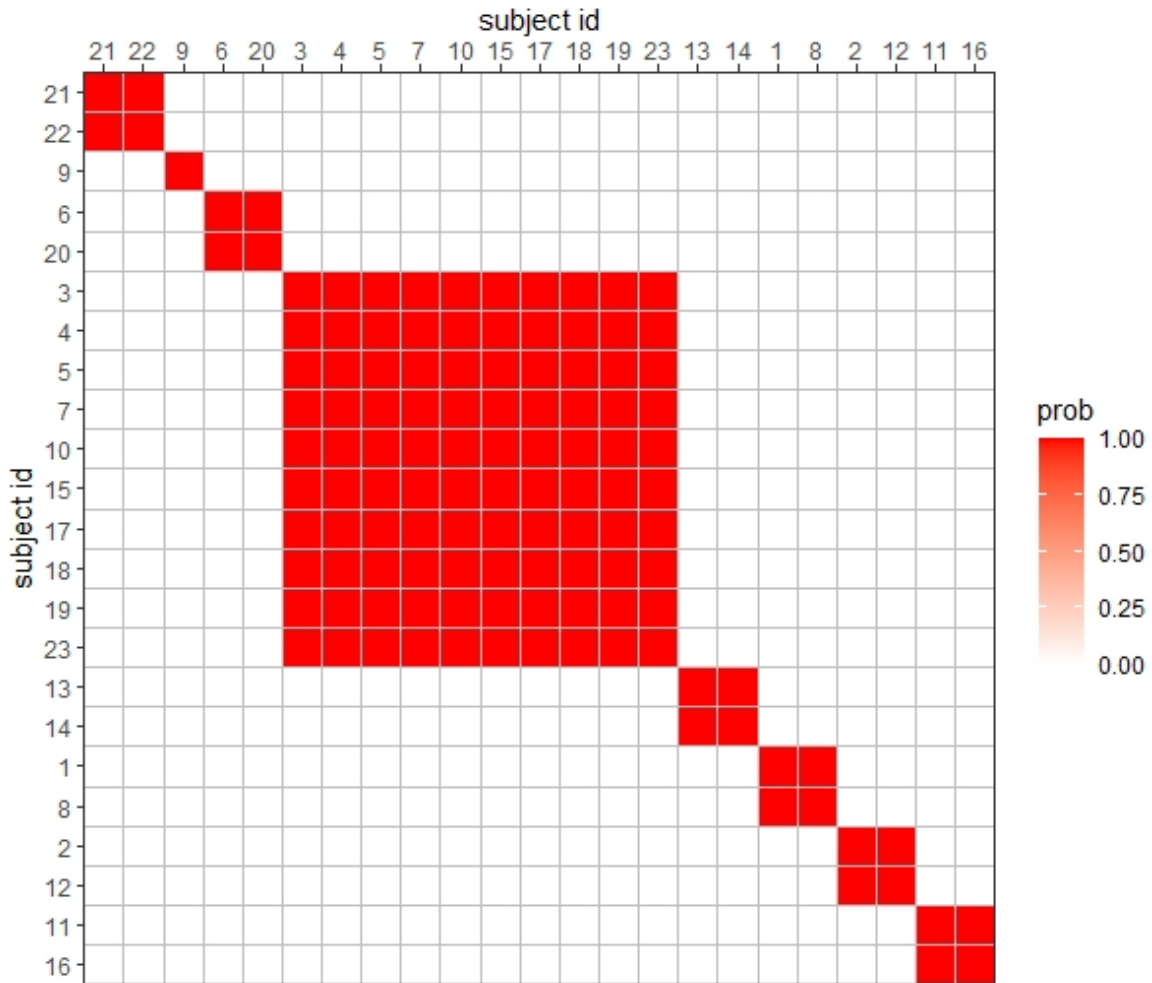


Figure 4.5: The posterior probabilities of subject level co-clustering. Each each i,j entry in this heatmap is given by $P(s_i = s_j|Y)$. Subjects generally fall into either the larger subject cluster 4, or cluster with a single other subject.




Figure 4.6: The regional cluster assignments for subject cluster 4. The color denotes the cluster assignment corresponding to the means in Table 4.3. The vertical dotted line indicates the raising and lowering of the subject’s arm respectively. We see the clusters are stable during the rest sections before reconfiguring during the active portions, aligning with inferences found in Section 2.3.

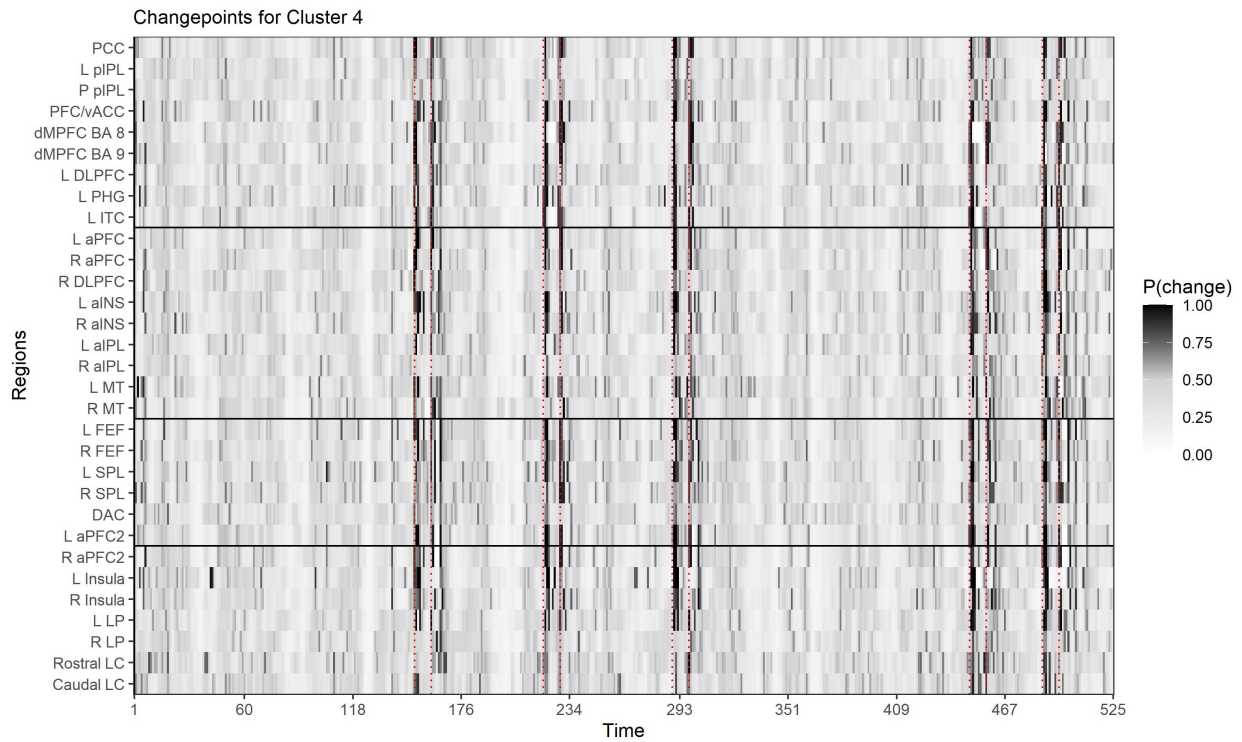


Figure 4.7: The change points for subject cluster 4. Time is along the horizontal and the ROI are along the vertical axis. Each entry is $P(\gamma_{rt} = 0|Y)$. The vertical dotted lines are when the subjects are raising and lowering their arms. We see the model is able to capture the expected change points corresponding to the arm movements.

Chapter 5

Future Directions

In the neurosciences, there has been a surge in interest in studying Dynamic functional connectivity in recent years. Diverse neuroimaging techniques and experimental settings have been used to develop approaches for investigating Dynamic functional connectivity. The projects developed in this Thesis can be expanded in new directions, some of which I outline below.

