

RESEARCH ARTICLE

A predictor-informed multi-subject bayesian approach for dynamic functional connectivity

Jaylen Lee¹, Sana Hussain², Ryan Warnick⁶, Marina Vannucci³, Isaac Menchaca², Aaron R. Seitz⁴, Xiaoping Hu², Megan A. K. Peters^{2,5}, Michele Guindani^{7*}

1 Department of Statistics, University of California, Irvine, Irvine, California, United States of America, **2** Department of Bioengineering, University of California, Riverside, Riverside, California, United States of America, **3** Department of Statistics, Rice University, Houston, Texas, United States of America, **4** Department of Psychology, University of California, Riverside, Riverside, California, United States of America, **5** Department of Cognitive Sciences, University of California, Irvine, Irvine, California, United States of America, **6** Microsoft Security Research, Microsoft, Redmond, Washington, United States of America, **7** Department of Biostatistics, University of California, Los Angeles, Los Angeles, California, United States of America

* mguindani@ucla.edu

OPEN ACCESS

Citation: Lee J, Hussain S, Warnick R, Vannucci M, Menchaca I, Seitz AR, et al. (2024) A predictor-informed multi-subject bayesian approach for dynamic functional connectivity. *PLoS ONE* 19(5): e0298651. <https://doi.org/10.1371/journal.pone.0298651>

Editor: Debo Cheng, University of South Australia, AUSTRALIA

Received: January 9, 2023

Accepted: January 30, 2024

Published: May 16, 2024

Copyright: © 2024 Lee et al. This is an open access article distributed under the terms of the [Creative Commons Attribution License](https://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Data Availability Statement: The fMRI dataset and pupillometry data are available on Dryad at <https://doi.org/10.7280/D1HQ3B>. The code for the proposed PIBDFC model can be downloaded from GitHub repository at <https://github.com/jayerule/PIBDFC>.

Funding: This study was supported by the National Science Foundation Graduate Research Fellowship in the form of a grant to JL [DGE-1839285]. MG was partially supported by the Karen Toffler Charitable Trust in the form of an award.

Abstract

Dynamic functional connectivity investigates how the interactions among brain regions vary over the course of an fMRI experiment. Such transitions between different individual connectivity states can be modulated by changes in underlying physiological mechanisms that drive functional network dynamics, e.g., changes in attention or cognitive effort. In this paper, we develop a multi-subject Bayesian framework where the estimation of dynamic functional networks is informed by time-varying exogenous physiological covariates that are simultaneously recorded in each subject during the fMRI experiment. More specifically, we consider a dynamic Gaussian graphical model approach where a non-homogeneous hidden Markov model is employed to classify the fMRI time series into latent neurological states. We assume the state-transition probabilities to vary over time and across subjects as a function of the underlying covariates, allowing for the estimation of recurrent connectivity patterns and the sharing of networks among the subjects. We further assume sparsity in the network structures via shrinkage priors, and achieve edge selection in the estimated graph structures by introducing a multi-comparison procedure for shrinkage-based inferences with Bayesian false discovery rate control. We evaluate the performances of our method vs alternative approaches on synthetic data. We apply our modeling framework on a resting-state experiment where fMRI data have been collected concurrently with pupillometry measurements, as a proxy of cognitive processing, and assess the heterogeneity of the effects of changes in pupil dilation on the subjects' propensity to change connectivity states. The heterogeneity of state occupancy across subjects provides an understanding of the relationship between increased pupil dilation and transitions toward different cognitive states.

Competing interests: The authors have declared that no competing interests exist.

1 Introduction

Functional connectivity (FC) has emerged as one of the most active research areas in functional magnetic resonance imaging (fMRI). The purpose of FC studies is to characterize the undirected statistical dependencies between brain regions and thus gain a greater understanding of brain functioning [1, 2]. Simple approaches to studying FC rely on readily available measures of temporal correlation, such as the partial correlations between two brain regions after conditioning upon all other regions [3, 4]. Such metrics assume that interactions between brain regions are constant in space and time throughout the fMRI session [5]. Rather, neuroscientists have become increasingly aware that functional connectivity is dynamic and fluctuating, i.e. non-stationary, and that studying the dynamics of FC is crucial for improving our understanding of human brain function [2, 6, 7]. The term “chronnectome” has been introduced to describe the growing focus on identifying time-varying, but reoccurring, patterns of coupling among brain regions [8].

Recent studies have highlighted how subjects are more likely to experience particular connectivity states when some underlying physiological conditions are present. For example, [9] have investigated the association between heart rate variations and FC. Similarly, in a sleep fMRI study, [10] have shown that transitions between connectivity states slow as subjects fall into deeper sleep stages. As a further example, [11] 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 functional connectivity states are modulated by internal or external conditions measured during the course of an experiment. In the experimental study we consider in this manuscript, we have available fMRI data collected together with pupillometry measurements. Pupil dilation has become increasingly popular in cognitive psychology to measure cognitive processing and resource allocation. It is believed that the changes in pupil diameter reflect brain state fluctuations driven by neuromodulatory systems [12]. For example, the pupil dilates more under conditions of higher attention [13]. Thus, pupil dilation measurements can be seen as an index of effort exertion, task demand, or difficulty in an fMRI experiment [14]. Thus, it is of interest to understand how pupil dilation is associated with an increased probability of particular connectivity states experienced by a subject during an experiment [15].

Many of the commonly used approaches for studying dynamic connectivity rely on multi-step inferences. For example, in [8] the fMRI time courses are first segmented by a sequence of sliding windows, and then precision matrices are estimated in each window. Finally, *k*-means clustering methods are used to identify re-occurring patterns of FC states. Post-hoc analyses may be employed to assess the association of the estimated dynamic connectivity states with other available measurements, like pupil dilation measurements [16]. More recent developments replace the *k*-means clustering process with deep neural network based clusters or latent factor models to estimate FC states. [17–19]. However, the arbitrary choice of the window length and the offset may lead to spurious dynamic profiles and poor estimates of correlations for each brain state [20, 21]. Alternative approaches were proposed by [22–24], who developed change point detection methods to recursively partition the fMRI time series into stable contiguous segments where networks of partial correlations are estimated by employing the graphical lasso of [25]. A more recent variation use a vine copula model to estimate the relationships between ROI [26]. These methods do not require pre-specifying the segment lengths and can detect the temporal persistence of the functional networks. However, they do not account for the possibility of states being revisited and hence do not conform to the idea that

the chronnectome exhibits recurrent patterns of dynamic coupling between brain regions of interest (ROIs).

Other model-based approaches to dynamic connectivity consider the set of ROIs as the nodes (or vertices) of an underlying graph and employ homogeneous hidden Markov models (HMMs) to detect state transitions and infer a discrete set of latent connectivity states over time. [27] develop a Bayesian HMM to model dynamic FC as the transition between state-specific graphs, each graph being related to others via an underlying super-graph. [28] directly model the evolving correlation matrix using an HMM. [29] use product HMMs to describe the evolution of the sliding-windows correlations and capture temporal dependencies across resting-state networks. [30] used a Bayesian HMM to estimate the dynamic structure of graph theoretical measures of whole-brain FC. Also, HMMs have been employed in time-varying vector autoregressive (VAR) modeling frameworks for whole-brain resting state connectivity, where the VAR coefficients and the innovation covariance matrix are allowed to change with the latent states [6, 31, 32]. However, these implementations of hidden Markov models typically assume that the probabilistic model underlying the state transitions is constant throughout an experiment. Crucially, such a homogeneity assumption does not allow to assess the modulatory effect of time-varying physiological factors on the transitions, e.g. how changes in vigilance measured via pupil dilation can impact state transitions [7].

In this paper, we develop a multi-subject Bayesian framework where the estimation of dynamic functional networks is informed by time-varying exogenous physiological covariates that are simultaneously recorded in each subject during the fMRI experiment. More specifically, we introduce a multi-subject non-homogeneous HMM modeling framework where the transition probabilities between states are shared between subjects and vary over time as a function of the covariates. Our setting allows for the estimation of unique connectivity state transitions for each subject. It also permits group-based inferences, via recurring connectivity patterns and sharing of networks among the subjects. With respect to the 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. [33] have recently proposed a two-step multi-subject fused-lasso approach for detecting change points in functional connectivity. Differently from 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 can therefore 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. This differs from the recently developed dynamic mixture model by [34] where the network edges are estimated from the task information, as opposed to the transitions like in our proposed approach. Our modeling approach further assumes sparsity in the network structures, by assuming a shrinkage prior on the connectivity networks. Additionally, we propose a strategy for edge selection that combines the posterior shrinkage-informed thresholding approach of [35] with the Bayesian False Discovery Rate controlling method of [36].

We apply our modeling framework to 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, [37] 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. Therefore, pupil dilation has been used as a proxy for a metric of a person's willingness to exert more effort and involve a greater mental effort to complete a task. 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 [38]. In our application, we demonstrate how the model can recover expected change points in dynamic FC 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 we describe our proposed method and edge selection procedure. We also give a brief synopsis of our Markov Chain Monte Carlo (MCMC) approach to posterior inference. In section 3 we showcase our model performance on simulated data and provide comparisons to the sliding window and homogenous HMM approaches. Lastly, in Section 4, we apply our model to the LC handgrip data with accompanying results and analysis. Section 5 concludes the paper with a discussion.

2 Methods

In this section, we describe our 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—simultaneously—estimates the predictor effects on those dynamics. We achieve this by constructing a non-homogeneous Hidden Markov Model (HMM) where the transition probabilities are functions of time-varying covariates.

2.1 An HMM model for dynamic functional 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 interest (ROIs), $t = 1, \dots, T$ on subject $i = 1, \dots, N$. We 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 [39] to capture activation patterns over time, as done for example in [27]. Here, however, since we are interested in capturing connectivity patterns through the modeling of the time-varying precision matrix, we 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 FC, 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 [40].

