



A Bayesian Multiplex Graph Classifier of Functional Brain Connectivity Across Diverse Tasks of Cognitive Control

Sharmistha Guha¹ · Jose Rodriguez-Acosta¹ · Ivo D. Dinov^{2,3}

Accepted: 22 May 2024 / Published online: 11 June 2024

© The Author(s), under exclusive licence to Springer Science+Business Media, LLC, part of Springer Nature 2024

Abstract

This article seeks to investigate the impact of aging on functional connectivity across different cognitive control scenarios, particularly emphasizing the identification of brain regions significantly associated with early aging. By conceptualizing functional connectivity within each cognitive control scenario as a graph, with brain regions as nodes, the statistical challenge revolves around devising a regression framework to predict a binary scalar outcome (aging or normal) using multiple graph predictors. Popular regression methods utilizing multiplex graph predictors often face limitations in effectively harnessing information within and across graph layers, leading to potentially less accurate inference and predictive accuracy, especially for smaller sample sizes. To address this challenge, we propose the *Bayesian Multiplex Graph Classifier (BMGC)*. Accounting for multiplex graph topology, our method models edge coefficients at each graph layer using bilinear interactions between the latent effects associated with the two nodes connected by the edge. This approach also employs a variable selection framework on node-specific latent effects from all graph layers to identify influential nodes linked to observed outcomes. Crucially, the proposed framework is computationally efficient and quantifies the uncertainty in node identification, coefficient estimation, and binary outcome prediction. BMGC outperforms alternative methods in terms of the aforementioned metrics in simulation studies. An additional BMGC validation was completed using an fMRI study of brain networks in adults. The proposed BMGC technique identified that sensory motor brain network obeys certain lateral symmetries, whereas the default mode network exhibits significant brain asymmetries associated with early aging.

Keywords Aging · Bayesian statistics · Functional brain connectivity · Multiplex graph classification · Variable selection

Introduction

In recent years, the study of cognitive control among aging adults has gained significance as a prominent and essential research focus. With the ongoing aging of the population, there is a growing need to comprehend the nuances of cognitive functions in older individuals. This understanding is crucial not just because of the healthcare costs involved but also due to the significant economic and social consequences linked with an aging populace. The field of cognitive neuroscience stands to acquire valuable insights into the aging brain by examining functional connectivity in the human brain during cognitive control tasks. This paper utilizes brain imaging data, representing functional connectivity graphs derived from functional magnetic resonance imaging (fMRI) (Dinov, 2023; Richiardi et al., 2011), to investigate the interplay between cognitive decline and brain regions. The fMRI study focused on two distinct cognitive

Sharmistha Guha and Jose Rodriguez-Acosta contributed equally to this work.

✉ Sharmistha Guha
sharmistha@tamu.edu

Jose Rodriguez-Acosta
jeroda7105@tamu.edu

Ivo D. Dinov
dinov@umich.edu

¹ Department of Statistics, Texas A&M University, 3143 TAMU, College Station 77843, TX, USA

² Statistics Online Computational Resource, University of Michigan, 426 N. Ingalls St., Ann Arbor 48109, MI, USA

³ Department of Computational Medicine and Bioinformatics, University of Michigan, 100 Washtenaw Avenue, Ann Arbor 48109, MI, USA

task conditions: inhibition and initiation (Rieck et al., 2021a). Inhibition is broadly conceptualized as the ability to suppress or countermand a thought, action, or feeling. Initiation, on the other hand, refers to engaging in purposeful goal-directed activities. For each participant, there are two graphs with nodes as brain regions of interest (ROIs), each corresponding to one of these task conditions. This dual perspective enables us to uncover how aging affects functional connectivity across various cognitive control scenarios, shedding light on the nuanced differences in brain network interactions between inhibition and initiation tasks in aging adults. Using the Schaefer 200 network of fMRI connectivity (Yang et al., 2023) for two task conditions (Z-transformed correlation matrices *input*), we examined the left and right hemisphere network differences in the brains of normal and aging adults, as quantified by the Mini-Mental State Examination score (*outcome*).

From a statistical standpoint, the main challenge lies in devising a regression framework that incorporates a binary outcome alongside a two-layer multiplex graph as the predictor. Our primary objective is to achieve classification of the binary outcome, identify the graph nodes significantly associated with the outcome, and estimate the regression impact of each layer of the multiplex graph on the outcome. The modeling framework proposed in this article enables us to tackle these challenges simultaneously. Although literature on generalized linear models involving a multiplex graph predictor is limited, some previous studies have addressed unsupervised modeling of multiplex graphs. However, these existing approaches do not adequately meet all our inferential objectives, as outlined below.

In regression involving a single-layer graph predictor, commonly used methods often entail transforming the graph into a high-dimensional set of edge weights (Craddock et al., 2009; Richiardi et al., 2011). Subsequent inference is drawn utilizing developments in generic high-dimensional regression architectures (Carvalho et al., 2010; Park & Casella, 2008; Tibshirani, 1996), or neural network (NN) models (Dinh & Ho, 2020; Dinov & Velez, 2021; Polson & Ročková, 2018). While these approaches can be extended to handle multiplex graphs straightforwardly, restructuring the multiplex graph using these methods might inadequately capture the effects of intricate interconnections among nodes within different layers on the outcome. This potential limitation could compromise the precision and interpretability of the regression model.

Recent advancements in the realm of object regressions, featuring a scalar outcome and a single-layer graph predictor, have demonstrated promising capabilities in leveraging the graph's structural information. Specifically, Guha and Rodriguez (2021, 2023) have introduced innovative Bayesian shrinkage priors designed to estimate the effects of nodes and edges within the graph on the

scalar outcome. By exploiting the topology of the graph predictor, these proposed methods exhibit accurate outcome prediction, robust parameter inference, and precise uncertainty estimates, especially in identifying influential graph nodes and edges. Moreover, Guha and Guhaniyogi (2021) have introduced a Bayesian generalized linear modeling framework tailored for brain connectome data, incorporating low-rankness and group sparsity to identify significant associations between a symmetric tensor response and scalar predictors. Additionally, Guha and Guhaniyogi (2024) have formulated a nonparametric Bayesian mixture modeling framework for clustering subjects based on their individual networks and covariates, accounting for differences in relationships across subject groups and identifying network nodes associated with each covariate within each cluster. More recently, there has been a growing body of work on interpretable graph neural network approaches that provide theoretically guaranteed predictions for scalar outcomes based on graph inputs. These approaches differ from earlier literature on graph neural networks (GNN) (Liu & Zhou, 2022; Zhou et al., 2020) in their ability to offer statistical interpretability for the constructed GNN (Zhou et al., 2023). Despite their successes, these approaches are tailored for single-layer graph predictors. Extending them to handle interactions between nodes within different layers for a multiplex graph predictor to offer inference on influential nodes presents nontrivial modeling challenges.