5.1 Further Data Fusion

In this thesis I introduced 3 state-space models to estimate time-varying FC, connecting FC dynamics in the state-space to observed information; physiological in Chapter 2 and experimental settings in Chapters 3 and 4. This approach opens the door to other forms of data fusion as well. Principle among them is the so-called Microbiota-Gut-Brain axis , the interaction between the diversity of the gut-microbiome and brain health (Cryan et al., 2019). There have been observed differences in FC comparing normally colonized mice and mice raised in a germ-free environment after ischemic stroke. Namely, the germ-free mice exhib-

ited stronger, denser, more global FC networks, indicating less local structural organization (Aswendt et al., 2021). Results like these suggest there are non-dynamic factors that play a role in the FC configurations that the brain can cycle through. An interesting extension to the PIBDFC model of chapter 2 could be to utilize subject level characteristics such as their gut microbiome (Aswendt et al., 2021) or even their gene expression (Arnatkeviciute et al., 2021) to describe between-subject FC heterogeneity. Hoff and Niu (2012) describes a model for the regression of a covariance matrix on covariates, which could be a way to pair together observed gut-microbiome or genetic information with Functional Connectivity.

5.2 Neural Dynamic Connectivity Modeling

The principle interest in studying FC is in gaining an understanding of the patterns of neural coupling among remote regions. One critical assumption that is made throughout this work is that the hemodynamic response function (hrf) is either known or takes a standard form. For example, in the PIBDFC of Chapter 2, I assume either the activation components of the BOLD time series have been removed by fitting the GLM, or the time series has been deconvolved in some form. Both approaches require an assumption about the shape of the hrf. The effects of assuming the incorrect shape can result in improper deconvolutions and resulting TVFC estimates (Lindquist et al., 2009; Gitelman et al., 2003). Similarly, the shape of the hrf is unknown for most regions and is highly variable across regions and people. This uncertainty is completely lost either by assuming the form of the hrf, or treating hemodynamic deconvolution as a pre-processing step. This can be alleviated by having a fully Bayesian process modeling both the shape of the hrf as well as the underlying FC. Such a process would have all uncertainty about the shape of the hrf propagate through to the estimated FC metrics. Some work has been done in Bayesian modeling the of the hrf in activation fMRI studies, most notably the work of Marrelec (2003), where the authors

specify a non-parametric basis representation of the hrf. A future direction involves utilizing a prior on the hrf and in-turn modeling FC as a function of the deconvolved neural activities of interest.

5.3 Scalability

Lastly, I discuss the scalability of the proposed methods. The models were fit using MCMC methods on the order of 30 ROI rather than on the thousands of voxel time series that is measured in fMRI. While this allows this allows gross statements about broad patterns of FC in defined regions of the brain, it does not take advantage of the spatial resolution allowed by fMRI. To allow for more precise modeling of spatial patterns of connectivity, I propose investigating variational inference for state space models as a source for a potentially massive speed-up in computation (Ghahramani and Hinton, 2000). Loaiza-Maya et al. (2022) propose a variational inference approach for non-linear state space models, applying their work to an economics data set of millions of sales over 20,000 customers. Leveraging such a computational approach would be useful for estimating TVFC among the thousands of voxels available.

Bibliography

- B. Alderson-Day, K. Diederer, C. Fernyhough, J. M. Ford, G. Horga, D. S. Margulies, S. McCarthy-Jones, G. Northoff, J. M. Shine, J. Turner, V. van de Ven, R. Van Lutterveld, F. Waters, R. Jardri, Y. Yu, X. Zhou, H. Wang, X. Hu, X. Zhu, L. Xu, C. Zhang, Z. Sun, M. Pannunzi, R. Hindriks, R. G. Bettinardi, E. Wenger, N. Lisofsky, J. Martensson, O. Butler, E. Filevich, M. Becker, M. Lochstet, S. Kühn, G. Deco, J. Wang, X. Zuo, Y. He, A. Anderson, and M. S. Cohen. Graph-based network analysis of resting-state functional MRI. *NeuroImage*, 10(5):1–14, 2016. ISSN 17451701. doi: 10.1093/schbul/sbw078. URL <http://dx.doi.org/10.1016/j.neuroimage.2017.06.006>.
- E. A. Allen, E. Damaraju, S. M. Plis, E. B. Erhardt, T. Eichele, and V. D. Calhoun. Tracking whole-brain connectivity dynamics in the resting state. *Cerebral Cortex*, 24(3):663–676, 2014. ISSN 10473211. doi: 10.1093/cercor/bhs352.
- D. Amar, D. Yekutieli, A. Maron-Katz, T. Hendler, and R. Shamir. A hierarchical Bayesian model for flexible module discovery in three-way time-series data. *Bioinformatics*, 31(12):i17–i26, 2015. ISSN 14602059. doi: 10.1093/bioinformatics/btv228.
- K. S. Ambrosen, K. J. Albers, T. B. Dyrby, M. N. Schmidt, and M. Morup. Nonparametric Bayesian clustering of structural whole brain connectivity in full image resolution. In *2014 International Workshop on Pattern Recognition in Neuroimaging*, pages 1–4. IEEE, 6 2014. ISBN 978-1-4799-4149-0. doi: 10.1109/PRNI.2014.6858507. URL <https://ieeexplore.ieee.org/document/6858507/>.
- A. Anderson and M. S. Cohen. Decreased small-world functional network connectivity and clustering across resting state networks in schizophrenia: An fMRI classification tutorial. *Frontiers in Human Neuroscience*, 7(SEP):1–18, 2013.
- A. Arnatkeviciute, B. D. Fulcher, M. A. Bellgrove, and A. Fornito. Where the genome meets the connectome: Understanding how genes shape human brain connectivity. *NeuroImage*, 244:118570, 12 2021. ISSN 1053-8119. doi: 10.1016/J.NEUROIMAGE.2021.118570.
- M. Aswendt, C. Green, R. Sadler, G. Llovera, L. Dzikowski, S. Heindl, M. Gomez de Agüero, M. Diedenhofen, S. Vogel, F. Wieters, D. Wiedermann, A. Liesz, and M. Hoehn. The gut microbiota modulates brain network connectivity under physiological conditions and after acute brain ischemia. *iScience*, 24(10):103095, 10 2021. ISSN 2589-0042. doi: 10.1016/J.ISCI.2021.103095.

- P. Babaeeghazvini, L. M. Rueda-Delgado, J. Gooijers, S. P. Swinnen, and A. Daffertshofer. Brain Structural and Functional Connectivity: A Review of Combined Works of Diffusion Magnetic Resonance Imaging and Electro-Encephalography. *Frontiers in Human Neuroscience*, 15, 10 2021. ISSN 1662-5161. doi: 10.3389/fnhum.2021.721206. URL <https://www.frontiersin.org/articles/10.3389/fnhum.2021.721206/full>.
- L. E. Baum, T. Petrie, G. Soules, and N. Weiss. A Maximization Technique Occurring in the Statistical Analysis of Probabilistic Functions of Markov Chains. *The Annals of Mathematical Statistics*, 41(1):164–171, 2 1970. ISSN 0003-4851. doi: 10.1214/aoms/1177697196. URL <https://projecteuclid.org/journals/annals-of-mathematical-statistics/volume-41/issue-1/A-Maximization-Technique-Occurring-in-the-Statistical-Analysis-of-Probabilistic/10.1214/aoms/1177697196.full><https://projecteuclid.org/journals/annals-of-mathematical-s>.
- R. Baumgartner, L. Ryner, W. Richter, R. Summers, M. Jarmasz, and R. Somorjai. Comparison of two exploratory data analysis methods for fMRI: Fuzzy clustering vs. principal Component analysis. *Magnetic Resonance Imaging*, 18(1):89–94, 2000. ISSN 0730725X. doi: 10.1016/S0730-725X(99)00102-2.
- M. J. Beal, Z. Ghahramani, and C. E. Rasmussen. The Infinite Hidden {Markov} Model. *NIPS*, 14:577–584, 2002.
- C. F. Beckmann, M. DeLuca, J. T. Devlin, and S. M. Smith. Investigations into resting-state connectivity using independent component analysis. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 360(1457):1001–1013, 2005. ISSN 09628436. doi: 10.1098/RSTB.2005.1634. URL <https://royalsocietypublishing.org/doi/10.1098/rstb.2005.1634>.
- A. Bhadra, J. Datta, Y. Li, N. G. Polson, and B. Willard. Prediction Risk for the Horseshoe Regression. *Journal of Machine Learning Research*, 20(78):1–39, 2019.
- A. Bitto and S. Frühwirth-Schnatter. Achieving shrinkage in a time-varying parameter model framework. *Journal of Econometrics*, 210(1):75–97, 2019. ISSN 0304-4076.
- D. M. Blei and M. I. Jordan. Variational Inference for {D}irichlet Process Mixtures. *Bayesian analysis*, 1(1):121–143, 2006.
- A. M. Bornstein and N. D. Daw. Dissociating hippocampal and striatal contributions to sequential prediction learning. *European Journal of Neuroscience*, 35(7):1011–1023, 4 2012. ISSN 1460-9568. doi: 10.1111/J.1460-9568.2011.07920.X. URL <https://onlinelibrary.wiley.com/doi/full/10.1111/j.1460-9568.2011.07920.x><https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1460-9568.2011.07920.x><https://onlinelibrary.wiley.com/doi/10.1111/j.1460-9568.2011.07920.x>.
- A. M. Bornstein and N. D. Daw. Cortical and hippocampal correlates of deliberation during model-based decisions for rewards in humans. *PLoS computational biology*, 9(12):e1003387, 2013.