We propose to model the dynamics of FC using an HMM framework with S latent states characterizing FC and the brain transitions during the fMRI experiment. Our 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's functional organization. However, we 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, we assume the data follow a Gaussian graphical model at time t of the type

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

with subject-level precision matrices which, at each time, are characterized by the values of one among S precision matrices, identifying which state is active at that time. Model (1) therefore implies connectivity networks that vary by subjects and by state.

2.2 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 [41]. For example, our 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.

Here, we propose to model the dynamics of FC by developing a non-homogeneous HMM framework where estimation is informed by subject-level time-varying exogenous physiological covariates, e.g. physiological factors like the pupillary data in our motivating application. More in detail, we assume that switches between states are regulated by transition probabilities that vary over time and across subjects as a function of B time-varying subject-level covariates as

$$Q_{rst}^i = P(s_{t+1} = s | s_t = r) = \frac{\exp(\zeta_{rs}^i + \mathbf{x}_t^{i^T} \boldsymbol{\rho}_s^i)}{\sum_{l=1}^S \exp(\zeta_{rl}^i + \mathbf{x}_t^{i^T} \boldsymbol{\rho}_l^i)}, \quad r, s = 1, \dots, S, \tag{2}$$

where \mathbf{x}_t^i denotes a $B \times 1$ vector of covariate values for subject i at time t , and $\boldsymbol{\rho}_s^i = (\rho_{s1}^i, \dots, \rho_{sB}^i)$ 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 ζ_{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, we define a state as reference. Without loss of generality, we set $s = 1$ as the reference state, and also set the coefficients $\rho_{1b}^i, b = 1, \dots, B$, and $\zeta_{1\cdot}^i, i = 1, \dots, N$ equal to zero. Thus, the state transition coefficients are interpreted with respect to the reference state, and we can re-express (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,

$$\log\left(\frac{Q_{rst}^i}{Q_{r1t}^i}\right) = \zeta_{rs}^i + \mathbf{x}_t^{i^T} \boldsymbol{\rho}_s^i, \quad r, s = 1, \dots, S. \tag{3}$$

In this formulation, the transition coefficients $\exp(\rho_{sb}^i), b = 1, \dots, B$, are more naturally interpreted as 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 as constant. Similarly, the coefficient $\exp(\zeta_{rs}^i)$ is interpreted as the expected odds of transitioning from state r to s compared to transitioning from state r to 1, when the time-varying covariates, \mathbf{x}_t^i , are 0 or at a baseline/average value.

We assume independent Gaussian priors for the transition parameters ρ and ξ . We 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, we assume that the individual coefficients ξ_{rs}^i and ρ_{sb}^i , $b = 1, \dots, B$, vary around population-level means, Z_{rs} and η_{sb} , as follows:

$$\begin{aligned}
 s_{t+1}^i | s_t^i &= r \sim \text{Multi}(Q_{r,s,t}^i) \quad t = 1, \dots, T, \\
 \xi_{rs}^i &\sim N(Z_{rs}, \sigma_\xi), \\
 \rho_{sb}^i &\sim N(\eta_{sb}, \sigma_\rho), \\
 Z_{rs} &\sim N(z_{rs}^0, \sigma_z), \\
 \eta_{sb} &\sim N(0, \sigma_\eta),
 \end{aligned}
 \tag{4}$$

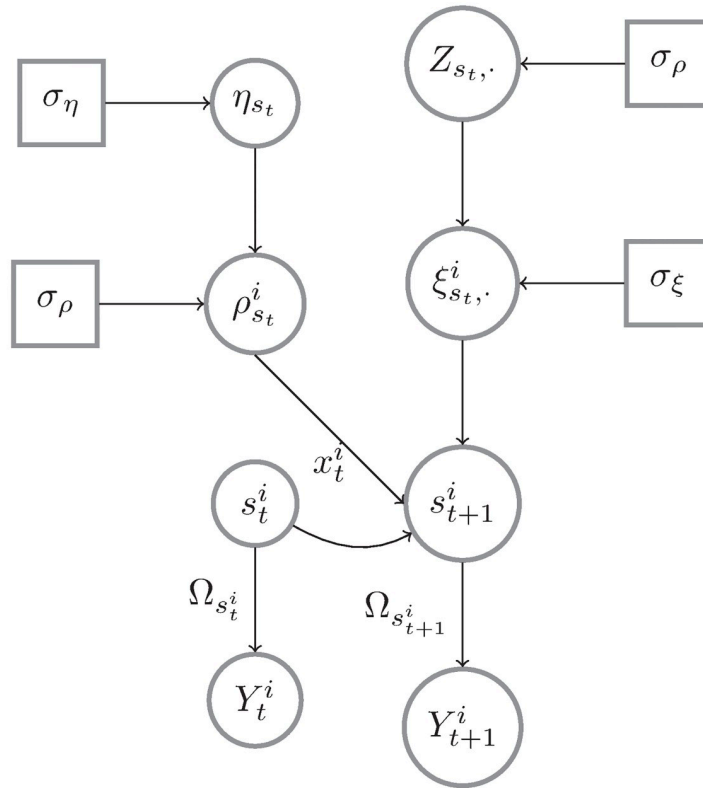
where $Q_{r,s,t}^i = (Q_{r,1,t}^i, \dots, Q_{r,S,t}^i)^T$, and $r, s = 1, \dots, S, b = 1, \dots, B$. By allowing each subject to have their own transition parameters the model allows for unique subject-level transition behavior while also capturing population-level estimates through the group level parameters. The interpretation of the group level parameters, η and Z , is similar to their single subject counterparts. The prior means z_{rs}^0 are usually set to 0 except for 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, we recommend setting the variance parameters of the priors for ξ, ρ, Z_{rs} and η_{sb} to some small positive value on the order of 0.1.

2.3 Modeling sparsity through a graphical horseshoe prior

Functional networks are thought to exhibit the so-called small world behavior, indicating a high degree of clustering and high efficiency in the estimated networks [42, 43]. 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 set to exact zeros. [27] achieve such a selection by using a G-Wishart prior to sample positive definite matrices according to estimated adjacency matrices that correspond to the FC 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 does not scale well to relatively large dimensions of the networks. Here, instead, we take a shrinkage-based approach and model the off-diagonal entries of the state-specific precision matrices $\Omega_s, s = 1, \dots$, in (1) by employing a graphical horseshoe prior [44]. Thus, we set

$$p(\Omega_s | \tau, \Lambda) \propto \prod_{j < k} N(\omega_{jks} | \lambda_{jk}^2 \tau^2) \prod_{j < k} C_+(\lambda_{jk} | 0, 1) I(\Omega_s \in S_R), \quad s = 1, \dots, S \tag{5}$$

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 (5), we further assume a non-informative flat prior for the diagonal elements, i.e. $\omega_{jjt} \propto 1$. 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 $\omega_{jkt}, j = 1, \dots, k - 1, k = 1, \dots, R$. 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



$$s_t^i \in \{1, \dots, S\}, t = 1, \dots, T, i = 1, \dots, N$$

Fig 1. 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 FC 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 $Q_{s_t^i}$. This entry is modeled according to Eq 3.

<https://doi.org/10.1371/journal.pone.0298651.g001>

maintains the magnitude of non-zero entries and thus reduces biases. Following [44], we assume a half-Cauchy prior on $\tau, \tau \sim C_+(\cdot; 0, \tau_0)$, 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 [35], where an edge j,k is selected if $E\left(\frac{1}{1+\lambda_{jk}\tau}\right) < 0.5$. We demonstrate such a process in S1 Fig in the S1 File. We find that a $\tau_0 = 1$ gives an expected edge density of approximately 50% while having the largest spread. Fig 1 provides a graphical representation of the proposed predictor-informed Bayesian dynamic FC model (PIBDFC).

2.4 Posterior inference

The posterior distribution for the parameters in the proposed model is not available in closed form. Hence, we revert to Markov Chain Monte Carlo (MCMC) techniques for posterior inferences. In particular, we follow [45] and employ Polya Gamma auxiliary variables [46] to sample the state transition parameters. Based on the sampled $Q_{s_t^i}$, we can construct a sequence of transition matrices based on Eq (3). After normalizing each row $Q_{s_t^i}^i$ so that it

sums to 1, we use a stochastic forward-backward algorithm to sample the state sequence [47]. Then, conditioned upon the estimated state sequence, it is possible to sample the precision matrix parameters using the blocked Gibbs algorithm presented in [44]. Other parameters in the hierarchical model for the states' transitions (4) are sampled via simple Gibbs steps. By iterating through the steps above, we obtain samples from the posterior. We provide a brief summary below:

1. **Sample $Q_{rs}^i, \zeta_{rs}^i, \rho^i$:** We can rewrite the likelihood for ζ_{rs}^i according to [48] to be in the form of Eq 6.

$$L(\zeta_{rs}^i) \propto \prod_{t:s_{t-1}^i=r} \frac{\exp(\zeta_{rs}^i - c_{rst}^i)^{I(s_t^i=s)}}{1 + \exp(\zeta_{rs}^i - c_{rst}^i)} \tag{6}$$

where $c_{rst}^i = \log \sum_{m \neq s} \exp(\zeta_{rm}^i + \mathbf{x}_t^i \rho_m^i - \mathbf{x}_t^i \rho_s^i)$. Using the Polya-Gamma augmented logistic regression technique of [46], we get the posterior of ζ_{rs}^i to be conditionally Gaussian.