An alternative strategy involves the aggregation of graphs from different layers to form a tensor, which is then employed to construct a regression framework with a scalar outcome and the tensor predictor. This approach can leverage recent advancements in tensor regression, encompassing penalized optimization (Fan et al., 2019; Zhou et al., 2013), low-rank methods (Ahmed et al., 2020; He et al., 2018), and the Bayesian multiway shrinkage literature (Guhaniyogi et al., 2017; Spencer et al., 2022). However, these approaches do not explicitly account for the symmetry constraint in individual layers of the multiplex graph. Moreover, their primary emphasis is on prediction and identifying significant edges or interconnections, rather than specifically identifying crucial nodes in each layer that influence the outcome.

Within the current body of literature, multiplex graphs are often regarded as random variables rather than being utilized as predictors in a regression framework. In this literature, researchers predominantly concentrate on establishing appropriate relationships among edges and various types of connections defining the multiple layers (Gollini & Murphy, 2016; Han et al., 2015; Heaney, 2014). These advancements have expanded the use of exponential random graph models (Frank & Strauss, 1986; Holland & Leinhardt, 1981) and latent variable models (Airoldi et al., 2008; Hoff et al., 2002;

Nowicki & Snijders, 2001), initially designed for a single graph, to facilitate inference in multiplex graphs. Furthermore, these frameworks have been extended to handle time-varying or dynamic multiplex graphs. In such scenarios, stochastic processes are employed to leverage information across edges and over time, as noted in previous research (Durante et al., 2017; Hoff, 2015; Snijders et al., 2013). Additionally, these frameworks have also been adapted to accommodate multiplex graphs with covariates associated with the graph nodes, as demonstrated in prior studies (Contisciani et al., 2020; Xu et al., 2023; Zhang et al., 2022). Nevertheless, these methods for multiplex graphs are fundamentally unsupervised and do not align with our specific inferential goals.

There is a burgeoning area of research in graph representation learning which focuses on efficiently encoding high-dimensional sparse graph-structured data into low-dimensional dense vectors. This task has been extensively investigated across various fields, including machine learning and data mining (Chen et al., 2020; Chen et al., 2021; Heidari & Iosifidis, 2021; Li et al., 2023; Ngoc et al., 2020; Wang et al., 2019; Zhao et al., 2022). Although much of this literature is oriented towards unsupervised learning, there is potential to adapt these methods to our context. In such a scenario, the low-dimensional features extracted from a multiplex graph can be regressed on a binary outcome, potentially leading to effective classification of the binary outcome. Inference on the influential graph nodes that are significantly associated with the binary outcome, while providing uncertainty estimates.

This article presents Bayesian Multiplex Graph Classifier (BMGC), which is a novel generalized linear modeling framework tailored for a binary outcome and a multiplex graph predictor. Within our framework, an edge coefficient within each layer is expressed as a function of latent effects associated with the nodes linked by that specific edge. Specifically, these edge coefficients are modeled as a bilinear interaction between node-specific latent effects, resulting in a low-rank structure for the coefficients corresponding to the multiplex graph at each layer. This low-rank structure achieves parsimony, particularly beneficial when dealing with multiplex graphs defined on a large number of nodes. To facilitate the identification of network nodes significantly associated with the outcome, we propose a variable selection framework that operates jointly on latent effects across all layers corresponding to a particular node. This structured approach not only ensures simplicity but also enables precise classification of the binary outcome, facilitates inference on network nodes and edges relevant to the outcome, and provides well-calibrated uncertainties for both inference and classification. We choose the Bayesian framework for inference due to its natural ability to provide uncertainty in inference, especially in identifying influential nodes. This is

particularly essential when dealing with a moderate sample size, with the number of graph edges exceeding the sample size by a significant margin. While our research originates from a neuro-scientific study, its versatility extends beyond this domain, making it broadly applicable in various scenarios involving multiplex graphs, such as genomics data.

The rest of the article proceeds as follows. “[Functional Connectivity Data during Cognitive Control for an Adult Lifespan Sample](#)” section discusses functional connectivity data during different tasks on cognitive control. “[Model and Prior Formulation](#)” section provides a detailed description of the regression framework with a multiplex graph predictor and prior formulation on model coefficients. “[Posterior Inference](#)” section discusses posterior computation. Simulation studies in “[Simulation Study](#)” section and the analysis of neuroimaging data in “[Functional Magnetic Resonance Imaging \(fMRI\) Study](#)” section demonstrate superior performance of the proposed approach compared to its competitors. Finally, “[Conclusion and Future Work](#)” section offers concluding remarks and possible future extensions. [Appendix](#) describes full conditional posterior distributions for all the model parameters. They are employed to construct Markov Chain Monte Carlo for model computation.

Functional Connectivity Data during Cognitive Control for an Adult Lifespan Sample

In this section, we describe the functional connectivity data, derived from functional magnetic resonance imaging (fMRI) of healthy aging adults, recorded under different cognitive tasks (Rieck et al., 2021b). In this study, 144 individuals (ages 20–86) were recruited from the greater Toronto area to participate in a study focusing on functional activity during cognitive control. The participants underwent scanning using a Siemens 3T MRI scanner as they performed cognitive control tasks designed to gauge functional activity during inhibition and initiation. Throughout the scanning process, participants reclined on the scanner bed with a mirror positioned in front of their face, reflecting a computer screen displaying the experimental stimuli. Participants used earplugs and wore headphones to mitigate scanner noise, and communication with the experimenter between scans was facilitated through a microphone. Holding a response box in their right hand, participants used their index and middle fingers to respond to stimuli presented on the screen (specific finger responses were counterbalanced among participants). In all tasks, letters were displayed in the center of the computer screen using the Courier New font against a dark gray background. Inhibition and initiation were measured employing a go/no-go paradigm, wherein participants were exposed to a