- A. M. Bornstein and K. A. Norman. Reinstated episodic context guides sampling-based decisions for reward. *Nature neuroscience*, 20(7):997–1003, 2017.
- F. Bowman. Spatiotemporal models for region of interest analyses of functional neuroimaging data. *Journal of the American Statistical Association*, 102(478):442–453, 2007.
- F. D. B. Bowman. Brain imaging analysis. *Annual Review of Statistics and Its Application*, 1:61–85, 2014. ISSN 2326831X. doi: 10.1146/annurev-statistics-022513-115611.
- K. Bush and J. Cisler. Decoding neural events from fMRI BOLD signal: A comparison of existing approaches and development of a new algorithm. *Magnetic Resonance Imaging*, 31(6):976–989, 7 2013. ISSN 0730725X. doi: 10.1016/j.mri.2013.03.015. URL <https://linkinghub.elsevier.com/retrieve/pii/S0730725X1300101X>.
- A. Cadonna, S. Frühwirth-Schnatter, and P. Knaus. Triple the gamma—a unifying shrinkage prior for variance and variable selection in sparse state space and TVP models. *Econometrics*, 8(2):0–36, 2020. ISSN 22251146. doi: 10.3390/econometrics8020020.
- V. D. Calhoun, R. Miller, G. Pearlson, and T. Adali. The Chronnectome: Time-Varying Connectivity Networks as the Next Frontier in fMRI Data Discovery. *Neuron*, 84(2):262–274, 2014. ISSN 10974199. doi: 10.1016/j.neuron.2014.10.015.
- C. M. Carvalho, N. G. Polson, and J. G. Scott. The horseshoe estimator for sparse signals. *Biometrika*, 97(2):465–480, 2010.
- G. Casella and R. L. Berger. *Statistical Inference*. Duxbury, Belmont, CA, 15 edition, 2002. ISBN 0534243126.
- N. K. Chandra and S. Bhattacharya. Non-marginal decisions: A novel bayesian multiple testing procedure. *Electronic Journal of Statistics*, 13(1):489–535, 2019. ISSN 19357524. doi: 10.1214/19-EJS1535.
- N. K. Chandra, P. Mueller, and A. Sarkar. Bayesian Scalable Precision Factor Analysis for Massive Sparse Gaussian Graphical Models, 2021.
- C. Chang and G. H. Glover. Time-frequency dynamics of resting-state brain connectivity measured with fMRI. *NeuroImage*, 50(1):81–98, 2010. ISSN 10538119. doi: 10.1016/j.neuroimage.2009.12.011.
- C. Chang and G. H. Glover. Time-frequency dynamics of resting-state brain connectivity measured with fMRI. *Neuroimage*, 50(1):81–98, 2011. doi: 10.1016/j.neuroimage.2009.12.011.Time-frequency.
- Y. Chen, Y. N. Liu, P. Zhou, X. Zhang, Q. Wu, X. Zhao, and D. Ming. The transitions between dynamic micro-states reveal age-related functional network reorganization. *Frontiers in Physiology*, 10(JAN):1–10, 2019.

- S. Chiang, E. R. Vankov, H. J. Yeh, M. Guindani, M. Vannucci, Z. Haneef, J. M. Stern, D. Ph, E. R. Vankov, D. Ph, H. J. Yeh, M. Guindani, D. Ph, M. Vannucci, D. Ph, Z. Haneef, and J. M. Stern. Temporal and spectral characteristics of dynamic functional connectivity between resting-state networks reveal information beyond static connectivity. *016672(1):1–25*, 2018. ISSN 19326203. doi: 10.7910/DVN/NOSYLW.
- R. C. Craddock, P. E. Holtzheimer, X. P. Hu, and H. S. Mayberg. Disease state prediction from resting state functional connectivity. *Magnetic Resonance in Medicine*, 62(6):1619–1628, 12 2009. ISSN 07403194. doi: 10.1002/mrm.22159. URL [www.interscience.https://onlinelibrary.wiley.com/doi/10.1002/mrm.22159](http://onlinelibrary.wiley.com/doi/10.1002/mrm.22159).
- I. Cribben, R. Haraldsdottir, L. Y. Atlas, T. D. Wager, and M. A. Lindquist. Dynamic connectivity regression: Determining state-related changes in brain connectivity. *NeuroImage*, 61(4):907–920, 2012. ISSN 10538119. doi: 10.1016/j.neuroimage.2012.03.070. URL <http://dx.doi.org/10.1016/j.neuroimage.2012.03.070>.
- J. F. Cryan, K. J. O, C. S. M Cowan, K. V. Sandhu, T. F. S Bastiaanssen, M. Boehme, M. G. Codagnone, S. Cusotto, C. Fulling, A. V. Golubeva, K. E. Guzzetta, M. Jaggar, C. M. Long-Smith, J. M. Lyte, J. A. Martin, A. Molinero-Perez, G. Moloney, E. Morelli, E. Morillas, J. S. Cruz-Pereira, V. L. Peterson, K. Rea, N. L. Ritz, E. Sherwin, S. Spichak, E. M. Teichman, M. van de Wouw, A. Paula Ventura-Silva, S. E. Wallace-Fitzsimons, N. Hyland, G. Clarke, and T. G. Dinan. The Microbiota-Gut-Brain Axis. *Physiol Rev*, 99:1877–2013, 2019. doi: 10.1152/physrev.00018.2018.-The. URL www.prv.org.
- E. Damaraju, E. A. Allen, A. Belger, J. M. Ford, S. McEwen, D. H. Mathalon, B. A. Mueller, G. D. Pearlson, S. G. Potkin, A. Preda, J. A. Turner, J. G. Vaidya, T. G. Van Erp, and V. D. Calhoun. Dynamic functional connectivity analysis reveals transient states of dysconnectivity in schizophrenia. *NeuroImage: Clinical*, 5(July):298–308, 2014.
- N. D. Daw and D. Shohamy. The Cognitive Neuroscience of Motivation and Learning. *Social Cognition*, 26(5):593–620, 10 2008. ISSN 0278-016X. doi: 10.1521/soco.2008.26.5.593. URL <http://guilfordjournals.com/doi/10.1521/soco.2008.26.5.593>.
- G. Deshpande, S. LaConte, G. A. James, S. Peltier, and X. Hu. Multivariate {Granger} causality analysis of {fMRI} data. *Hum. Brain Mapp.*, 30(4):1361–1373, 4 2009.
- G. Deshpande, P. Santhanam, and X. Hu. Instantaneous and causal connectivity in resting state brain networks derived from functional MRI data. *NeuroImage*, 54(2):1043–1052, 1 2011. ISSN 10538119. doi: 10.1016/j.neuroimage.2010.09.024.
- I. S. Dhillon. Co-clustering documents and words using bipartite spectral graph partitioning. *Proceedings of the Seventh ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 269–274, 2001. doi: 10.1145/502512.502550.
- X. Di, R. C. Reynolds, and B. B. Biswal. Imperfect (de)convolution may introduce spurious psychophysiological interactions and how to avoid it. *Human Brain Mapping*, 38(4):1723–1740, 4 2017. ISSN 10659471. doi: 10.1002/hbm.23413. URL <https://onlinelibrary.wiley.com/doi/10.1002/hbm.23413>.