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

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 $PG(1, \zeta_{rs}^i - c_{rst}^i)$. We use a similar strategy to update ρ_{rb}^i , the logistic component for subject i for state r and covariate b , achieving the posterior:

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

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

2. **Sample s_t^i :** We sample the sequence of states by adapting the stochastic forward-backward algorithm presented by [47].
3. **Sample the matrices $\Omega_s^i, s = 1, \dots, S$:** The conditional posterior for Ω_s is as follows:

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

For MCMC inference purposes, [44] adopt auxiliary variables v_λ and ξ_r , in order to modify the Gibbs sampling procedure presented by [49]. This procedure is performed for a column-wise update in a fashion similar to [50]. For each state, we 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 Z_{rs}, η_b :** These conditional posteriors follow the typical normal-normal update:

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

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

2.5 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, [44] 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 does not provide a multiplicity control for false discoveries.

We 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 we consider the posterior estimates of Ω_s for each state $s = 1, \dots, S$, separately, in the following we 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 [51], for real numbers $c_1, c_2 > 0$, we can then determine the optimal set of decisions $\boldsymbol{\delta} = \{\delta_{12}, \delta_{13}, \dots, \delta_{R-1 R}\}$ by minimizing the following loss function:

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

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,

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

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,

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

[44] 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 [52]. 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 [44], and related discussions in [53]. The result extends trivially to the folded normal distribution characterizing $|\omega_{jk}|$. Note that $\kappa_{jk} \in (0, 1)$, and that larger values of $E(\kappa_{jk})$ indicate stronger shrinkage of the posterior estimates toward zero.

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, [35] 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 [54] 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 [55], who consider inference on the partial correlation matrix derived from Ω .

3 Simulation study

In this Section, we present three sets of simulated datasets that aim at measuring the performance of our model with respect to the detection of non-zero connectivities and the estimation of the latent connectivity states over time. More specifically, in the first two simulation studies, we 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 [56], and the Bayesian Dynamic Functional Connectivity (BDFC) model developed by [27]. The Tapered SW represents a standard approach in the FC literature, whereas BDFC uses a homogeneous HMM to model latent connectivity state dynamics. The 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 were developed for single-subject inference. We 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, we simulated 5,000 posterior samples after 5,000 burn-in draws. When fitting PIBDFC, we set the hyperparameters $\tau_0 = 1$, $\sigma_\xi = \sigma_\rho = \sigma_z = \sigma_\eta = 0.1$, following the motivations of Section 2.

We assess the performance of our model on 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 regions j and k . Following the discussion in Section 2, let δ_{jk} indicate the decision after the model fit. Then we 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, we 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 $\sum \{I(s_t^i = s)I(\hat{s}_t^i = s)\} / \sum I(s_t^i = s)$.

3.1 Simulation Study 1

In our first study, we investigate the performance of our model in an ideal setting where the data generation process matches the model closely. We set $T = 300$ time points, $R = 16$ ROIs, $N = 30$ subjects, and $S = 3$ connectivity states. In this setting, we 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 Fig 2. In order to study the effect of the predictor information on the estimation of the transition probabilities and the FC dynamics, we 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, we set the transition probabilities for the latent state trajectories as follows

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

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. We then simulate different state sequences for each subject using Eq (3), and replicated the process over 30 independent simulated data sets. In order to assess the performance of the methods for different levels of signal strength, we repeated the simulation experiment using different precision matrices Ω_s , $s = 1, 2, 3$ of the same structure of the top row of Fig 2 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 random non-zero entries. This output matrix is then normalized to a partial correlation matrix. Thus, we obtained a total of six sets of precision matrices to learn the structure of. We show the aggregated results in Fig 3. 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 regard to true positive rate and F1 score, across all levels of partial correlations. The BDFC appears as the most conservative, as highlighted by the large true negative rates, but low true positive rates. Based on the results above, the PIBDFC displays the best balance of finding true non-zero partial correlations while controlling for false positives.

In the following, we illustrate the inferential analyses enabled by the proposed PIBDFC approach by showcasing a single replicate. In Fig 2 (bottom row) we 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. The model is also able to recover the most likely state transition sequence for each subject, as determined by the *maximum a posteriori* state estimate at each time point. See Fig 4. It is also important to assess the ability of the

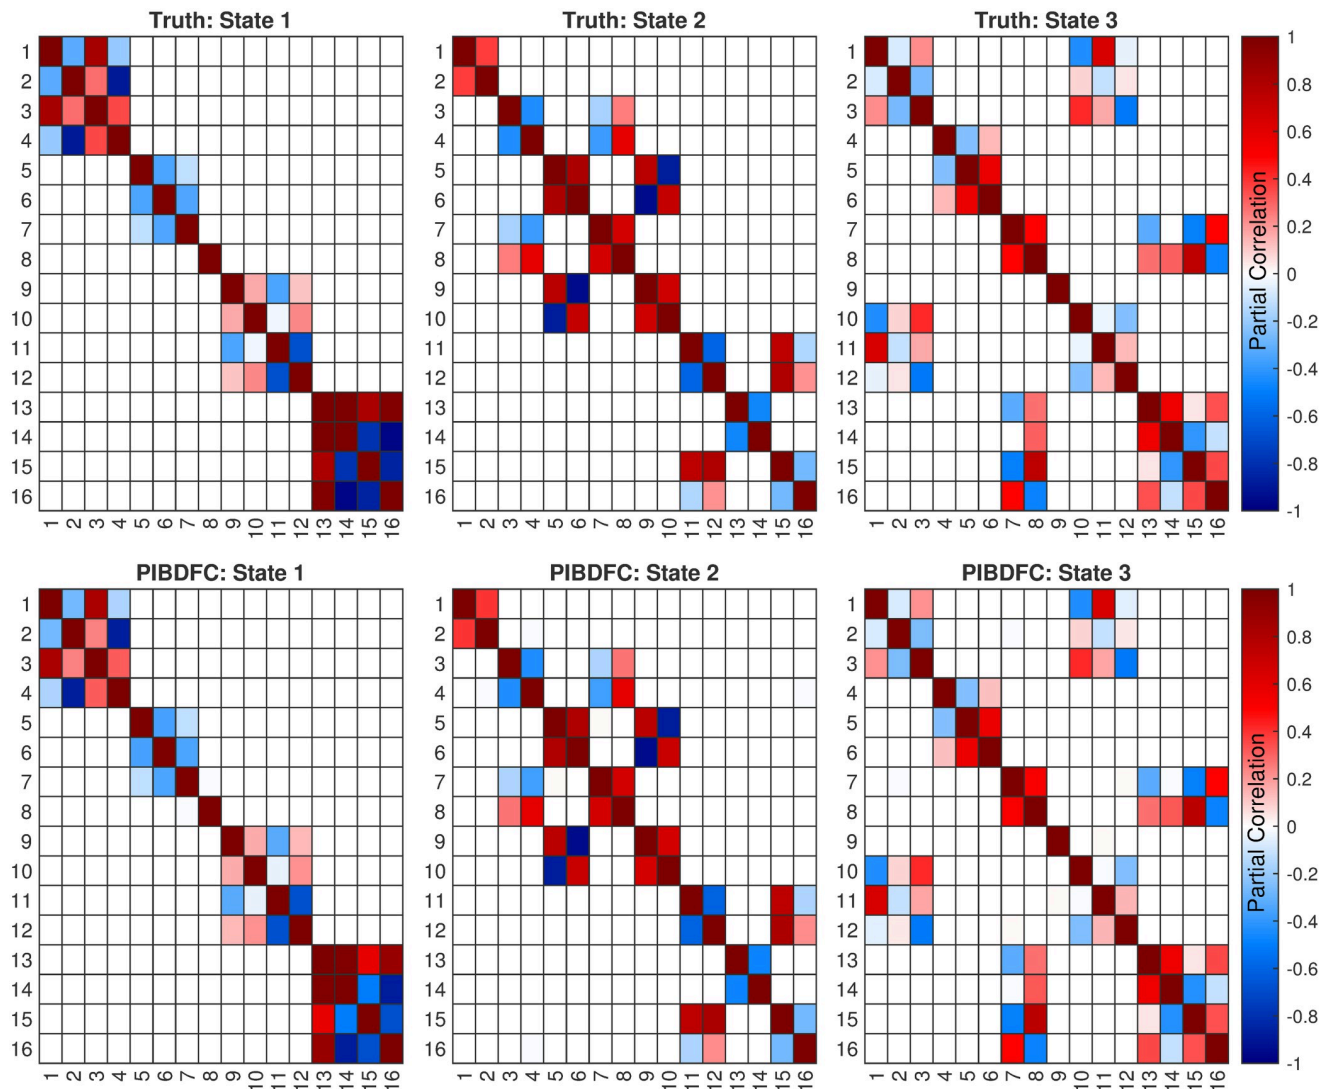


Fig 2. Example estimation of state related partial correlation matrices on simulated data. *Top:* The true partial correlation matrices for each state responsible for generating the simulation data in the Simulation Study 1. *Bottom:* The estimated partial correlation matrix from the proposed PIBDFC 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 and 3 for details.

<https://doi.org/10.1371/journal.pone.0298651.g002>

method to identify true change points in the connectivity structure. Fig 5 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 1. PIBDFC is also able to draw as many posterior draws in a third of the computation time.