sequence of uppercase letters. Their instruction was to respond (termed “go”) when encountering the letter “X” and refrain from responding (termed “no-go”) for all other letters. Non-X stimuli (i.e., “no-go” stimuli) were selected randomly from a pool of 20 other letters: A, B, C, D, E, F, G, H, I, J, L, M, N, O, P, Q, R, S, T, and U. The task was divided into an “inhibition” block, characterized by a higher number of go trials than no-go trials (120 go, 40 no-go), and a shorter “initiation” block, where there were more no-go trials than go trials (20 go, 60 no-go). The sequence of these two blocks was randomized among participants. Letter stimuli were displayed for a duration of 400 ms, accompanied by an average inter-stimulus interval (marked by a fixation cross) of 1200 ms, with random jittering ranging between 900 and 1500 ms. The overall duration of the go/no-go task amounted to 6 min and 24 s. Blood-oxygen-level dependent (BOLD) fMRI data were collected with a 12-channel head coil using an echo-planar imaging sequence with 40 axial slices acquired parallel to the anterior-posterior commissure. A total of 216 volumes were collected for the go/no-go task. High resolution anatomical scans used for warping the BOLD images to MNI space were acquired with a T1-weighted MP-RAGE sequence in which 160 axial slices were collected. Functional data for each task were preprocessed with a mix of AFNI functions as well as Octave and MATLAB scripts using the Optimizing of Preprocessing Pipelines for Neuro-Imaging software package (OPPNI) (Churchill et al., 2017). For the current dataset, a fixed pipeline for all participants was conducted with the following steps: (1) rigid-body alignment of the timeseries to correct for movement; (2) removal and interpolation of outlier volumes; (3) correction for physiological (i.e., cardiac and respiratory) noise; (4) slice timing correction; (5) spatial smoothing with a 6 mm smoothing kernel; (6) temporal detrending; (7) regression of six motion parameter estimates (X, Y, and Z translation and rotation) on the timeseries; (8) regression of signal in tissue of no interest (white matter, vessels and cerebrospinal fluid) on the time-series; and finally (9) warping to MNI space and resampling to 4 mm 3 isotropic voxel.

For each participant, there are two correlation matrices corresponding to the two distinct task conditions: inhibition and initiation. Assessment of functional connectivity (quantified with time-series correlations between different brain regions) were computed using the Schaefer 200 parcel 17 network atlas (Schaefer et al., 2018; Yeo et al., 2011). Functional connectivity was computed for the fMRI task using the CONN toolbox in MATLAB (Whitfield-Gabrieli & Nieto-Castanon, 2012). Functional time courses for each task condition were extracted from the 200 parcel nodes that comprised 17 networks, resulting in 200×200 connectivity matrices. We used the bilateral 100 brain regions of interest (ROIs) collectively

representing 200 nodes covering 17 different known brain-networks. Z-transformed functional connectivity correlation matrices were computed for these Schaefer-100 nodes to examine the underlying neurobiology of aging.

The Mini-Mental State Examination (MMSE) score (Mega et al., 1996; Moon et al., 2023) represents the *binary cognitive outcome measure* of interest we want to predict using the our novel Bayesian multiplex graph modelling framework. In this supervised prediction study, the outcome variable represents the integer-valued MMSE measure binarized into *normal* ($0 \leq \text{MMSE} \leq 29$) and *aging* ($30 = \text{MMSE}$) cohorts.

Model and Prior Formulation

Bayesian Multiplex Graph Classifier

For each subject, we observe a binary outcome $y \in \{0, 1\}$, a p -dimensional predictor $\mathbf{x}$, and a L -layer multiplex graph predictor $\{\mathcal{G}_\alpha\}_{\alpha=1}^L$. A L -layer multiplex graph is defined as a sequence of L graphs: $\{\mathcal{G}_\alpha\}_{\alpha=1}^L = \{(\mathcal{N}_\alpha, \mathcal{E}_\alpha)\}_{\alpha=1}^L$, with the α th graph consisting of nodes $\mathcal{N}_\alpha$ and edges $\mathcal{E}_\alpha \subset \mathcal{N}_\alpha \times \mathcal{N}_\alpha$. We assume that the node sets are same across different layers, i.e., $\mathcal{N}_\alpha = \mathcal{N}_{\alpha'} = \mathcal{N}$, for any $1 \leq \alpha \neq \alpha' \leq L$. This common set of nodes is denoted by $\mathcal{N} = \{\mathcal{N}^{(1)}, \dots, \mathcal{N}^{(V)}\}$, where $|\mathcal{N}| = V$ represents the number of nodes. We construct a $V \times V$ adjacency matrix $\mathbf{W}^{(\alpha)} \in \mathbb{R}^{V \times V}$ from the graph predictor at the α -th layer. The entry at position (v, v') of this matrix signifies the strength of association between nodes $\mathcal{N}^{(v)}$ and $\mathcal{N}^{(v')}$, $v, v' = 1, \dots, V$. For the functional connectivity data, this entry represents the Z-transformed correlation coefficient between the functional time series observed at nodes $\mathcal{N}^{(v)}$ and $\mathcal{N}^{(v')}$, under a specific task condition. This article concentrates specifically on undirected graphs without self-relationships within each layer, a characteristic considered scientifically meaningful in the functional connectivity data. This implies that the adjacency matrix $\mathbf{W}^{(\alpha)}$ is symmetric, and its diagonal entries are zero. In what follows, we represent the L -layer multiplex graph by the sequence of adjacency matrices $\mathbf{W} = \{\mathbf{W}^{(\alpha)}\}_{\alpha=1}^L$.

For each subject, the proposed high-dimensional generalized linear model representing the relationship between the binary outcome y and the multiplex graph predictor is given by,

$$P(y = 1 | \mathbf{x}, \mathbf{W}) = G^{-1} \left(\mu + \mathbf{x}^T \boldsymbol{\beta}_x + \frac{1}{2} \sum_{\alpha=1}^L \langle \mathbf{W}^{(\alpha)} | \boldsymbol{\Gamma}^{(\alpha)} \rangle \right). \quad (1)$$

Here, $G(\cdot)$ represents the link function, with popular choices being logit, probit, or log-log functions. Here, μ is the intercept. The coefficient $\boldsymbol{\beta}_x$ is the p -dimensional vector corresponding to $\mathbf{x}$, $\boldsymbol{\Gamma}^{(\alpha)}$ is the coefficient matrix of dimension

$V \times V$ corresponding to the α -th layer of the graph adjacency matrix $\mathbf{W}^{(\alpha)}$, and $\langle \mathbf{W}^{(\alpha)} | \mathbf{\Gamma}^{(\alpha)} \rangle = \text{Trace}(\mathbf{\Gamma}^{(\alpha)T} \mathbf{W}^{(\alpha)})$ denotes the Frobenius inner product of the matrices $\mathbf{W}^{(\alpha)}$ and $\mathbf{\Gamma}^{(\alpha)}$. The Frobenius inner product serves as an extension of the dot product, transitioning from vector spaces to matrix spaces, and it naturally represents the inner product in the space of matrices. Similar to $\mathbf{W}^{(\alpha)}$, we presume that the coefficient matrix $\mathbf{\Gamma}^{(\alpha)}$ is symmetric and has zero entries along its diagonal.