- X. Di, Z. Zhang, and B. B. Biswal. Understanding psychophysiological interaction and its relations to beta series correlation. *Brain Imaging and Behavior*, 15(2):958–973, 2021.
- B. B. Doll, D. Shohamy, and N. D. Daw. Multiple memory systems as substrates for multiple decision systems. *Neurobiology of learning and memory*, 117:4–13, 2015.
- H. Eavani, T. D. Satterthwaite, R. C. R. E. Gur, R. C. R. E. Gur, and C. Davatzikos. Unsupervised learning of functional network dynamics in resting state fMRI. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 7917 LNCS:426–437, 2013. ISSN 03029743. doi: 10.1007/978-3-642-38868-2_{_}36. URL http://link.springer.com/10.1007/978-3-642-38868-2_36.
- M. El-Baba, D. J. Lewis, Z. Fang, A. M. Owen, S. M. Fogel, and J. B. Morton. Functional connectivity dynamics slow with descent from wakefulness to sleep. *PLoS ONE*, 14(12): 1–13, 2019. ISSN 19326203. doi: 10.1371/journal.pone.0224669. URL <http://dx.doi.org/10.1371/journal.pone.0224669>.
- R. Epstein and N. Kanwisher. A cortical representation of the local visual environment. *Nature*, 392(6676):598–601, 1998.
- E. B. Erhardt, E. A. Allen, Y. Wei, T. Eichele, and V. D. Calhoun. SimTB, a simulation toolbox for fMRI data under a model of spatiotemporal separability. *NeuroImage*, 4(59): 4160–4167, 2012.
- V. Essen and G. Tononi. *An Introduction to Brain Networks*. 2016.
- D. A. Fair, B. L. Schlaggar, A. L. Cohen, F. M. Miezin, N. U. F. Dosenbach, K. K. Wenger, M. D. Fox, A. Z. Snyder, M. E. Raichle, and S. E. Petersen. A method for using blocked and event-related fMRI data to study “resting state” functional connectivity. *NeuroImage*, 35(1):396–405, 2007.
- T. S. Ferguson. A Bayesian Analysis of Some Nonparametric Problems. *The Annals of Statistics*, 1(2), 3 1973. ISSN 0090-5364. doi: 10.1214/aos/1176342360.
- A. Fornito, A. Zalesky, and M. Breakspear. Graph analysis of the human connectome: Promise, progress, and pitfalls. *NeuroImage*, 80:426–444, 2013. ISSN 10538119. doi: 10.1016/j.neuroimage.2013.04.087. URL <http://dx.doi.org/10.1016/j.neuroimage.2013.04.087>.
- E. B. Fox, E. B. Sudderth, M. I. Jordan, and A. S. Willsky. A sticky HDP-HMM with application to speaker diarization. *Annals of Applied Statistics*, 5(2 A):1020–1056, 2011. ISSN 19326157. doi: 10.1214/10-AOAS395.
- J. Friedman, T. Hastie, and R. Tibshirani. Sparse covariance estimation. *Biostatistics*, 9 (3):432–441, 2008. doi: 10.1093/biostatistics/kxm045. URL <http://statweb.stanford.edu/~tibs/ftp/graph.pdf>.

- K. J. Friston. Functional and effective connectivity in neuroimaging: a synthesis. *Human Brain Mapping*, 2:56–78, 1994.
- K. J. Friston. Functional and Effective Connectivity: A Review. *Brain Connectivity*, 1(1): 13–36, 2011.
- K. J. Friston, C. Buechel, G. R. Fink, J. Morris, E. Rolls, and J. Dolan. Psychophysiological and Modulatory Interactions in Neuroimaging. *NeuroImage*, 6(6):218–229, 1997.
- S. Frühwirth-Schnatter and G. Malsiner-Walli. From here to infinity: sparse finite versus Dirichlet process mixtures in model-based clustering. *Advances in Data Analysis and Classification*, 13(1):33–64, 3 2019. ISSN 1862-5347. doi: 10.1007/s11634-018-0329-y. URL <https://doi.org/10.1007/s11634-018-0329-y><http://link.springer.com/10.1007/s11634-018-0329-y>.
- S. Frühwirth-Schnatter and H. Wagner. Stochastic model specification search for Gaussian and partial non-Gaussian state space models. *Journal of Econometrics*, 154(1):85–100, 2010. ISSN 03044076. doi: 10.1016/j.jeconom.2009.07.003. URL <http://dx.doi.org/10.1016/j.jeconom.2009.07.003>.
- E. I. George and R. E. McCulloch. Variable selection via Gibbs sampling. *Journal of the American Statistical Association*, 88(423):881–889, 1993. ISSN 1537274X. doi: 10.1080/01621459.1993.10476353. URL <http://www.jstor.org/stable/2290777?origin=crossref>.
- E. I. George and R. E. McCulloch. Approaches for {Bayesian} Variable Selection. *Statistica Sinica*, 7(2):339–373, 1997.
- Z. Ghahramani and G. E. Hinton. Variational Learning for Switching State-Space Models. *Neural Computation*, 12(4):831–864, 4 2000. ISSN 0899-7667. doi: 10.1162/089976600300015619.
- D. R. Gitelman, W. D. Penny, J. Ashburner, and K. J. Friston. Modeling regional and psychophysiological interactions in fMRI: The importance of hemodynamic deconvolution. *NeuroImage*, 19(1):200–207, 2003. ISSN 10538119. doi: 10.1016/S1053-8119(03)00058-2.
- X. Golay, S. Kollias, G. Stoll, D. Meier, A. Valavanis, and P. Boesiger. A new correlation-based fuzzy logic clustering algorithm for FMRI. *Magnetic Resonance in Medicine*, 40(2):249–260, 8 1998. ISSN 07403194. doi: 10.1002/mrm.1910400211. URL <https://onlinelibrary.wiley.com/doi/10.1002/mrm.1910400211>.
- M. B. Gordy. A generalization of generalized beta distributions. Technical Report 1998-18, Board of Governors of the Federal Reserve System (U.S.), 1998.
- S. T. Grafton, J. Sutton, W. Couldwell, M. Lew, and C. Waters. Network Analysis of Motor System Connectivity in Parkinson’s Disease: Modulation of Thalamocortical Interactions After Pallidotomy. *Human Brain Mapping*, 2:4, 1994.