3.2 Simulation Study 2

In this second simulation study, we measure the performance of our approach with synthetic data that are similar to real fMRI data. More specifically, we use the Matlab simulation toolbox SimTB of [57] and follow the simulation approach of [27]. The SimTB

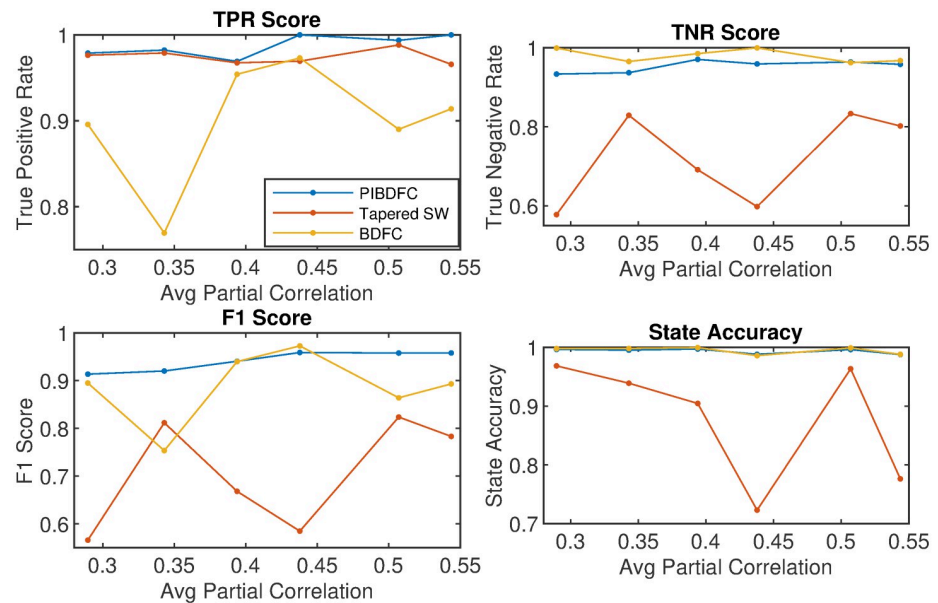


Fig 3. Edge detection performance by PIBDFC, BDFC, and Tapered SW. True Positive Rate, True Negative Rate, F1 Score, and state accuracy metrics for the PIBDFC, BDFC, and Tapered SW approaches over different settings of the correlation structure. Along each horizontal axis is the average strength of the non-zero partial correlations for each state, corresponding to different levels of signal strength.

<https://doi.org/10.1371/journal.pone.0298651.g003>

toolbox implements a canonical hemodynamic response function [58], 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. FC 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. We then simulate the state sequence over $T = 150$ time points with x_t being 0 for the first 75 time points and 1 for the last 75 among all subjects. Similar to Simulation Study 1, we use the exact same Q_t among all subjects. We repeat this process for $N = 30$ subjects over 30 simulation replicates.

In Table 2 we show the results to the application on the SimTB data. PIBDFC does a good job at detecting the connectivities between the simulated regions, despite a misspecified likelihood. The performance in both graph and state estimation appears to decline slightly in comparison to the Simulation 1 setting, which is expected. The Tapered SW approach suffers from low specificity. Compared to the standard HMM of BDFC, the proposed PIBDFC performs slightly better at detecting changes in state transitions, thus improving graph estimation performance as a result. This is likely due to the distortion introduced in 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 the approach of [27], despite allowing for individual differences in state dynamics among the 30 subjects.

3.3 Simulation Study 3

In this simulation setting, we compare the performances of our model and the Connectivity Change Point Detection (CCPD) model of [33] on edge- and change-point detection. Contrary to our model, the CCPD model employs a two-stage approach for estimating dynamic FC. In

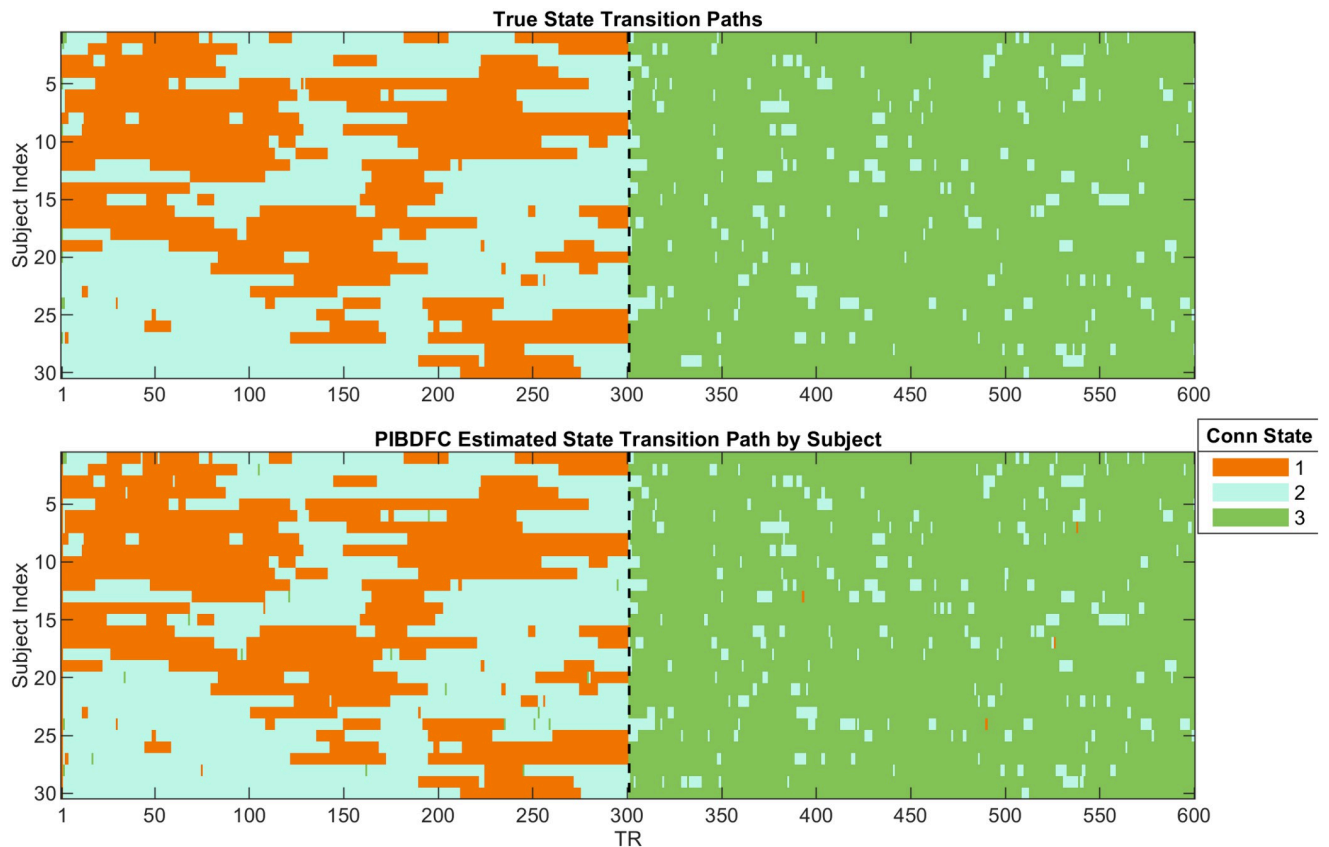


Fig 4. Example estimation of state transitions on simulated data. *Top:* The true state transition path for each subject (vertical axis) across each time point (horizontal axis). The color in each cell identifies which precision matrix in Fig 2 generated the simulated the data for each subject-time point pairs. *Bottom:* The *maximum a posteriori* estimated state trajectories from PIBDFC.

<https://doi.org/10.1371/journal.pone.0298651.g004>

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 we simulate data under the CCPD assumption of common change points. More specifically, we 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. We 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, does not infer common change points across individuals; therefore, we report the average number of change points across all subjects.

In Table 3, we 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 overestimate the number of change points on average, as it sometimes estimates very sudden changes of state for a brief collection of time points in some subjects. In contrast, in simulation studies 1 and 2, the change points are generated from a process that truly follows a hidden Markov model, leading to more accurate