To achieve flexible modeling of the effect of the multiplex graph on the outcome, each graph edge coefficient is expressed as an interaction between the nodes connected by the specific edge. More precisely, let $\xi_v^{(\alpha)} \in \mathbb{R}^H$ denote the H -dimensional coordinate in the latent space, representing the latent effect of the node $\mathcal{N}^{(v)}$ specific to the α th layer, $1 \leq v \leq V$. The coefficient corresponding to the (v, v') th edge at the α -th layer is modeled as a weighted inner product of latent effects for nodes $\mathcal{N}^{(v)}$ and $\mathcal{N}^{(v')}$, given by, $\gamma_{v,v'}^{(\alpha)} = \langle \xi_v^{(\alpha)} | \xi_{v'}^{(\alpha)} \rangle_{\Theta^{(\alpha)}} = \xi_v^{(\alpha)T} \Theta^{(\alpha)} \xi_{v'}^{(\alpha)}$, $1 \leq v \neq v' \leq V$, $\alpha = 1, \dots, L$. This mechanism allows nodes with coordinates aligned in the same direction to positively impact the outcome, whereas nodes with coordinates in opposite directions exert a negative influence on the outcome. The interaction effect is dependent on the $H \times H$ weighing matrix $\Theta^{(\alpha)}$ which is a diagonal matrix with the h -th diagonal entry $\theta_h^{(\alpha)} \in \{-1, 0, 1\}$. The diagonal elements indicate the effect of the h -th dimension of the latent effect $\xi_v^{(\alpha)}$ on the graph coefficients. A value of 0 for $\theta_h^{(\alpha)}$ signifies that the h -th dimension of the latent effect specific to the α -th layer is not informative of the outcome. The aforementioned modeling of the graph coefficient leads to a low-order spectral decomposition of $\mathbf{\Gamma}^{(\alpha)}$, given by, $\mathbf{\Gamma}^{(\alpha)} = \Xi^{(\alpha)T} \Theta^{(\alpha)} \Xi^{(\alpha)}$, where $\Xi^{(\alpha)} = [\xi_1^{(\alpha)} : \dots : \xi_V^{(\alpha)}]$ denotes a $H \times V$ matrix which constitutes node-specific latent effects at the α th layer. The assumed low-order spectral representation massively reduces the count of estimable parameters from $LV(V-1)/2$ to $LVH + HL = HL(V+1)$, with the typical condition that $H \ll V$.

Prior Distribution on Coefficients

Given that one of the primary focuses of inference is the identification of influential nodes within $\mathcal{N}$ for predicting the outcome, we employ a variable selection framework on the node-specific latent effects. Node $\mathcal{N}^{(v)}$ has no effect from any of the layers on the outcome if all coefficients corresponding to the edges connected to the node $\mathcal{N}^{(v)}$ are uninformative about the outcome, i.e., $\gamma_{v,v'}^{(\alpha)} = 0$ for all $v' \neq v$ and $\alpha = 1, \dots, L$. Given the bi-linear representation of $\gamma_{v,v'}^{(\alpha)}$, this is achieved when the v th node specific latent effects at all layers are not informative about the outcome, which leads to $\xi_v = (\xi_v^{(1)T}, \dots, \xi_v^{(L)T})^T = \mathbf{0}$. In order to infer on the influential nodes, a spike-and-slab

mixture distribution prior (Ishwaran & Rao, 2005) is assigned on ξ_v . More specifically, we set

$$\xi_v \stackrel{iid}{\sim} \eta_v N(\mathbf{0}, \mathbf{K}) + (1 - \eta_v) \Delta(\mathbf{0}), \quad \eta_v \sim \text{Ber}(\delta), \quad (2)$$

where $\Delta(\mathbf{0})$ is the Dirac-delta function at $\mathbf{0}$ and $\mathbf{K}$ is a $LH \times LH$ covariance matrix. The parameter δ corresponds to the probability of the nonzero mixture component and η_v is a binary indicator set to 0 if $\xi_v = \mathbf{0}$. Thus, the posterior distributions of η_v 's enable identification of nodes related to the outcome. To account for multiplicity in multiple variable selection, we assign $\delta \sim \text{Beta}(a, b)$, following the popular literature (Scott & Berger, 2010). The covariance matrix is assigned $\mathbf{K} \sim IW(v, \mathbf{I}_{LH})$ where $IW(v, \mathbf{I}_{LH})$ denotes an Inverse-Wishart distribution with a $LH \times LH$ identity matrix $\mathbf{I}_{LH}$ and degrees of freedom v .

To learn which components of $\xi_v^{(\alpha)}$ are informative, we assign a hierarchical prior on $\theta_h^{(\alpha)}$ given by,

$$\begin{aligned} \theta_h^{(\alpha)} &= \Delta(1)\pi_{h,1}^{(\alpha)} + \Delta(0)\pi_{h,2}^{(\alpha)} + \Delta(-1)\pi_{h,3}^{(\alpha)}, \\ \alpha &= 1, \dots, L, \quad h = 1, \dots, H, \\ (\pi_{h,1}^{(\alpha)}, \pi_{h,2}^{(\alpha)}, \pi_{h,3}^{(\alpha)}) &\sim \text{Dir}(1, h^\zeta, 1), \quad \zeta > 1. \end{aligned}$$

The hyper-parameters of the Dirichlet distribution are chosen to introduce increasing shrinkage on $\theta_h^{(\alpha)}$ as h grows to prevent overfitting due to choice of a large H . Specifically, $\hat{H}^{(\alpha)} = \sum_{h=1}^H |\theta_h^{(\alpha)}|$ estimate the dimensions of $\xi_v^{(\alpha)}$ needed for effective modeling, also referred to as the *effective dimension* for $\xi_v^{(\alpha)}$. The prior specification is completed by setting $\beta_x \sim N(\mathbf{0}, \mathbf{I}_p)$ and setting a non-informative prior on μ .

Posterior Inference

Let $\mathbf{w}^{(\alpha)}$ and $\boldsymbol{\gamma}^{(\alpha)}$ denote the upper triangular vectors of dimension $V(V-1)/2 \times 1$ extracted from $\mathbf{W}^{(\alpha)}$ and $\mathbf{\Gamma}^{(\alpha)}$, respectively. Considering that both $\mathbf{w}^{(\alpha)}$ and $\boldsymbol{\gamma}^{(\alpha)}$ are symmetric matrices with zero diagonal entries, the Frobenius inner product between $\mathbf{W}^{(\alpha)}$ and $\mathbf{\Gamma}^{(\alpha)}$ can be simplified as $\langle \mathbf{W}^{(\alpha)} | \mathbf{\Gamma}^{(\alpha)} \rangle = 2\mathbf{w}^{(\alpha)T} \boldsymbol{\gamma}^{(\alpha)}$. This simplification enables us to establish a connection between the modeling framework with a multiplex graph predictor in (1) and a binary regression framework, as represented by the following equation,

$$P(y = 1 | \{\mathbf{w}^{(\alpha)}\}_{\alpha=1}^L, \mathbf{x}) = G^{-1} \left(\mu + \mathbf{x}^T \boldsymbol{\beta}_x + \sum_{\alpha=1}^L \mathbf{w}^{(\alpha)T} \boldsymbol{\gamma}^{(\alpha)} \right). \quad (3)$$

This article concentrates on employing the logit link function for $G(\cdot)$ to implement Markov chain Monte Carlo (MCMC) sampler for posterior inference on parameters.