- M. D. Greicius, G. Srivastava, A. L. Reiss, and V. Menon. Default-mode network activity distinguishes Alzheimer’s disease from healthy aging: Evidence from functional MRI. *Proceedings of the National Academy of Sciences*, 101(13):4637–4642, 3 2004. ISSN 0027-8424. doi: 10.1073/pnas.0308627101.
- R. Guigourès, Marc Boullé, F. Rossi, and M. Boullé. Discovering patterns in time-varying graphs: a triclustering approach. *Adv Data Anal Classif*, 12:509–536, 2018. doi: 10.1007/s11634-015-0218-6. URL <https://doi.org/10.1007/s11634-015-0218-6>.
- S. Herculano-Houzel. The human brain in numbers: A linearly scaled-up primate brain. *Frontiers in Human Neuroscience*, 3(NOV):1–11, 2009. ISSN 16625161. doi: 10.3389/neuro.09.031.2009.
- P. D. Hoff and X. Niu. A Covariance Regression Model. *Statistica Sinica*, 22(2):729–753, 2012.
- C. C. Holmes and L. Held. Bayesian auxiliary variable models for binary and multinomial regression. *Bayesian Analysis*, 1(1):145–168, 2006. ISSN 19360975. doi: 10.1214/06-BA105.
- T. Holsclaw, A. M. Greene, A. W. Robertson, and P. Smyth. Bayesian nonhomogeneous Markov models via Pólya-Gamma data augmentation with applications to rainfall modeling. *Annals of Applied Statistics*, 11(1):393–426, 2017.
- J. M. Houck, M. S. Çetin, J. R. Bustillo, C. Aine, J. Cañive, N. Perrone-Bizzozero, R. J. Thoma, J. M. Houck, M. S. Çetin, A. R. Mayer, J. Stephen, C. Aine, V. D. Calhoun, J. R. Bustillo, J. Cañive, R. J. Thoma, M. J. Brookes, C. Aine, M. J. Brookes, M. S. Çetin, and N. Perrone-Bizzozero. Magnetoencephalographic and functional MRI connectomics in schizophrenia via intra- and inter-network connectivity. *NeuroImage*, 145(April 2016): 96–106, 2017.
- S. Hussain, I. Menchaca, M. A. Shalchy, K. Yaghoubi, J. Langley, A. R. Seitz, X. P. Hu, and M. A. K. Peters. Locus coeruleus neuromelanin predicts ease of attaining and maintaining neural states of arousal. *bioRxiv*, 2022.
- R. M. Hutchison, T. Womelsdorf, E. A. Allen, P. A. Bandettini, V. D. Calhoun, M. Corbetta, S. Della Penna, J. H. Duyn, G. H. Glover, J. Gonzalez-Castillo, D. A. Handwerker, S. Keilholz, V. Kiviniemi, D. A. Leopold, F. de Pasquale, O. Sporns, M. Walter, and C. Chang. Dynamic functional connectivity: Promise, issues, and interpretations. *NeuroImage*, 80:360–378, 2013. ISSN 10538119. doi: 10.1016/j.neuroimage.2013.05.079. URL <http://dx.doi.org/10.1016/j.neuroimage.2013.05.079>.
- A. Iraj, T. P. Deramus, N. Lewis, M. Yaesoubi, J. M. Stephen, E. Erhardt, A. Belger, J. M. Ford, S. McEwen, D. H. Mathalon, B. A. Mueller, G. D. Pearson, S. G. Potkin, A. Preda, J. A. Turner, J. G. Vaidya, T. G. M. van Erp, and V. D. Calhoun. The spatial chronnectome reveals a dynamic interplay between functional segregation and integration. *Human Brain Mapping*, 40(10):3058–3077, 2019. ISSN 10970193. doi: 10.1002/hbm.24580.

- A. Johnson and A. D. Redish. Neural ensembles in CA3 transiently encode paths forward of the animal at a decision point. *Journal of Neuroscience*, 27(45):12176–12189, 2007.
- A. Johnson, M. A. A. van der Meer, and A. D. Redish. Integrating hippocampus and striatum in decision-making. *Current opinion in neurobiology*, 17(6):692–697, 2007.
- S. Joshi and J. I. Gold. Context-dependent relationships between locus coeruleus firing patterns and coordinated neural activity in the anterior cingulate cortex. *eLife*, 11:e63490, 2022.
- S. Joshi, Y. Li, R. M. Kalwani, and J. I. Gold. Relationships between {Pupil} {Diameter} and {Neuronal} {Activity} in the {Locus} {Coeruleus}, {Colliculi}, and {Cingulate} {Cortex}. *Neuron*, 89(1):221–234, 2016.
- S. Kim, P. Smyth, and H. Stern. A Bayesian Mixture Approach to Modeling Spatial Activation Patterns in Multisite fMRI Data. *IEEE Transactions on Medical Imaging*, 29(6):1260–1274, 6 2010. ISSN 0278-0062. doi: 10.1109/TMI.2010.2044045.
- P. Knaus, A. Bitto-Nemling, A. Cadonna, and S. Frühwirth-Schnatter. Shrinkage in the Time-Varying Parameter Model Framework Using the R Package shrinkTVP. *Journal of Statistical Software*, 100(13), 2021. ISSN 15487660. doi: 10.18637/JSS.V100.I13.
- D. R. Kowal, D. S. Matteson, and D. Ruppert. Dynamic shrinkage processes. *Journal of the Royal Statistical Society. Series B: Statistical Methodology*, 81(4):781–804, 2019. ISSN 14679868. doi: 10.1111/rssb.12325.
- A. Kucyi, M. J. Hove, M. Esterman, R. M. Hutchison, and E. M. Valera. Dynamic Brain Network Correlates of Spontaneous Fluctuations in Attention. *Cerebral cortex (New York, N.Y. : 1991)*, 27(3):1831–1840, 2017. ISSN 14602199. doi: 10.1093/cercor/bhw029.
- S. Kundu, J. Ming, J. Pierce, J. McDowell, and Y. Guo. Estimating dynamic brain functional networks using multi-subject fMRI data. *NeuroImage*, 183(5):635–649, 12 2018. ISSN 10538119. doi: 10.1016/j.neuroimage.2018.07.045. URL <https://linkinghub.elsevier.com/retrieve/pii/S1053811918306608>.
- A. R. Laird, J. L. Lancaster, and P. T. Fox. {BrainMap}: {The} {Social} {Evolution} of a {Human} {Brain} {Mapping} {Database}. *Neuroinformatics*, 3(1):65–78, 2005.
- J. L. Lancaster, D. Tordesillas-Gutiérrez, M. Martinez, F. Salinas, A. Evans, K. Zilles, J. C. Mazziotta, and P. T. Fox. Bias between {MNI} and {Talairach} coordinates analyzed using the {ICBM}-152 brain template. *Hum. Brain Mapp.*, 28(11):1194–1205, 11 2007.
- J. Langley, S. Hussain, J. J. Flores, I. J. Bennett, and X. Hu. Characterization of age-related microstructural changes in locus coeruleus and substantia nigra pars compacta. *Neurobiology of Aging*, 87:89–97, 3 2020.
- S. L. Lauritzen. *Graphical Models*. Oxford University Press, Oxford, 1996.