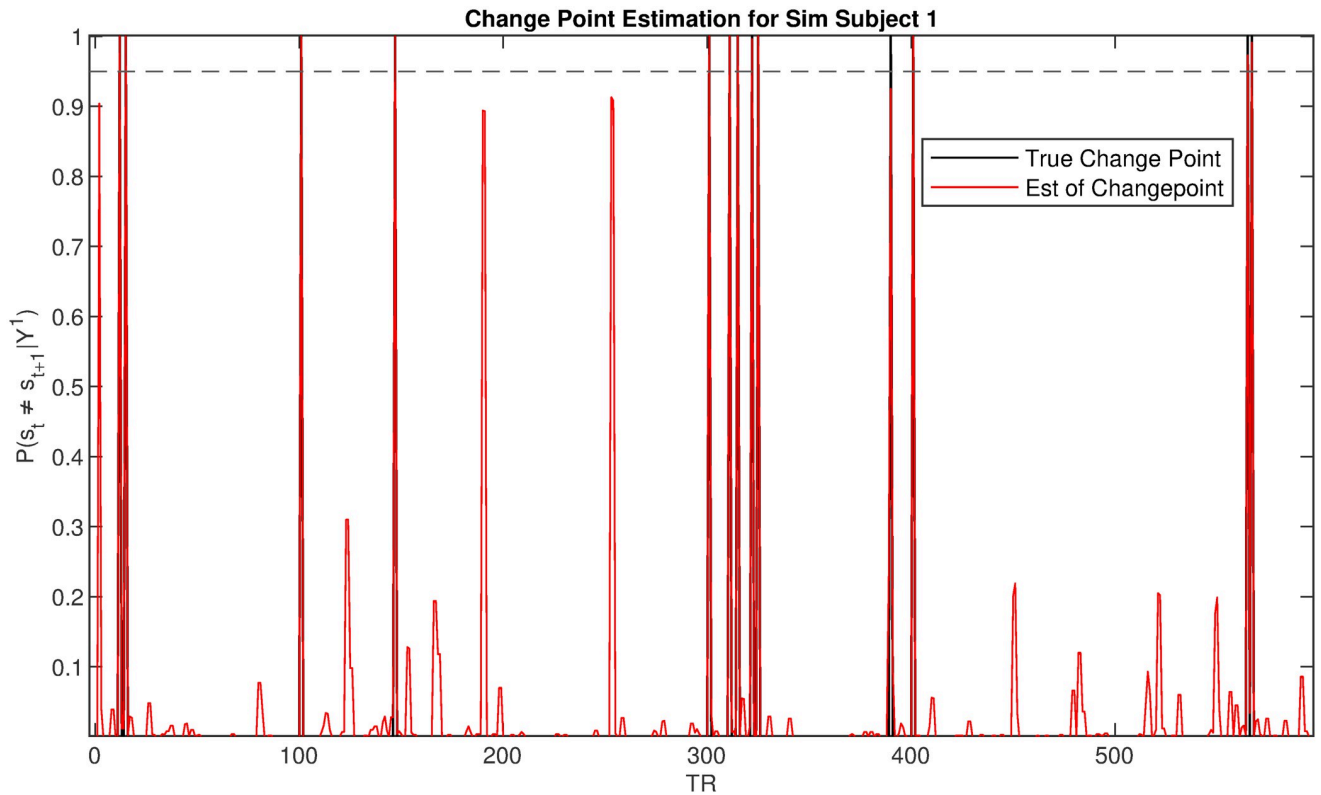


Fig 5. Example estimation of state changepoints on simulated data. Estimation of the connectivity change points in a representative subject. The horizontal axis indicates the time points while the vertical axis reports the posterior probability $P(s_t \neq s_{t-1} | Y_{1:T}^1)$. The posterior probabilities of a change point are in red, whereas the black spikes represent the true change points for the subject. We also display a horizontal dotted line at 0.95 to reflect the informal rule of declaring a change-point if $P(s_t \neq s_{t-1} | Y_{1:T}^1) > 0.95$.

<https://doi.org/10.1371/journal.pone.0298651.g005>

Table 1. Simulation Study 1: Results over 30 repetitions. We 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.

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

<https://doi.org/10.1371/journal.pone.0298651.t001>

Table 2. Simulation Study 2: Results over 30 repetitions. We 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 shown in brackets. The proposed method maintains the best balance between sensitivity and specificity as well as latent state estimation accuracy.

Metric	Method	State 1	State 2	State 3
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)
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)
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)
Time (min)	PIBDFC		161.23 (29.493)	
	Tapered SW		1.9241 (0.31)	
	BDFC		500.57 (18.11)	

<https://doi.org/10.1371/journal.pone.0298651.t002>

estimates. 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.

4 Case study

We apply the proposed PIBDFC model to the motivating dataset. In our application, we demonstrate how the model can recover expected change points in dynamic FC states, as those states align quite well with the experimental events regulated by the behavioral task. We are also able to estimate the effect of pupil dilation on the subjects' propensity to change states.

4.1 Experimental design and data collection

In this experiment, subjects performed a handgrip task adapted from [59]. Thirty-one participants (18 females, mean age 25 years \pm 4 years) enrolled in this study at the University of

Table 3. Simulation Study 3: Results over 30 repetitions. We show the entry-wise true positive and true negative rates for the estimated graphs for the corresponding states. We also show 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.

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

<https://doi.org/10.1371/journal.pone.0298651.t003>

California, Riverside Center for Advanced Neuroimaging, but one was excluded due to a history of attention deficit hyperactive disorder resulting in a total of $N = 30$ subjects. 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 [59]. 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 fMRI data underwent a standard preprocessing pipeline using 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 [60, 61]. Finally, all data were transformed from the individual subject space to the Montreal Neurological Institute (MNI) standard space using standard procedure in FSL [60, 61].

Pupillometry data were collected during the scans, using a sampling rate of 2kHz, preprocessed using the ET-remove artifacts toolbox (github.com/EmotionCognitionLab/ET-remove-artifacts), and downsampled to match the temporal resolution of the fMRI data [59]. 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 the acquisition of pupil dilation measurements, resulting in $N = 27$ for all subsequent analyses.

Since we 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 [62]. The fronto-parietal control network (FPCN) was chosen because it is linked to DAN and regulates perceptual attention. Salience network (SN) was selected because it determines which stimuli in our environment are most deserving of attention [59, 63]. Talariach coordinates for regions of interest (ROIs) within DMN, FPCN, and DAN were taken from [64] and converted to MNI coordinates while SN MNI coordinates were taken directly from [64–67]. 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 [68]. Blood oxygen level-dependent (BOLD) signal from each voxel within an ROI were extracted and averaged to represent the overall signal for an ROI. We 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 isotopic sphere [69, 70]. See the [S1 File](#) for a list of the ROIs and corresponding MNI stereotaxic space coordinates and networks.

4.2 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 $t = 1, \dots, 1050$. We also included concurrently recorded pupillometry data as a proxy for quantifying the effect of LC engagement on the dynamics of FC [71].

We fit our model with different number of total states, i.e., $S = 3, 4, 5, 6$. However, when assuming more than 3 states, the fit simply degenerated to 3 states in the posterior inference, with no observations assigned to additional states. This result indicates no posterior support for models with $S > 3$. Thus, here we present the model specification for 3 states with the following settings for the hyperparameters in (2). We set the group level baseline relative transition prior means $z_r^0 = 2$ for $r = 2, 3$ while all other elements of z_r^0 are set to 0. We also set the prior spread of the baseline transitions and pupillary effects $\sigma_z, \sigma_\eta = 0.05$. This combination of settings is used to encourage self-transitions, as they correspond to preferring smoother state sequences a priori among all subjects. We 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 (see S1 Fig in the [S1 File](#)).

4.3 Results and discussion

[Fig 6](#) plots the estimated connectivity networks for each of the three states. Nodes represent ROIs and edges identify the estimated non-zero partial correlations between pairs of nodes.

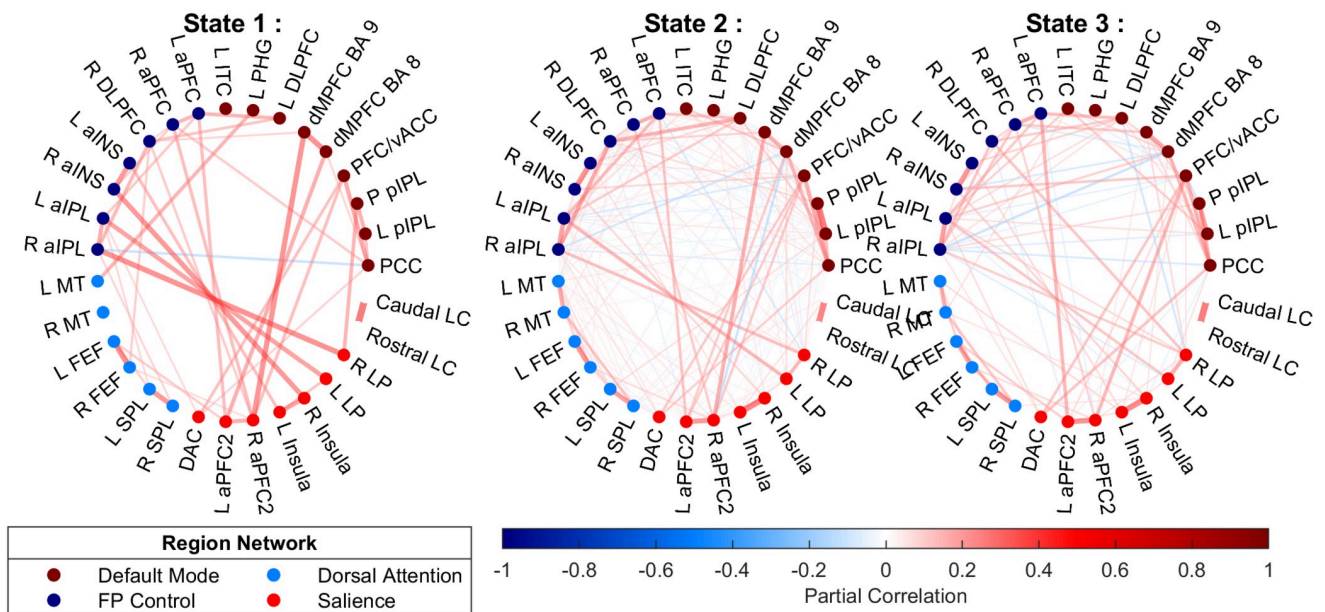


Fig 6. Real Data Analysis: The estimated connectivity networks for the ROIs. Nodes represent ROIs and the edges denote the partial correlations between the connected nodes. The edge colors correspond to the directionality of the partial correlations and the width corresponds to the magnitude. Node colors identify clusters of regions into a priori defined networks. See Section 4 and S1 Table in the [S1 File](#).

<https://doi.org/10.1371/journal.pone.0298651.g006>

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 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), Salience Network (SN), and Locus Coeruleus (LC). Fig 7 shows the *maximum a posteriori* (MAP) estimated state sequences from our model for all 27 subjects. The subjects' rows are ordered by the similarity of the estimated state trajectories as captured by a hierarchical clustering using Euclidean distance.