We utilize the result discussed in Theorem 1 of Polson et al. (2013) to obtain

$$\begin{aligned}
 p(y|\mathbf{x}, \{\mathbf{w}^{(\alpha)}\}_{\alpha=1}^L) &= \frac{\exp\left\{y(\mu + \mathbf{x}^T \boldsymbol{\beta}_x + \sum_{\alpha=1}^L \mathbf{w}^{(\alpha)T} \boldsymbol{\gamma}^{(\alpha)})\right\}}{1 + \exp\left\{\mu + \mathbf{x}^T \boldsymbol{\beta}_x + \sum_{\alpha=1}^L \mathbf{w}^{(\alpha)T} \boldsymbol{\gamma}^{(\alpha)}\right\}} \\
 &= \exp\left\{(y - 0.5)\left(\mu + \mathbf{x}^T \boldsymbol{\beta}_x + \sum_{\alpha=1}^L \mathbf{w}^{(\alpha)T} \boldsymbol{\gamma}^{(\alpha)}\right)\right\} \\
 &\quad \times \int \exp\left[-\frac{\omega}{2}\left(\mu + \mathbf{x}^T \boldsymbol{\beta}_x + \sum_{\alpha=1}^L \mathbf{w}^{(\alpha)T} \boldsymbol{\gamma}^{(\alpha)}\right)^2\right] p(\omega) d\omega,
 \end{aligned} \quad (4)$$

where $p(\omega)$ is the density of PG(1,0) distribution. With the outcome, ordinary vector predictor and multiplex graph predictor obtained for n subjects, $\{y_i, \mathbf{x}_i, \mathbf{W}_i\} : i = 1, \dots, n$, we use (4) and the data augmentation approach outlined in Polson et al. (2013) for effective posterior computation. For subjects $i = 1, \dots, n$, we introduce latent variables $\omega_1, \dots, \omega_n$ in the likelihood. More precisely, the augmented likelihood function comes in the following form,

$$\begin{aligned}
 p(\mathbf{y} | \boldsymbol{\beta}_x, \boldsymbol{\Gamma}^{(1)}, \dots, \boldsymbol{\Gamma}^{(L)}, \boldsymbol{\omega}) &\propto \prod_{i=1}^n p(y_i | \boldsymbol{\beta}_x, \boldsymbol{\Gamma}^{(1)}, \dots, \boldsymbol{\Gamma}^{(L)}, \omega_i) \\
 &\propto \prod_{i=1}^n \exp\left\{(y_i - 0.5)\left(\mu + \mathbf{x}_i^T \boldsymbol{\beta}_x + \sum_{\alpha=1}^L \mathbf{w}_i^{(\alpha)T} \boldsymbol{\gamma}^{(\alpha)}\right) - \frac{\omega_i}{2}\left(\mu + \mathbf{x}_i^T \boldsymbol{\beta}_x + \sum_{\alpha=1}^L \mathbf{w}_i^{(\alpha)T} \boldsymbol{\gamma}^{(\alpha)}\right)^2\right\} \\
 &\propto \prod_{i=1}^n \exp\left\{-\frac{\omega_i}{2}\left[\frac{(y_i - 0.5)}{\omega_i} - (\mu + \mathbf{x}_i^T \boldsymbol{\beta}_x + \sum_{\alpha=1}^L \mathbf{w}_i^{(\alpha)T} \boldsymbol{\gamma}^{(\alpha)})\right]^2\right\}.
 \end{aligned}$$

While the original conditional posterior distributions for the parameters are not available in closed forms, the augmented full conditional distributions belong to standard families. Appendix provides details of the full conditional distributions of the parameters.

Let $\eta_{v,1}, \dots, \eta_{v,S}$ be the S post burn-in MCMC samples corresponding to η_v . The posterior probability of node $\mathcal{N}^{(v)}$ being significantly associated with the outcome can be empirically estimated as $P(\widehat{\eta_v = 1} | \mathbf{y}) = \frac{1}{S} \sum_{s=1}^S \eta_{v,s}$. In what follows, we cluster logit-transformed posterior probabilities into three groups using a Bayesian mixture model (Fraleigh et al., 2012).

The nodes in the cluster with the highest and lowest cluster means are considered as influential and uninfluential with respect to the outcome, respectively. We classify these nodes as “activated” and “inactivated” nodes, respectively. The analysis is less conclusive about the nodes in the third cluster, which are referred to as “inconclusive.” These estimated probabilities also provide insight into the uncertainty surrounding node inference, implying that nodes with posterior probabilities of being influential close to 0 or 1 are associated with less uncertainty in inference. Estimating uncertainty is vital in neuroscientific settings, particularly in presence of moderate sample sizes (n) and low signal-to-noise ratios. Moreover, the proposed framework is utilized for the classification of an unobserved binary outcome.

Let $\boldsymbol{\Gamma}_1^{(\alpha)}, \dots, \boldsymbol{\Gamma}_S^{(\alpha)}, \boldsymbol{\beta}_{x,1}, \dots, \boldsymbol{\beta}_{x,S}$ and $\mu_1, \dots, \mu_S$ be the S post burn-in MCMC samples for $\boldsymbol{\Gamma}^{(\alpha)}, \boldsymbol{\beta}_x$ and μ , respectively, obtained after suitable burn-in and thinning. To classify an unobserved outcome $y_* \in \{0, 1\}$ associated with a newly observed L -layer multiplex graph predictor $\mathbf{W}_*$, and with predictor $\mathbf{x}_* \in \mathbb{R}^p$, we compute the estimated classification probability given a post burn-in iterate $\{\mu_s, \boldsymbol{\beta}_{x,s}, \boldsymbol{\Gamma}_s^{(1)}, \dots, \boldsymbol{\Gamma}_s^{(L)}\}$, denoted by $P(y_* = 1 | \mu_s, \boldsymbol{\beta}_{x,s}, \boldsymbol{\Gamma}_s^{(\alpha)}, \mathbf{W}_*, \mathbf{x}_*)$, from (1) using $G(\cdot)$ as the logit link function, for $s = 1, \dots, S$. The unobserved outcome y_* is predicted to be 1 or 0 if $\frac{1}{S} \sum_{s=1}^S P(y_* = 1 | \mu_s, \boldsymbol{\beta}_{x,s}, \boldsymbol{\Gamma}_s^{(\alpha)}, \mathbf{W}_*, \mathbf{x}_*)$ is greater than or less than a selected cut-off t_c , respectively. To judge sensitivity to the choice of the cut-off, we compute ROC curves constructed with True Positive Rates (TPR) and False Positive Rates (FPR) of classification corresponding to a range of cut-off values t_c . We present area under the ROC curve (AUC) which ranges between 0.5 to 1. A value of 0.5 signifies random classification, while 1 signifies perfect classification. A higher AUC denotes superior classification performance by a method.