- K.-J. Lee, G. L. Jones, B. S. Caffo, and S. S. Bassett. Spatial Bayesian Variable Selection Models on Functional Magnetic Resonance Imaging Time-Series Data. *Bayesian Analysis*, 9(3):699–732, 2014. doi: 10.1214/14-BA873.
- B. C. Lehmann, S. R. White, R. N. Henson, Cam-CAN, and L. Geerligs. Assessing dynamic functional connectivity in heterogeneous samples. *NeuroImage*, 157:635–647, 8 2017. ISSN 1053-8119. doi: 10.1016/J.NEUROIMAGE.2017.05.065.
- N. Leonardi and D. Van De Ville. On spurious and real fluctuations of dynamic functional connectivity during rest, 2015. ISSN 10959572.
- F. Li and N. R. Zhang. Bayesian variable selection in structured high-dimensional covariate spaces with applications in genomics. *Journal of the American Statistical Association*, 105(491):1202–1214, 2010. ISSN 01621459. doi: 10.1198/jasa.2010.tm08177. URL <https://doi.org/10.1198/jasa.2010.tm08177>.
- K. Li, L. Guo, J. Nie, G. Li, and T. Liu. Review of methods for functional brain connectivity detection using fMRI. *Computerized Medical Imaging and Graphics*, 33(2):131–139, 2009. ISSN 08956111. doi: 10.1016/j.compmedimag.2008.10.011.
- Y. Li, B. A. Craig, and A. Bhadra. The Graphical Horseshoe Estimator for Inverse Covariance Matrices. *Journal of Computational and Graphical Statistics*, 28(3):747–757, 2019.
- Y. O. Li, T. Adali, and V. D. Calhoun. Estimating the number of independent components for functional magnetic resonance imaging data. *Human Brain Mapping*, 28(11):1251–1266, 2007. ISSN 10659471. doi: 10.1002/hbm.20359.
- M. A. Lindquist. The Statistical Analysis of fMRI Data. *Statistical Science*, 23(4):439–464, 2008. ISSN 08834237. doi: 10.1214/09-STS282.
- M. A. Lindquist, J. M. Loh, L. Y. Atlas, and T. D. Wager. Modeling the Hemodynamic Response Function in {fMRI}: Efficiency, Bias, and Mis-modeling. *NeuroImage*, 45:187–198, 2009.
- M. A. Lindquist, Y. Xu, M. B. Nebel, and B. S. Caffo. Evaluating Dynamic Bivariate Correlations in Resting-state fMRI: A comparison study and a new approach. *NeuroImage*, 101(1):531–546, 2014. ISSN 15378276. doi: 10.1016/j.neuroimage.2014.06.052.
- R. Loaiza-Maya, M. S. Smith, D. J. Nott, and P. J. Danaher. Fast and accurate variational inference for models with many latent variables. *Journal of Econometrics*, 230(2):339–362, 10 2022. ISSN 03044076. doi: 10.1016/j.jeconom.2021.05.002.
- D. Lombardo, C. Cassé-Perrot, J. P. Ranjeva, A. Le Troter, M. Guye, J. Wirsich, P. Payoux, D. Bartrés-Faz, R. Bordet, J. C. Richardson, O. Felician, V. Jirsa, O. Blin, M. Didic, and D. Battaglia. Modular slowing of resting-state dynamic functional connectivity as a marker of cognitive dysfunction induced by sleep deprivation. *NeuroImage*, 222(May), 2020. ISSN 10959572. doi: 10.1016/j.neuroimage.2020.117155.

- M. J. Lowe, M. D. Phillips, J. T. Lurito, D. Mattson, M. Dzemidzic, and V. P. Mathews. Multiple sclerosis: Low-frequency temporal blood oxygen level-dependent fluctuations indicate reduced functional connectivity - Initial results. *Radiology*, 224(1):184–192, 2002.
- J. Lukemire, Y. Wang, A. Verma, and Y. Guo. HINT: A Hierarchical Independent Component Analysis Toolbox for Investigating Brain Functional Networks using Neuroimaging Data. pages 1–24, 2018. URL <http://arxiv.org/abs/1803.07587>.
- H. Lv, Z. Wang, E. Tong, L. M. Williams, G. Zaharchuk, M. Zeineh, A. N. Goldstein-Piekarski, T. M. Ball, C. Liao, and M. Wintermark. Resting-State Functional MRI: Everything That Nonexperts Have Always Wanted to Know. *American Journal of Neuroradiology*, 39(8):1390–1399, 2018. ISSN 0195-6108.
- M.-e. Lynall, D. S. Bassett, R. Kerwin, P. J. Mckenna, U. Müller, and E. Bullmore. Functional connectivity and brain networks in schizophrenia. *Journal of Neuroscience*, 30(28):9477–9487, 2010. ISSN 1529-2401. doi: 10.1523/JNEUROSCI.0333-10.2010.Functional.
- S. Madeira and A. Oliveira. Biclustering algorithms for biological data analysis: a survey. *IEEE/ACM Transactions on Computational Biology and Bioinformatics*, 1(1):24–45, 1 2004. ISSN 1545-5963. doi: 10.1109/TCBB.2004.2. URL <http://ieeexplore.ieee.org/document/1324618/>.
- E. Makalic and D. F. Schmidt. A simple sampler for the horseshoe estimator. *IEEE Signal Processing Letters*, 23(1):179–182, 2016a. ISSN 10709908. doi: 10.1109/LSP.2015.2503725.
- E. Makalic and D. F. Schmidt. High-Dimensional Bayesian Regularised Regression with the BayesReg Package. 11 2016b. URL <http://arxiv.org/abs/1611.06649>.
- G. Malsiner-Walli, S. Frühwirth-Schnatter, B. Grün, G. Malsiner-Walli, B. Grün, B. Grün, and S. Frühwirth-Schnatter. Model-based clustering based on sparse finite Gaussian mixtures. *Statistics and Computing*, 26:303–324, 2016. doi: 10.1007/s11222-014-9500-2.
- M. Markicevic, I. Savvateev, C. Grimm, and V. Zerbi. Emerging imaging methods to study whole-brain function in rodent models. *Translational Psychiatry*, 11(1), 2021. ISSN 21583188. doi: 10.1038/s41398-021-01575-5. URL <http://dx.doi.org/10.1038/s41398-021-01575-5>.
- G. Marrelec. Bayesian estimation of the hemodynamic response function in functional MRI. pages 229–247, 2003.
- M. Mather, R. Huang, D. Clewett, S. E. Nielsen, R. Velasco, K. Tu, S. Han, and B. L. Kennedy. Isometric exercise facilitates attention to salient events in women via the noradrenergic system. *NeuroImage*, 210:116560, 4 2020.
- D. G. McLaren, M. L. Ries, G. Xu, and S. C. Johnson. A generalized form of context-dependent psychophysiological interactions (gPPI): A comparison to standard approaches. *NeuroImage*, 61(4):1277–1286, 2012. ISSN 10538119. doi: 10.1016/j.neuroimage.2012.03.068. URL <http://dx.doi.org/10.1016/j.neuroimage.2012.03.068>.

- V. Menon and L. Q. Uddin. Saliency, switching, attention and control: a network model of insula function. *Brain Struct Funct*, 214(5-6):655–667, 6 2010.
- R. J. Meszlényi, K. Buza, and Z. Vidnyánszky. Resting State fMRI Functional Connectivity-Based Classification Using a Convolutional Neural Network Architecture. *Frontiers in Neuroinformatics*, 11(October):1–12, 2017. ISSN 1662-5196. doi: 10.3389/fninf.2017.00061. URL <http://journal.frontiersin.org/article/10.3389/fninf.2017.00061/full>.
- P. Müller, G. Parmigiani, and K. Rice. *{FDR} and {B}ayesian multiple comparisons rules*. bepress, 2006.
- P. Müller, G. Parmigiani, and K. Rice. *{FDR} and {B}ayesian Multiple Comparisons Rules*. In J. M. Bernardo, M. J. Bayarri, J. O. Berger, A. P. Dawid, D. Heckerman, A. F. M. Smith, and M. West, editors, *Bayesian Statistics 8*. Oxford, UK: Oxford University Press., 2007.
- M. A. Newton, A. Noueir, D. Sarkar, and P. Ahlquist. Detecting differential gene expression with a semiparametric hierarchical mixture method. *Biostatistics*, 5:155–176, 2004.
- J. A. Nielsen, B. A. Zielinski, P. T. Fletcher, A. L. Alexander, N. Lange, E. D. Bigler, J. E. Lainhart, and J. S. Anderson. Multisite functional connectivity MRI classification of autism: ABIDE results. *Frontiers in Human Neuroscience*, 7, 2013.
- G. L. Page, F. A. Quintana, and D. B. Dahl. Dependent Modeling of Temporal Sequences of Random Partitions. *Journal of Computational and Graphical Statistics*, 0(0):1–29, 2021. ISSN 15372715. doi: 10.1080/10618600.2021.1987255. URL <https://doi.org/10.1080/10618600.2021.1987255>.
- V. Pariyadath, E. A. Stein, and T. J. Ross. Machine learning classification of resting state functional connectivity predicts smoking status. *Frontiers in Human Neuroscience*, 8 (June):1–10, 2014.
- C. Peterson, F. C. Stingo, and M. Vannucci. *{Bayesian} Inference of Multiple Gaussian Graphical Models*. *Journal of the American Statistical Association*, 110(509):159–174, 2015. ISSN 1537274X. doi: 10.1080/01621459.2014.896806.
- R. A. Poldrack, J. Clark, E. a. Paré-Blagoev, D. Shohamy, J. Creso Moyano, C. Myers, and M. A. Gluck. Interactive memory systems in the human brain. *Nature*, 414(6863):546–550, 2001.
- N. G. Polson, J. G. Scott, and J. Windle. Bayesian inference for logistic models using Pólya-Gamma latent variables. *Journal of the American Statistical Association*, 108(504): 1339–1349, 2013. ISSN 1537274X. doi: 10.1080/01621459.2013.829001.
- E. Premi, S. Gazzina, M. Diano, A. Girelli, V. D. Calhoun, A. Iraj, Q. Gong, K. Li, F. Cauda, R. Gasparotti, A. Padovani, B. Borroni, and M. Magoni. Enhanced dynamic functional connectivity (whole-brain chronnectome) in chess experts. *Scientific Reports*, 10(1):1–8,