By inspecting Fig 6, it is apparent that state 1 shows relatively sparser connectivity than the other two states. In state 1, we 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, Fig 7 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, we 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-

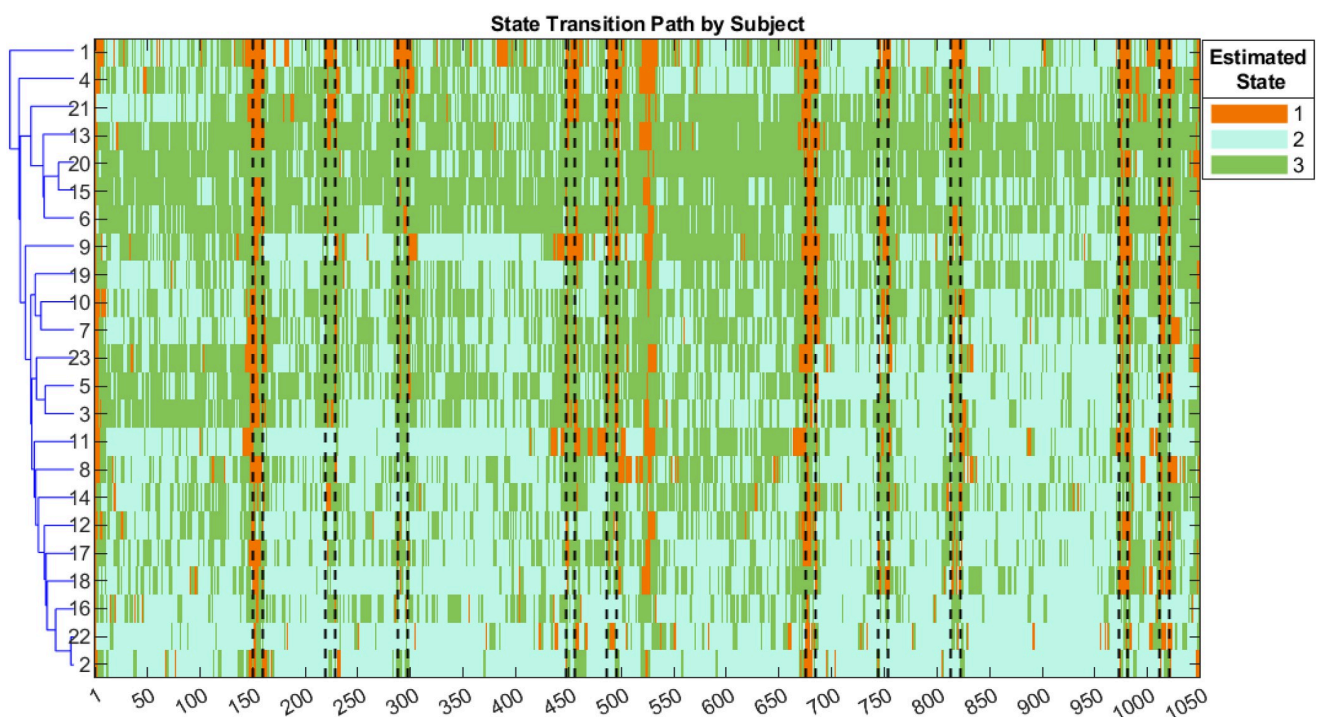


Fig 7. Real Data Analysis: Estimated states' transition path for each subject. The horizontal axis indicates the TR with vertical dotted lines indicating portions where the subject raises their arm. Subject sequences are aligned so that the first 525 time points show sequences from the sham condition and the time points 526–1050 show sequences from the active condition. The vertical axis displays the subject indices, ordered by similarity in state trajectory according to a hierarchical clustering (based on the Euclidean distance) of their MAP transition behavior.

<https://doi.org/10.1371/journal.pone.0298651.g007>

connectivity—between DMN and FPCN. Interestingly, these parallel some of the strongest connectivities from state 1. The relative occupancy in state 2 appears higher in the active condition (Fig 7, right half) than the sham condition (Fig 7, left half), suggesting subjects tended to occupy this relatively strong, broadly-connected state more often when periodically engaging in actively squeezing the 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. We 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, we also see many more connections with SN from DAN (light blue) than in either of the other two states. We 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 Fig 7): this state traded off with state 2 for relative percentage occupancy across the subjects.

Finally, a unique feature of our model is that it allows the investigation of how pupillary dilation modulates state transitions. Fig 8 provides the posterior distribution of the group (e^η , left) and individual (e^ρ) effects of pupil dilation on state dynamics. We start by assessing the relationship 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

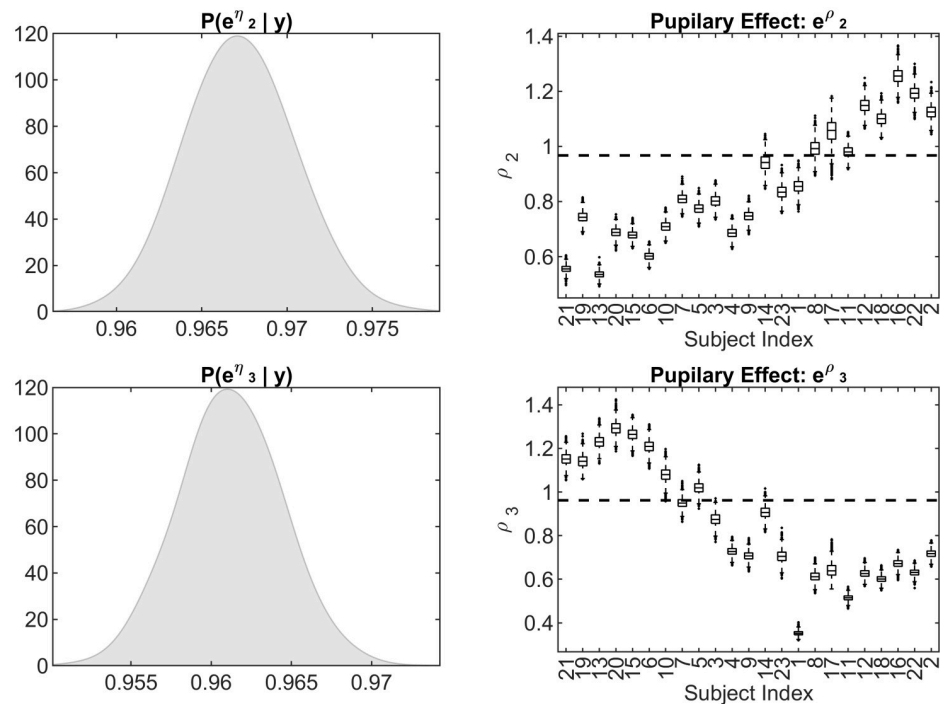


Fig 8. Estimated effects of pupil dilation on state transition probability. Real Data Analysis: The posterior distribution of the group effect of pupillary dilation e^η (left), and individual effects of pupillary dilation e^ρ . Rows indicate the propensity for transitioning into states 2 and 3 respectively. For the individual effects, subjects are identically ordered as in Fig 7. The horizontal dotted line is the posterior mean for the group-level effects, $\eta_2 = 0.687$ and $\eta_3 = 0.651$ respectively.

<https://doi.org/10.1371/journal.pone.0298651.g008>

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

Further inspection of the right column of Fig 8 shows that the posterior distributions of the individual effects of pupil dilation e^{p_1} 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 clustering, based on the Euclidean distance (similarly as in Fig 7). 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^{p_1} to e^{p_2} : trending downwards and upwards respectively. Quite importantly, the correspondence between the groupings observed in Figs 7 and 8 is a result of the posterior inference, not necessarily implied by the structure of our model. The differences in state trajectories between subjects lie in the state occupancy when pupil dilation is not higher than the reference, despite all subjects tending to transition to state 1 when raising their arm.

More specifically, subjects clustered in the first half of Fig 8 (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 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.

5 Conclusion

We have proposed a multi-subject Bayesian approach for estimating dynamic FC 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 effects of time-varying covariates on the connectivity dynamics. We have applied our model to multi-subject resting state fMRI data with pupillary physiological data and we 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 we 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) and (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, in general, the number of states could be learned by assuming a Bayesian-nonparametric specification where the number of FC states is learned directly from the data (see, for example [72, 73]). However, the computational complexity of the inferential algorithm would increase considerably. Variational Bayes approaches could be implemented to obtain approximate inferences 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.

Supporting information

S1 File.
(PDF)

Author Contributions

Conceptualization: Ryan Warnick, Marina Vannucci, Michele Guindani.

Data curation: Jaylen Lee, Sana Hussain, Isaac Menchaca, Aaron R. Seitz, Xiaoping Hu, Megan A. K. Peters, Michele Guindani.

Formal analysis: Jaylen Lee, Ryan Warnick, Megan A. K. Peters, Michele Guindani.

Funding acquisition: Aaron R. Seitz, Xiaoping Hu, Megan A. K. Peters, Michele Guindani.

Investigation: Jaylen Lee, Aaron R. Seitz, Michele Guindani.

Methodology: Jaylen Lee, Aaron R. Seitz, Xiaoping Hu, Megan A. K. Peters, Michele Guindani.

Project administration: Michele Guindani.

Resources: Megan A. K. Peters, Michele Guindani.

Software: Ryan Warnick, Michele Guindani.

Supervision: Marina Vannucci, Michele Guindani.

Validation: Jaylen Lee, Sana Hussain, Megan A. K. Peters.

Visualization: Jaylen Lee, Megan A. K. Peters.

Writing – original draft: Jaylen Lee, Sana Hussain, Ryan Warnick, Marina Vannucci, Megan A. K. Peters, Michele Guindani.

Writing – review & editing: Jaylen Lee, Sana Hussain, Ryan Warnick, Marina Vannucci, Megan A. K. Peters, Michele Guindani.