Results

In this section, we apply the proposed *Bayesian Multiplex Graph Classifier (BMGC)* method to perform inference on nodes and classification on multiplex graphs using both simulated data and real biomedical imaging data (brain fMRI study).

Simulation Study

First, we investigate the performance of the BMGC technique through simulation and compare it against other classification methods. Our evaluation involves examining BMGC’s capability to identify influential nodes and then comparing its performance in coefficient estimation and predictive inference with competing models. These

competitors encompass penalized optimization, Bayesian shrinkage methods, Bayesian tensor regression and deep neural networks. In each simulation scenario, the number of layers (L) is fixed at 2, mirroring the number of layers in the neuroimaging dataset under evaluation.

Simulated Data Generation

For the i th sample, $i = 1, \dots, n$, we generate a multiplex graph $\mathbf{W}_i$ by simulating symmetric adjacency matrices $\mathbf{W}_i^{(1)}, \dots, \mathbf{W}_i^{(L)}$ for L layers. The upper triangular vector $\mathbf{w}_i^{(\alpha)}$ of $\mathbf{W}_i^{(\alpha)}$ is independently drawn from the normal distribution $N(\mathbf{0}, \mathbf{I})$, and all diagonal entries of $\mathbf{W}_i^{(\alpha)}$ are set to 0. Let the true intercept and the true multiplex graph coefficient be given by μ_0 and $\mathbf{\Gamma}_0$, respectively. The response y_i for the i th sample is generated according to model (1) with $G(\cdot)$ as the logit link function. No predictor $\mathbf{x}_i$ is included in these simulations. In all scenarios, we set $\mu_0 = 0.5$. We use $n = 400$ samples for model fitting and reserve $n_* = 50$ samples for prediction.

Let δ_0 be the probability of a node being influential w.r.t. the outcome. We refer to $(1 - \delta_0)$ as the node sparsity. For the purpose of generating the true activation pattern of nodes in relation to the outcome, we simulate node-specific activation indicators $\eta_{1,0}, \dots, \eta_{V,0} \stackrel{i.i.d.}{\sim} \text{Ber}(\delta_0)$. To simulate the α th layer of the multiplex graph coefficient $\mathbf{\Gamma}_0^{(\alpha)}$, we first simulate the LH_0 -dimensional latent effects corresponding to each of the V nodes, $\xi_{v,0} = (\xi_{v,0}^{(1)}, \dots, \xi_{v,0}^{(L)})^T$, from $\xi_{v,0} \stackrel{ind.}{\sim} \eta_{v,0} N(\mathbf{0}, \mathbf{I}_{LH_0}) + (1 - \eta_{v,0}) \Delta(\mathbf{0})$, for $v = 1, \dots, V$. Using these latent effects for nodes, the true coefficients corresponding to the (v, v') th edge at the α th layer is constructed using the low-rank formulation $\gamma_{v,v',0}^{(\alpha)} = \xi_{v,0}^{(\alpha)T} \xi_{v',0}^{(\alpha)}$, for $v < v'$. We set $\gamma_{v,v',0}^{(\alpha)} = \gamma_{v',v,0}^{(\alpha)}$ and $\gamma_{v,v,0}^{(\alpha)} = 0$ to satisfy the symmetry and zero diagonals the graph coefficient at each layer, respectively. In using this setup, we assume that both the fitted and true graph coefficients have a low-rank structure. This assumption allows us to evaluate how well the proposed method captures and represents the underlying relationships between the graph nodes and the outcome variable. With this setup, we simulate data with $V = 20$, at five different levels of node sparsity, setting $\delta_0 = 0.2, 0.3, 0.4, 0.5, 0.6$, respectively.

Competing Methods and Performance Metrics

We compare the performance of the proposed BMGC model vis-a-vis two sets of competitors. The first set of competitors treats the edges between nodes in the multiplex graph predictor as a “long vector of predictors” and perform binary regression of the outcome y_i on the vector $(\mathbf{w}_i^{(\alpha)T} : \alpha = 1, \dots, L)^T$ of dimension $LV(V - 1)/2$. This approach overlooks the relational nature of the multiplex

graph predictor, potentially limiting the ability of such models to capture the effect of intricate interconnections between nodes on the outcome. To this end, we employ the horseshoe prior (Carvalho et al., 2010) and the Bayesian Lasso (Park & Casella, 2008) on the regression coefficients due to their state-of-the-art empirical performance in regressions with a high-dimensional predictor vector. We refer to these competitors as BHS and BLasso, respectively, and implement them using our own codes. As a frequentist high-dimensional binary regression competitor, we adopt a penalized optimization framework with the Lasso penalty on the predictor coefficients (Tibshirani, 1996). The Lasso is implemented using the `glmnet` (Friedman et al., 2010) package in R, with the penalty parameter chosen through ten-fold cross-validation. Finally, we also implement a two-layer neural net (NN) as a competitor. NN models are fit using the `tensorflow` (Allaire & Chollet, 2023) and `keras` (Allaire & Tang, 2023) packages in R. The inputs to the NN models are composed of the upper-triangular entries of each graph layer, which are concatenated to create an input vector of dimension $LV(V - 1)/2$. The models are fit with two hidden layers, each having the Rectified Linear Unit (ReLU) activation function (Nair & Hinton, 2010). The first of these two hidden layers includes L_1 and L_2 regularization to perform shrinkage on the coefficients from the input layer. The output layer of each model uses the sigmoid activation function (Narayan, 1997). The L_1 and L_2 regularization parameters and learning rate for the models are tuned through five-fold cross validation. For training each NN, 90 percent of the training samples are used to fit the model, while 10 percent are used for validation. The batch size is equal to half of the number of samples used to fit the model. The second set of competitors treats every layer of the multiplex graph as a $V \times V$ tensor and uses a binary tensor regression approach with L tensor predictors and a binary response. We employ the framework in Guhaniyogi et al. (2017) to estimate the tensor coefficients. Unlike the first set of competitors, this approach envisions the matrix structure of each graph layer. However, it does not take into account the symmetry of each undirected graph layer. We refer to this approach as TensorReg.