2020. ISSN 20452322. doi: 10.1038/s41598-020-63984-8. URL <http://dx.doi.org/10.1038/s41598-020-63984-8>.
- M. G. Preti, T. A. Bolton, and D. Van De Ville. The dynamic functional connectome: State-of-the-art and perspectives. *NeuroImage*, 2017. ISSN 10959572. doi: 10.1016/j.neuroimage.2016.12.061.
- M. E. Raichle. The {Restless} {Brain}. *Brain Connectivity*, 1(1):3–12, 1 2011.
- M. Ramezani, A. Heidari, E. Fatemizadeh, and H. Soltanian-Zadeh. Spectral clustering approach with sparsifying technique for functional connectivity detection in the resting brain. In *2010 International Conference on Intelligent and Advanced Systems*, pages 1–5. IEEE, 6 2010. ISBN 978-1-4244-6623-8. doi: 10.1109/ICIAS.2010.5716164.
- V. Rockova and K. McAlinn. Dynamic Variable Selection with Spike-and-Slab Process Priors. *Bayesian Analysis*, 16(1):233–269, 2020. ISSN 19316690. doi: 10.1214/20-BA1199.
- S. L. Scott. Bayesian Methods for Hidden Markov Models: Recursive Computing in the 21st Century. *Journal of the American Statistical Association*, 97(457):337–351, 2002.
- S. Shakil, C. H. Lee, and S. D. Keilholz. Evaluation of sliding window correlation performance for characterizing dynamic functional connectivity and brain states. *NeuroImage*, 133:111–128, 6 2016. ISSN 1053-8119. doi: 10.1016/J.NEUROIMAGE.2016.02.074.
- A. Shenhav, M. M. Botvinick, and J. D. Cohen. The expected value of control: an integrative theory of anterior cingulate cortex function. *Neuron*, 79(2):217–240, 2013.
- K. S. Smith and A. M. Graybiel. A dual operator view of habitual behavior reflecting cortical and striatal dynamics. *Neuron*, 79(2):361–374, 2013.
- M. Smith and L. Fahrmeir. Spatial Bayesian variable selection with application to functional magnetic resonance imaging. *Journal of the American Statistical Association*, 102(478):417–431, 2007. ISSN 01621459. doi: 10.1198/016214506000001031.
- S. M. Smith, M. Jenkinson, M. W. Woolrich, C. F. Beckmann, T. E. J. Behrens, H. Johansen-Berg, P. R. Bannister, M. De Luca, I. Drobnjak, D. E. Flitney, R. K. Niazy, J. Saunders, J. Vickers, Y. Zhang, N. De Stefano, J. M. Brady, and P. M. Matthews. Advances in functional and structural {MR} image analysis and implementation as {FSL}. *NeuroImage*, 23:S208–S219, 1 2004.
- R. Stilla, G. Deshpande, S. LaConte, X. Hu, and K. Sathian. Posteromedial {Parietal} {Cortical} {Activity} and {Inputs} {Predict} {Tactile} {Spatial} {Acuity}. *Journal of Neuroscience*, 27(41):11091–11102, 10 2007.
- F. C. Stingo, Y. A. Chen, M. G. Tadesse, and M. Vannucci. Incorporating biological information into linear models: A Bayesian approach to the selection of pathways and genes. *Annals of Applied Statistics*, 5(3):1978–2002, 2011. ISSN 19326157. doi: 10.1214/11-AOAS463.

- F. C. Stingo, M. D. Swartz, and M. Vannucci. A Bayesian approach to identify genes and gene-level SNP aggregates in a genetic analysis of cancer data. *Statistics and Its Interface*, 8(2):137–151, 2015. ISSN 19387989. doi: 10.4310/SII.2015.v8.n2.a2. URL <http://www.intlpress.com/site/pub/pages/journals/items/sii/content/vols/0008/0002/a002/>.
- F. T. Sun, L. M. Miller, and M. D’Esposito. Measuring interregional functional connectivity using coherence and partial coherence analyses of fMRI data. *NeuroImage*, 21(2):647–658, 2004. ISSN 10538119. doi: 10.1016/j.neuroimage.2003.09.056.
- Y. Tang, D. Chen, and X. Li. Dimensionality reduction methods for brain imaging data analysis. *ACM Computing Surveys*, 54(4), 2021. ISSN 15577341. doi: 10.1145/3448302.
- P. Tewarie, M. D. Steenwijk, M. J. Brookes, B. M. J. Uitdehaag, J. J. G. Geurts, C. J. Stam, and M. M. Schoonheim. Explaining the heterogeneity of functional connectivity findings in multiple sclerosis: An empirically informed modeling study. *Human Brain Mapping*, 39(6):2541–2548, 6 2018. ISSN 1065-9471. doi: 10.1002/hbm.24020. URL <https://onlinelibrary.wiley.com/doi/10.1002/hbm.24020>.
- R. Tibshirani and P. Wang. Spatial smoothing and hot spot detection for CGH data using the fused lasso. *Biostatistics*, 9(1):18–29, 2008. doi: 10.1093/biostatistics/kxm013. URL <https://academic.oup.com/biostatistics/article/9/1/18/253824>.
- M. P. Van Den Heuvel and A. Fornito. Brain networks in schizophrenia. *Neuropsychology Review*, 24(1):32–48, 2014.
- A. Venkataraman, K. R. Van Dijk, R. L. Buckner, and P. Golland. Exploring functional connectivity in fMRI via clustering. In *2009 IEEE International Conference on Acoustics, Speech and Signal Processing*, pages 441–444. IEEE, 4 2009. ISBN 978-1-4244-2353-8. doi: 10.1109/ICASSP.2009.4959615. URL <http://ieeexplore.ieee.org/document/4959615/>.
- D. Vidaurre, S. M. Smith, and M. W. Woolrich. Brain network dynamics are hierarchically organized in time. *Proceedings of the National Academy of Sciences*, 114(48):201705120, 2017. ISSN 0027-8424. doi: 10.1073/pnas.1705120114. URL <http://www.pnas.org/lookup/doi/10.1073/pnas.1705120114>.
- A. Viterbi. Error bounds for convolutional codes and an asymptotically optimum decoding algorithm. *IEEE Transactions on Information Theory*, 13(2):260–269, 4 1967. ISSN 0018-9448. doi: 10.1109/TIT.1967.1054010. URL <http://ieeexplore.ieee.org/document/1054010/>.
- H. Wang. Bayesian graphical lasso models and efficient posterior computation. *Bayesian Analysis*, 7(4):867–886, 2012.
- S. Wang, S. F. Feng, and A. M. Bornstein. Mixing memory and desire: How memory reactivation supports deliberative decision-making. *Wiley Interdisciplinary Reviews: Cognitive Science*, 13(2):e1581, 2022.