References

1. Friston KJ, Jezzard P, Turner R. Analysis of Functional MRI Time-Series. *Human Brain Mapping*. 1994; 1(2):153–171. <https://doi.org/10.1002/hbm.460010207>
2. Hutchison RM, Womelsdorf T, Allen EA, Bandettini PA, Calhoun VD, Corbetta M, et al. Dynamic functional connectivity: Promise, issues, and interpretations. *NeuroImage*. 2013; 80:360–378. <https://doi.org/10.1016/j.neuroimage.2013.05.079> PMID: 23707587
3. Fornito A, Zalesky A, Breakspear M. Graph analysis of the human connectome: Promise, progress, and pitfalls. *NeuroImage*. 2013; 80:426–444. <https://doi.org/10.1016/j.neuroimage.2013.04.087> PMID: 23643999
4. Friston KJ. Functional and Effective Connectivity: A Review. *Brain Connectivity*. 2011; 1(1):13–36. <https://doi.org/10.1089/brain.2011.0008> PMID: 22432952
5. Li J, Wang Z, Palmer S, McKeown M. Dynamic Bayesian network modeling of fMRI: a comparison of group-analysis methods. *NeuroImage*. 2008; 41(2):398–407. <https://doi.org/10.1016/j.neuroimage.2008.01.068> PMID: 18406629

6. Vidaurre D, Smith SM, Woolrich MW. Brain network dynamics are hierarchically organized in time. *Proceedings of the National Academy of Sciences*. 2017; 114(48):12827–12832. <https://doi.org/10.1073/pnas.1705120114> PMID: 29087305
7. Lurie DJ, Kessler D, Bassett DS, Betzel RF, Breakspear M, Kheilholz S, et al. Questions and controversies in the study of time-varying functional connectivity in resting fMRI. *Network neuroscience* (Cambridge, Mass). 2020; 4(1):30–69. https://doi.org/10.1162/netn_a_00116 PMID: 32043043
8. Calhoun VD, Miller R, Pearlson G, Adali T. The Chronnectome: Time-Varying Connectivity Networks as the Next Frontier in fMRI Data Discovery. *Neuron*. 2014; 84(2):262–274. <https://doi.org/10.1016/j.neuron.2014.10.015> PMID: 25374354
9. Chand T, Li M, Jamalabadi H, Wagner G, Lord A, Alizadeh S, et al. Heart Rate Variability as an Index of Differential Brain Dynamics at Rest and After Acute Stress Induction. *Frontiers in Neuroscience*. 2020; 14. <https://doi.org/10.3389/fnins.2020.00645> PMID: 32714132
10. El-Baba M, Lewis DJ, Fang Z, Owen AM, Fogel SM, Morton JB. Functional connectivity dynamics slow with descent from wakefulness to sleep. *PLoS ONE*. 2019; 14(12):1–13. <https://doi.org/10.1371/journal.pone.0224669> PMID: 31790422
11. Kucyi A, Hove MJ, Esterman M, Hutchison RM, Valera EM. Dynamic Brain Network Correlates of Spontaneous Fluctuations in Attention. *Cerebral cortex* (New York, NY: 1991). 2017; 27(3):1831–1840. <https://doi.org/10.1093/cercor/bhw029> PMID: 26874182
12. Sobczak F, Pais-Roldán P, Takahashi K, Yu X. Decoding the brain state-dependent relationship between pupil dynamics and resting state fMRI signal fluctuation. *eLife*. 2021; 10. <https://doi.org/10.7554/eLife.68980> PMID: 34463612
13. Siegle GJ, Steinhauer SR, Stenger VA, Konecky R, Carter CS. Use of concurrent pupil dilation assessment to inform interpretation and analysis of fMRI data. *NeuroImage*. 2003; 20(1):114–124. [https://doi.org/10.1016/S1053-8119\(03\)00298-2](https://doi.org/10.1016/S1053-8119(03)00298-2) PMID: 14527574
14. van der Wel P, van Steenbergen H. Pupil dilation as an index of effort in cognitive control tasks: A review. *Psychonomic Bulletin & Review*. 2018; 25(6):2005–2015. <https://doi.org/10.3758/s13423-018-1432-y> PMID: 29435963
15. Martin CG, He BJ, Chang C. State-related neural influences on fMRI connectivity estimation. *NeuroImage*. 2021; 244:118590. <https://doi.org/10.1016/j.neuroimage.2021.118590> PMID: 34560268
16. Haimovici A, Tagliazucchi E, Balenzuela P, Laufs H. On wakefulness fluctuations as a source of BOLD functional connectivity dynamics. *Scientific reports*. 2017; 7(1):5908–5908. <https://doi.org/10.1038/s41598-017-06389-4> PMID: 28724928
17. Li L, Pluta D, Shahbaba B, Fortin N, Ombao H, Baldi P. Modeling dynamic functional connectivity with latent factor Gaussian processes. *Adv Neural Inf Process Syst*. 2019; 32:8263–8273. PMID: 33041607
18. Spencer APC, Goodfellow M. Using deep clustering to improve fMRI dynamic functional connectivity analysis. *NeuroImage*. 2022; 257(119288):119288. <https://doi.org/10.1016/j.neuroimage.2022.119288> PMID: 35551991
19. Hussain S, Langley J, Seitz AR, Hu XP, Peters MAK. A Novel Hidden Markov Approach to Studying Dynamic Functional Connectivity States in Human Neuroimaging. *Brain Connectivity*. 2023; 13(3):154–163. <https://doi.org/10.1089/brain.2022.0031> PMID: 36367193
20. Lindquist MA, Xu Y, Nebel MB, Caffo BS. Evaluating Dynamic Bivariate Correlations in Resting-state fMRI: A comparison study and a new approach. *NeuroImage*. 2014; 101(1):531–546. <https://doi.org/10.1016/j.neuroimage.2014.06.052> PMID: 24993894
21. Shakil S, Lee CH, Keilholz SD. Evaluation of sliding window correlation performance for characterizing dynamic functional connectivity and brain states. *NeuroImage*. 2016; 133:111–128. <https://doi.org/10.1016/j.neuroimage.2016.02.074> PMID: 26952197
22. Cribben I, Haraldsdottir R, Atlas LY, Wager TD, Lindquist MA. Dynamic connectivity regression: Determining state-related changes in brain connectivity. *NeuroImage*. 2012; 61(4):907–920. <https://doi.org/10.1016/j.neuroimage.2012.03.070> PMID: 22484408
23. Cribben I, Wager TD, Lindquist MA. Detecting functional connectivity change points for single-subject fMRI data. *Frontiers in Computational Neuroscience*. 2013; 7(October):1–15. <https://doi.org/10.3389/fncom.2013.00143> PMID: 24198781
24. Xu Y, Lindquist MA. Dynamic Connectivity Detection: An algorithm for determining functional connectivity change points in fMRI data. *Frontiers in Neuroscience*. 2015; 9(JUL). <https://doi.org/10.3389/fnins.2015.00285> PMID: 26388711
25. Friedman J, Hastie T, Tibshirani R. Sparse Inverse Covariance Estimation with the Graphical Lasso. *Biostatistics*. 2008; 9(3):432–441. <https://doi.org/10.1093/biostatistics/kxm045> PMID: 18079126

26. Xiong X, Cribben I. Beyond Linear Dynamic Functional Connectivity: A Vine Copula Change Point Model. *Journal of Computational and Graphical Statistics*. 2023; 32(3):853–872. <https://doi.org/10.1080/10618600.2022.2127738>
27. Warnick R, Guindani M, Erhardt E, Allen E, Calhoun V, Vannucci M. A Bayesian Approach for Estimating Dynamic Functional Network Connectivity in fMRI Data. *Journal of the American Statistical Association*. 2018; 113(521):134–151. <https://doi.org/10.1080/01621459.2017.1379404> PMID: 30853734
28. Long Z, Liu X, Niu Y, Shang H, Lu H, Zhang J, et al. Improved dynamic functional connectivity estimation with an alternating hidden Markov model. *Cognitive Neurodynamics*. 2022; 17(5):1381–1398. <https://doi.org/10.1007/s11571-022-09874-3> PMID: 37786659
29. Sourty M, Thoraval L, Roquet D, Armspach JP, Foucher J, Blanc F. Identifying dynamic functional connectivity changes in dementia with lewy bodies based on product hidden Markov models. *Frontiers in Computational Neuroscience*. 2016; 10(Jun):1–11. <https://doi.org/10.3389/fncom.2016.00060> PMID: 27445778
30. Chiang S, Cassese A, Guindani M, Vannucci M, Yeh HJ, Haneef Z, et al. Time-Dependence of Graph Theory Metrics in Functional Connectivity Analysis. *NeuroImage*. 2015; 125:601–615. <https://doi.org/10.1016/j.neuroimage.2015.10.070> PMID: 26518632
31. Ting C, Ombao H, Sh-Hussein. Estimating Dynamic Connectivity States in fMRI Using Regime-Switching Factor Models. *IEEE Transactions on Medical Imaging*. 2018; 37:1011–1023. <https://doi.org/10.1109/TMI.2017.2780185> PMID: 29610078
32. Ombao H, Fiecas M, Ting C, Low Y. Statistical Models for Brain Signals with Properties that Evolve Across Trials. *NeuroImage*. 2018; 180:609–618. <https://doi.org/10.1016/j.neuroimage.2017.11.061> PMID: 29223740
33. Kundu S, Ming J, Pierce J, McDowell J, Guo Y. Estimating dynamic brain functional networks using multi-subject fMRI data. *NeuroImage*. 2018; 183:635–649. <https://doi.org/10.1016/j.neuroimage.2018.07.045> PMID: 30048750
34. Kundu S, Ming J, Nocera J, McGregor KM. Integrative learning for population of dynamic networks with covariates. *NeuroImage*. 2021; 236:118181. <https://doi.org/10.1016/j.neuroimage.2021.118181> PMID: 34022384
35. Carvalho CM, Polson NG, Scott JG. The horseshoe estimator for sparse signals. *Biometrika*. 2010; 97(2):465–480. <https://doi.org/10.1093/biomet/asq017>
36. Müller P, Parmigiani G, Rice K. FDR and Bayesian multiple comparisons rules. *bepress*; 2006.
37. Joshi S, Li Y, Kalwani RM, Gold JI. Relationships between Pupil Diameter and Neuronal Activity in the Locus Coeruleus, Colliculi, and Cingulate Cortex. *Neuron*. 2016; 89(1):221–234. <https://doi.org/10.1016/j.neuron.2015.11.028> PMID: 26711118
38. Hussain S, Menchaca I, Shalchy MA, Yaghoubi K, Langley J, Seitz AR, et al. Locus coeruleus neuromelanin predicts ease of attaining and maintaining neural states of arousal. *bioRxiv*. 2022;.
39. Friston KJ. Functional and effective connectivity in neuroimaging: a synthesis. *Human Brain Mapping*. 1994; 2:56–78. <https://doi.org/10.1002/hbm.460020107>
40. Fair DA, Schlaggar BL, Cohen AL, Miezin FM, Dosenbach NUF, Wenger KK, et al. A method for using blocked and event-related fMRI data to study “resting state” functional connectivity. *NeuroImage*. 2007; 35(1):396–405. <https://doi.org/10.1016/j.neuroimage.2006.11.051> PMID: 17239622
41. Wilson KA, James GA, Kilts CD, Bush KA. Combining Physiological and Neuroimaging Measures to Predict Affect Processing Induced by Affectively Valent Image Stimuli. *Scientific Reports*. 2020; 10(1):9298. <https://doi.org/10.1038/s41598-020-66109-3> PMID: 32518277
42. Wang J, Zuo X, He Y. Graph-based network analysis of resting-state functional MRI. *Frontiers in Systems Neuroscience*. 2010; 4(June):1–14. <https://doi.org/10.3389/fnsys.2010.00016> PMID: 20589099
43. Essen V, Tononi G. An Introduction to Brain Networks. In: Fornito A, Zalesky A, Bullmore ET, editors. *Fundamentals of Brain Network Analysis*. San Diego: Academic Press; 2016. p. 1–35. Available from: <https://www.sciencedirect.com/science/article/pii/B9780124079083000017>.
44. Li Y, Craig BA, Bhadra A. The Graphical Horseshoe Estimator for Inverse Covariance Matrices. *Journal of Computational and Graphical Statistics*. 2019; 28(3):747–757. <https://doi.org/10.1080/10618600.2019.1575744>
45. Holsclaw T, Greene AM, Robertson AW, Smyth P. Bayesian nonhomogeneous Markov models via Pólya-Gamma data augmentation with applications to rainfall modeling. *Annals of Applied Statistics*. 2017; 11(1):393–426. <https://doi.org/10.1214/16-AOAS1009>
46. Polson NG, Scott JG, Windle J. Bayesian inference for logistic models using Pólya-Gamma latent variables. *Journal of the American Statistical Association*. 2013; 108(504):1339–1349. <https://doi.org/10.1080/01621459.2013.829001>