To infer on the performance of BMGC in terms of identifying nodes significantly associated with the outcome, we present the estimated posterior probability of a node being influential across the layers, i.e., $P(\eta_v = 1 | \text{Data})$. The performance of the competitors in estimating multiplex graph coefficient $\mathbf{\Gamma}$ is evaluated using the scaled mean squared error (MSE). The MSE is defined as $\frac{\|\hat{\mathbf{\Gamma}} - \mathbf{\Gamma}_0\|^2}{\|\mathbf{\Gamma}_0\|^2}$, where $\hat{\mathbf{\Gamma}}$ represents a suitable point estimate of $\mathbf{\Gamma}$. For instance, we use the frequentist point estimate for Lasso, while BMGC, BHS, TensorReg and BLasso employ the posterior mean of $\mathbf{\Gamma}$. To assess predictive power of the competitors, we use the held-out test samples to compute the AUC of the predicted

Table 1 Node selection results across five simulation scenarios, each depicted in a separate column, with varying probabilities of node inclusion (δ_0). Every column displays the posterior probability of the $V = 20$ nodes in the simulation setup being active, corresponding

Node	$\delta_0 = 0.2$	$\delta_0 = 0.3$	$\delta_0 = 0.4$	$\delta_0 = 0.5$	$\delta_0 = 0.6$
1	0.9800	0.8737	0.8553	0.7837	0.6230
2	0.1003	0.3243	0.1480	0.2700	0.2357
3	0.9807	0.9283	0.8143	0.8997	0.7960
4	0.0343	0.0617	0.1280	0.2030	0.2227
5	0.0380	0.9307	0.5587	0.7590	0.8277
6	0.0250	0.1157	0.0740	0.2440	0.2267
7	0.0653	0.0330	0.0757	0.2343	0.2047
8	0.9770	0.8007	0.9263	0.6573	0.6943
9	0.0083	0.0400	0.3840	0.7753	0.3637
10	0.0090	0.8343	0.9260	0.6807	0.6743
11	0.0087	0.0937	0.8603	0.7623	0.5460
12	0.9697	0.8960	0.8763	0.9057	0.7413
13	0.0107	0.0290	0.0307	0.0950	0.0933
14	0.0047	0.0113	0.0487	0.1620	0.8287
15	0.0060	0.0170	0.0367	0.1313	0.8203
16	0.0110	0.0307	0.0223	0.8750	0.6587
17	0.0143	0.0173	0.1127	0.5250	0.7757
18	0.0077	0.0567	0.0477	0.1003	0.0993
19	0.0073	0.0580	0.1203	0.1797	0.1130
20	0.0073	0.0213	0.0720	0.0683	0.1397

probabilities, $P(y_* = 1 | W_*)$, of each model compared against the ground truth labels, $y_{*,1}, \dots, y_{*,n_*}$. For the Lasso, this quantity corresponds to the point estimate from the fitted model, while for the Bayesian competitors, this is computed using the mean of the posterior predictive distribution. All comparisons presented are averaged over five simulation replicates.

Simulation Results

Accuracy of Influential Node Identification Table 1 displays a matrix corresponding to five different simulation settings under different true node sparsity ($1 - \delta_0$) levels for $V = 20$. In each matrix, the colored and white cells in the l th column represent the truly influential and truly uninfluential nodes for the l th simulation scenario, respectively, for $l = 1, \dots, 5$. Overlaid on these matrices are the estimated posterior probabilities $P(\eta_v = 1 | Data)$ for $v = 1, \dots, 20$.

Table 1 shows remarkably accurate identification of truly influential nodes, with the estimated posterior probabilities of all influential nodes in most simulations being above 0.5 in all five simulation scenarios. With high-node sparsity ($\delta_0 = 0.2$), the estimated posterior probabilities tend to be very close to 1 or 0, correctly classifying nodes as influential or not, with very minimal uncertainty. However, as sparsity decreases, the uncertainty in node identification deteriorates. In fact, when $\delta_0 = 0.6$, the estimated posterior probabilities of a few influential nodes becomes close to 0.5 and a few un-influential nodes becomes close to 0.25, displaying more uncertainty in node identification.

with various levels of node-sparsity in the truth. White and colored cells denote uninfluential and influential nodes, respectively. BMGC excels in identifying nodes, as indicated by high probabilities in colored cells and low probabilities in white cells

It is crucial to underscore that one of BMGC's key strengths in inference lies in its capability to identify significant nodes in multiplex graphs, considering the interactions across layers while quantifying uncertainty. This feature directly aligns with the inferential goals of our scientific study, which revolves around identification of brain regions of interest associated with accelerated aging. In contrast, prevalent frequentist or Bayesian high-dimensional regression techniques, tensor regression approach or neural networks lack the capability for node identification with uncertainty in the current context. For the assessment of point estimation in multiplex graph coefficients, Table 2 showcases the scaled Mean Squared Error (MSE) for each of the competitors across the five simulation scenarios. In all instances, the results consistently highlight the superior performance of BMGC. The performance gap between BMGC

Table 2 Scaled mean squared error (MSE) of estimating multiplex graph coefficients for all competitors under all five simulation scenarios. Deep neural network (NN) only offers predictive inference, hence the results from NN are not shown in the table. The result corresponding to the best performer under all simulation scenarios is boldfaced

Method	$\delta_0 = 0.2$	$\delta_0 = 0.3$	$\delta_0 = 0.4$	$\delta_0 = 0.5$	$\delta_0 = 0.6$
BMGC	0.9207	0.9693	0.9859	0.9910	0.9946
Lasso	1.2514	1.0781	1.0306	1.0150	1.0082
BHS	2.6233	1.6423	1.5982	1.8199	1.8967
BLasso	1.5864	1.2707	2.5810	1.5385	1.3124
TensorReg	45.4238	10.9493	11.4556	7.3167	4.4012




Fig. 1 Comparison of the performance of the BMGC method (blue) against alternative classification techniques (different color lines) using the area under the ROC curve (AUC) across different levels of the sparsity parameter δ_0

and its competitors is substantial with high-node density (e.g., smaller values of δ_0). As node sparsity decreases, the performance of BMGC deteriorates and it becomes close to Lasso when $\delta_0 = 0.6$.

To assess the out-of-sample predictive performance of the different models, Fig. 1 displays the area under the receiver operating characteristic curves (AUCs) obtained using various classification thresholds t_c (as discussed in “Posterior




Fig. 2 Comparison of the performance of the BMGC method (blue) against alternative classification techniques (different color lines) using the F1-Score across different levels of the sparsity parameter δ_0

Inference” section). We also show the F1-scores for all the competitors in Fig. 2 using $t_c = 0.5$ as the cut-off. Both figures draw similar conclusions regarding comparative performance of the competitors. With high-node sparsity (small values of δ_0), the graph structure is weak, causing all competitors to essentially exhibit behavior similar to that of penalized regression. Thus, with high-node sparsity, all competitors show similar out-of-sample classification performance. As sparsity decreases, classification performance of all competitors deteriorates, but BMGC consistently maintains an advantage over the others. This indicates that leveraging the multiplex graph structure of the predictor in BMGC can improve classification accuracy. Among the methods that disregard the graph structure associated with the predictor, NN exhibits inferior performance.

Functional Magnetic Resonance Imaging (fMRI) Study

Next, we will apply the BMGC technique to the functional magnetic resonance imaging (fMRI) brain dataset. Figure 3 shows a heatmap of the initial correlations between the ROI nodes in the left (columns) and right (rows) hemispheres.