- R. Warnick, M. Guindani, E. Erhardt, E. Allen, V. Calhoun, and M. Vannucci. A Bayesian Approach for Estimating Dynamic Functional Network Connectivity in fMRI Data. *Journal of the American Statistical Association*, 113(521):134–151, 2018.
- K. A. Wilson, G. A. James, C. D. Kilts, and K. A. Bush. Combining Physiological and Neuroimaging Measures to Predict Affect Processing Induced by Affectively Valent Image Stimuli. *Scientific Reports*, 10(1):9298, 2020.
- M. W. Woolrich, S. Jbabdi, B. Patenaude, M. Chappell, S. Makni, T. Behrens, C. Beckmann, M. Jenkinson, and S. M. Smith. Bayesian analysis of neuroimaging data in {FSL}. *NeuroImage*, 45(1):S173–S186, 3 2009.
- K. J. Worsley, J. I. Chen, J. Lerch, and A. C. Evans. Comparing functional connectivity via thresholding correlations and singular value decomposition. *Philosophical transactions of the Royal Society of London. Series B, Biological sciences*, 360(1457):913–920, 2005. ISSN 0962-8436. doi: 10.1098/RSTB.2005.1637. URL <https://pubmed.ncbi.nlm.nih.gov/16087436/>.
- G. R. Wu, W. Liao, S. Stramaglia, J. R. Ding, H. Chen, and D. Marinazzo. A blind deconvolution approach to recover effective connectivity brain networks from resting state fMRI data. *Medical Image Analysis*, 17(3):365–374, 2013. ISSN 13618415. doi: 10.1016/j.media.2013.01.003. URL <http://dx.doi.org/10.1016/j.media.2013.01.003>.
- X. Wu, R. Zurita-Milla, E. Izquierdo Verdiguier, and M. J. Kraak. Triclustering Georeferenced Time Series for Analyzing Patterns of Intra-Annual Variability in Temperature. *Annals of the American Association of Geographers*, 108(1):71–87, 1 2018. ISSN 24694460. doi: 10.1080/24694452.2017.1325725.
- Y. Xu and M. A. Lindquist. Dynamic Connectivity Detection: An algorithm for determining functional connectivity change points in fMRI data. *Frontiers in Neuroscience*, 9(JUL), 2015. ISSN 1662453X. doi: 10.3389/fnins.2015.00285.
- Q. Yu, E. B. Erhardt, J. Sui, Y. Du, H. He, D. Hjelm, M. S. Cetin, S. Rachakonda, R. L. Miller, G. Pearlson, and V. D. Calhoun. Assessing dynamic brain graphs of time-varying connectivity in fMRI data: Application to healthy controls and patients with schizophrenia. *NeuroImage*, 107:345–355, 2015a.
- Y. Yu, X. Zhou, H. Wang, X. Hu, X. Zhu, L. Xu, C. Zhang, and Z. Sun. Small-world brain network and dynamic functional distribution in patients with subcortical vascular cognitive impairment. *PLoS ONE*, 10(7):1–14, 2015b. ISSN 19326203. doi: 10.1371/journal.pone.0131893.
- G. Zhang, B. Cai, A. Zhang, J. M. Stephen, T. W. Wilson, V. D. Calhoun, and Y.-P. W. Wang. Estimating Dynamic Functional Brain Connectivity with a Sparse Hidden Markov Model. *IEEE Transactions on Medical Imaging*, 0062(c):1–1, 2019. ISSN 0278-0062. doi: 10.1109/tmi.2019.2929959.

L. Zhang, M. Guindani, F. Versace, and M. Vannucci. A spatio-temporal nonparametric Bayesian variable selection model of fMRI data for clustering correlated time courses. *NeuroImage*, 95:162–175, 2014. ISSN 10959572. doi: 10.1016/j.neuroimage.2014.03.024. URL <http://dx.doi.org/10.1016/j.neuroimage.2014.03.024>.

Appendix A

Appendix

Appendix 1

The following table reports the list of ROIs employed in the case study along with corresponding MNI stereotaxic space coordinates and their classification in *a priori* defined networks. See Section 2.3.3 for details.

Network	Abbreviation	Full Name	MNI Coordinates
Default Mode Network	PCC	Posterior Cingulate Cortex	(2, 54, 16)
	L pIPL	Left Posterior Inferior Parietal Lobule	(-46, -72, 28)
	R pIPL	Right Posterior Inferior Parietal Lobule	(50, -64, 26)
	PFC/vACC	Orbitofrontal Cortex/Ventral Anterior Cingulate Cortex	(4, 30, 26)
	dMPFC BA 8	Dorsomedial Prefrontal Cortex Brodmann Area 8	(-14, 54, 34)
	dMPFC BA 9	Dorsomedial Prefrontal Cortex Brodmann Area 9	(22, 58, 26)
	L DLPFC	Dorsolateral Prefrontal Cortex	(-50, 20, 34)
	L PHG	Parahippocampal Gyrus	(-10, -38, -2)
	L ITC	Inferolateral Temporal Cortex	(-60, -20, -18)
Fronto- Parietal Control Network	L aPFC	Left Anterior Prefrontal Cortex	(-36, 56, 10)
	R aPFC	Right Anterior Prefrontal Cortex	(34, 52, 10)
	dACC	Dorsal Anterior Cingulate Cortex	N/A
	L DLPFC	Left Dorsolateral Prefrontal Cortex	N/A
	R DLPFC	Right Dorsolateral Prefrontal Cortex	(46, 14, 42)
	L aINS	Left Anterior Insula	(-30, 20, -2)
	R aINS	Right Anterior Insula	(32, 22, -2)
	L aIPL	Left Anterior Inferior parietal Lobule	(-52, -50, 46)
	R aIPL	Right Anterior Inferior Parietal Lobule	(52, -46, 46)
Dorsal Attention Network	L MT	Left MidThalamus	(-44, -64, -2)
	R MT	Right MidThalamus	(50, -70, -4)
	L FEF	Left Frontal Eye Field	(-24, -8, 50)
	R FEF	Right Frontal Eye Field	(28, -10, 50)
	L SPL	Left Superior Parietal Lobule	(-26, -52, 56)
	R SPL	Right Superior Parietal Lobule	(24, -56, 54)
Salience Network	DAC	Dorsal Anterior Cingulate	(0, -22, 36)
	L aPFC	Left Anterior PFC	(-34, 44, 30)
	R aPFC	Right Anterior PFC	(32, 44, 30)
	L Insula	Left Insula	(-40, 2, 6)
	R Insula	Right Insula	(42, 2, 6)
	L LP	Left Lateral Parietal	(-62, -46, 30)
	R LP	Right Lateral Parietal	(62, -46, 30)
Locus Coeruleus	R LC	Rostral Locus Coeruleus	Probabilistic Atlas
	C LC	Caudal Locus Coeruleus	Probabilistic Atlas

Table A.1: The ROIs used in the case study along with *apriori* defined networks.

Method	PIBDFC			Tapered SW			BDFC (Warnick)		
State	1	2	3	1	2	3	1	2	3
Edge TPR	0.9814 (0.015)	1.0000 (0)	0.9806 (0.010)	0.9779 (0.018)	0.9676 (0.077)	0.9776 (0.015)	0.9221 (0.064)	0.9435 (0.082)	0.8326 (0.093)
Edge TNR	0.9672 (0.007)	0.9585 (0.007)	0.9351 (0.013)	0.7623 (0.074)	0.700 (0.107)	0.7034 (0.104)	0.9737 (0.039)	0.9835 (0.031)	0.9822 (0.034)
Edge F1 Score	0.9493 (0.019)	0.9585 (0.007)	0.9170 (0.020)	0.7459 (0.072)	0.6839 (0.141)	0.6888 (0.105)	0.9020 (0.090)	0.9330 (0.101)	0.8242 (0.108)
State Seq. Acc	0.9967 (0.001)	0.9880 (0.002)	0.9959 (0.001)	0.9340 (0.084)	0.7496 (0.323)	0.9538 (0.113)	0.9993 (0.001)	0.9871 (0.005)	0.9980 (0.001)
Comp Time (min)		197.57 (24.788)			0.6573 (0.085)			1015.5 (58.922)	

Table A.2: Simulation Study 1: results over 30 repetitions. I report sensitivity and specificity metrics for the estimated graphs of the corresponding states, together with the overall accuracy of the estimated state sequences. Standard deviations across the 30 simulations are showed in brackets. The proposed method maintains the best balance between sensitivity and specificity as well as latent state estimation accuracy.