47. Scott SL. Bayesian Methods for Hidden Markov Models: Recursive Computing in the 21st Century. *Journal of the American Statistical Association*. 2002; 97(457):337–351. <https://doi.org/10.1198/016214502753479464>
48. Holmes CC, Held L. Bayesian auxiliary variable models for binary and multinomial regression. *Bayesian Analysis*. 2006; 1(1):145–168.
49. Makalic E, Schmidt DF. A simple sampler for the horseshoe estimator. *IEEE Signal Processing Letters*. 2016; 23(1):179–182. <https://doi.org/10.1109/LSP.2015.2503725>
50. Wang H. Bayesian graphical lasso models and efficient posterior computation. *Bayesian Analysis*. 2012; 7(4):867–886. <https://doi.org/10.1214/12-BA729>
51. Müller P, Parmigiani G, Rice K. FDR and Bayesian Multiple Comparisons Rules. In: Bernardo JM, Bayarri MJ, Berger JO, Dawid AP, Heckerman D, Smith AFM, et al., editors. *Bayesian Statistics 8*. Oxford, UK: Oxford University Press.; 2007.
52. Gordy MB. A generalization of generalized beta distributions. Board of Governors of the Federal Reserve System (U.S.); 1998. 1998-18.
53. Bhadra A, Datta J, Li Y, Polson NG, Willard B. Prediction Risk for the Horseshoe Regression. *Journal of Machine Learning Research*. 2019; 20(78):1–39.
54. Newton MA, Noueiry A, Sarkar D, Ahlquist P. Detecting differential gene expression with a semiparametric hierarchical mixture method. *Biostatistics*. 2004; 5:155–176. <https://doi.org/10.1093/biostatistics/5.2.155> PMID: 15054023
55. Chandra NK, Mueller P, Sarkar A. Bayesian Scalable Precision Factor Analysis for Massive Sparse Gaussian Graphical Models; 2021.
56. Allen EA, Damaraju E, Plis SM, Erhardt EB, Eichele T, Calhoun VD. Tracking whole-brain connectivity dynamics in the resting state. *Cerebral Cortex*. 2014; 24(3):663–676. <https://doi.org/10.1093/cercor/bhs352> PMID: 23146964
57. Erhardt EB, Allen EA, Wei Y, Eichele T, Calhoun VD. SimTB, a simulation toolbox for fMRI data under a model of spatiotemporal separability. *NeuroImage*. 2012; 4(59):4160–4167. <https://doi.org/10.1016/j.neuroimage.2011.11.088>
58. Lindquist MA, Loh JM, Atlas LY, Wager TD. Modeling the Hemodynamic Response Function in fMRI: Efficiency, Bias, and Mis-modeling. *NeuroImage*. 2009; 45:187–198. <https://doi.org/10.1016/j.neuroimage.2008.10.065> PMID: 19084070
59. Mather M, Huang R, Clewett D, Nielsen SE, Velasco R, Tu K, et al. Isometric exercise facilitates attention to salient events in women via the noradrenergic system. *NeuroImage*. 2020; 210:116560. <https://doi.org/10.1016/j.neuroimage.2020.116560> PMID: 31978545
60. Smith SM, Jenkinson M, Woolrich MW, Beckmann CF, Behrens TEJ, Johansen-Berg H, et al. Advances in functional and structural MR image analysis and implementation as FSL. *NeuroImage*. 2004; 23:S208–S219. <https://doi.org/10.1016/j.neuroimage.2004.07.051> PMID: 15501092
61. Woolrich MW, Jbabdi S, Patenaude B, Chappell M, Makni S, Behrens T, et al. Bayesian analysis of neuroimaging data in FSL. *NeuroImage*. 2009; 45(1):S173–S186. <https://doi.org/10.1016/j.neuroimage.2008.10.055> PMID: 19059349
62. Greicius MD, Menon V. Default-Mode Activity during a Passive Sensory Task: Uncoupled from Deactivation but Impacting Activation. *Journal of Cognitive Neuroscience*. 2004; 16(9):1484–1492. <https://doi.org/10.1162/0898929042568532> PMID: 15601513
63. Menon V, Uddin LQ. Saliency, switching, attention and control: a network model of insula function. *Brain Struct Funct*. 2010; 214(5-6):655–667. <https://doi.org/10.1007/s00429-010-0262-0> PMID: 20512370
64. Deshpande G, Santhanam P, Hu X. Instantaneous and causal connectivity in resting state brain networks derived from functional MRI data. *NeuroImage*. 2011; 54(2):1043–1052. <https://doi.org/10.1016/j.neuroimage.2010.09.024> PMID: 20850549
65. Raichle ME. The Restless Brain. *Brain Connectivity*. 2011; 1(1):3–12. <https://doi.org/10.1089/brain.2011.0019> PMID: 22432951
66. Laird AR, Lancaster JL, Fox PT. BrainMap: The Social Evolution of a Human Brain Mapping Database. *Neuroinformatics*. 2005; 3(1):065–078. <https://doi.org/10.1385/NI:3:1:065> PMID: 15897617
67. Lancaster JL, Tordesillas-Gutiérrez D, Martínez M, Salinas F, Evans A, Zilles K, et al. Bias between MNI and Talairach coordinates analyzed using the ICBM-152 brain template. *Hum Brain Mapp*. 2007; 28(11):1194–1205. <https://doi.org/10.1002/hbm.20345> PMID: 17266101
68. Langley J, Hussain S, Flores JJ, Bennett IJ, Hu X. Characterization of age-related microstructural changes in locus coeruleus and substantia nigra pars compacta. *Neurobiology of Aging*. 2020; 87:89–97. <https://doi.org/10.1016/j.neurobiolaging.2019.11.016> PMID: 31870645

69. Deshpande G, LaConte S, James GA, Peltier S, Hu X. Multivariate Granger causality analysis of fMRI data. *Hum Brain Mapp.* 2009; 30(4):1361–1373. <https://doi.org/10.1002/hbm.20606> PMID: 18537116
70. Stilla R, Deshpande G, LaConte S, Hu X, Sathian K. Posteromedial Parietal Cortical Activity and Inputs Predict Tactile Spatial Acuity. *Journal of Neuroscience.* 2007; 27(41):11091–11102. <https://doi.org/10.1523/JNEUROSCI.1808-07.2007> PMID: 17928451
71. Joshi S, Gold JI. Context-dependent relationships between locus coeruleus firing patterns and coordinated neural activity in the anterior cingulate cortex. *eLife.* 2022; 11:e63490. <https://doi.org/10.7554/eLife.63490> PMID: 34994344
72. Beal MJ, Ghahramani Z, Rasmussen CE. The Infinite Hidden Markov Model. *NIPS.* 2002; 14:577–584.
73. Fox EB, Sudderth EB, Jordan MI, Willsky AS. A Sticky HDP-HMM with Application to Speaker Diarization. *The Annals of Applied Statistics.* 2011; 5:1020–1056. <https://doi.org/10.1214/10-AOAS395>