Figure 4 illustrates a partial brain connectogram indicating specific associations between some of the left hemisphere brain areas. Clearly some brain regions are highly connected, whereas other are more isolated. For simplicity, we only

show a fragment of connectogram reflecting the following 10 left hemisphere (LH) *central* and *peripheral* (striatal and extra-striatal) visual cortex brain regions, see Table 3.

Figure 5 shows a scatter plot of the *derived tripartite class labels*, (“inactivated”, “inconclusive”, or “activated”) computed by our algorithm. The glyph *colors* and *shapes* of scatter points represent the respective probabilities of the left (x-axis) and right (y-axis) regions of interest (i.e., nodes in the brain graph network). The point 2D Cartesian coordinates represent the corresponding probabilities that the left (x-axis) and right (y-axis) brain hemispherical regions are in one of the three derived class labels. Points along near line bisecting the plane first-quadrant correspond to ROIs with identical derived class labels for the left and right region. Hence, these ROIs correspond to *brain symmetries*, i.e., bilaterally synchronously associations with memory decline (according to the MMSE outcome binary label). Conversely, scatter points off the main diagonal correspond to *brain asymmetries*; left and right brain regions whose derived labels are inconsistent.

Some of these brain asymmetries have been previously reported in normal aging, mild cognitive impairment, and dementia (Lubben et al. 2021; Mandal et al. 2012; Thompson et al. 1998). For instance, we observed brain symmetries involving networks with left and right *somatosensory motor area* (ROI: SomMotA.3), where both left and right probabilities are high. This indicates that in aging, sensory motor brain function obeys certain lateral




Fig. 3 Heatmap of the correlations between the ROI nodes in the left and right hemispheres




Fig. 4 Connectogram showing partial correlations between several different brain regions in the left hemisphere

symmetries. On the other hand, in the *default mode network* (ROI: DefaultC.Rsp.1) the difference in the derived class labels for the left (inactivated) and right (activated) suggest significant brain asymmetries associated with pathological aging.

In this fMRI study, to evaluate the BMGC classification performance against alternative methods, we computed the probability of the outcome being 1 (normal aging) for each sample (inputs are the individual correlation matrices of participants' brain networks). This allowed us to construct

ROC curves based on pairs of True Positive Rate (TPR) and False Positive Rate (FPR) values, obtained by varying the threshold of these estimated probabilities. The AUC estimated for each competitor corresponds to the area under their respective ROC curve. We also compute F1-scores for all the competitors considering 0.5 as the threshold probability. Table 4 presents a comparison of the classification performance between the new BMGC technique and various alternatives (described in “[Competing Methods and Performance Metrics](#)” section).

Table 3 The 10 left-hemisphere (LH) central and (striatal and extrastriatal) peripheral visual cortex brain regions in Fig. 4

Region Acronym	Region Name
LH.VisCent.ExStr.1.VisCent	Left-Hemisphere Extrastriatal Central Visual Cortex Region 1
LH.VisCent.ExStr.2.VisCent	Left-Hemisphere Extrastriatal Central Visual Cortex Region 2
LH.VisCent.Striate.1.VisCent	Left-Hemisphere Striatal Central Visual Cortex Region 1
LH.VisCent.ExStr.3.VisCent	Left-Hemisphere Extrastriatal Central Visual Cortex Region 3
LH.VisCent.ExStr.4.VisCent	Left-Hemisphere Extrastriatal Central Visual Cortex Region 4
LH.VisCent.ExStr.5.VisCent	Left-Hemisphere Extrastriatal Central Visual Cortex Region 5
LH.VisPeri.ExStrInf.1.VisPeri	Left-Hemisphere Extrastriatal Inferior Peripheral Visual Cortex Region 1
LH.VisPeri.ExStrInf.2.VisPeri	Left-Hemisphere Extrastriatal Inferior Peripheral Visual Cortex Region 2
LH.VisPeri.ExStrInf.3.VisPeri	Left-Hemisphere Extrastriatal Inferior Peripheral Visual Cortex Region 3
LH.VisPeri.StriCal.1.VisPeri	Left-Hemisphere Striatal Caudal Peripheral Visual Cortex Region 1

BMGC demonstrates superior performance compared to all other competitors, likely because it leverages the multiplex graph architecture within the modeling framework. Although Neural Networks (NN) enable capturing complex non-linear relationships between the outcome and the multiplex graph, Bayesian Horseshoe (BHS) and Bayesian Lasso (BLasso) exhibit slightly improved performance over NN, indicating the presence of sparsity in the data structure. TensorReg also performs close to BLasso and BHS, perhaps for not exploiting the symmetry in the graph layers. While the

AUC values show BMGC offering more than 10% improvement in performance, the F1-scores suggest slight outperformance of BMGC over Lasso, BHS and Blasso. Digging a bit deeper, we identify that the probabilities for an outcome being “1” or “0” lie closely around 0.5 for Lasso, BHS and Blasso. However, these probabilities are more separated from 0.5 for BMGC. Thus the probability of an outcome being 1 is much higher than 0.5, or the probability of an outcome being 0 is much lower than 0.5 for BMGC in the case of a correctly classified outcome. This explains the




Fig. 5 Bivariate plot of the derived tripartite class labels computed by the multilayer network method. The glyph colors and shapes of each scatter point represent the regions of interest (nodes in the brain graph network). The point locations represent the corresponding

probabilities that the left (x-axis) and right (y-axis) brain hemispherical regions are in one of the three derived class labels (“inactivated”, “inconclusive”, or “activated”)

Table 4 The predictive performance, as measured by the AUC and F1-score, of BMGC and its competitors on the fMRI data. The best performer is highlighted in bold

Model	BMGC	Lasso	BHS	BLasso	NN	TensorReg
AUC	0.6623	0.5315	0.6035	0.5993	0.5708	0.5833
F1-Score	0.8382	0.7881	0.8033	0.8030	0.6968	0.7084

competitors' lower AUC values, but competitive F1-scores. It is worth mentioning that none of the competitors have the capability for node selection, a feature exclusive to BMGC.

Conclusion and Future Work

This article introduces an innovative Bayesian framework designed to address a classification problem characterized by a binary outcome and a multiplex graph predictor, with each layer representing an undirected graph. We model an edge coefficient in a specific layer using a bilinear interaction of latent effects corresponding to nodes connected by that particular edge. A group variable selection framework is employed on node-specific latent effects for all layers to identify nodes influential in the classification of the binary outcome. A Bayesian framework is adopted to allow uncertainty quantification in inference and prediction, especially in the identification of influential nodes. Our empirical results, derived from simulation studies, establish the superiority of our method in scenarios where the regression coefficients indeed exhibit a multiplex graph structure.

Our functional magnetic resonance imaging (fMRI) study suggests that the proposed BMGC technique may be useful to identify brain-network asymmetries involving left and right cortical regions according to their BMGC posterior predictive class probabilities. For each brain area, these likelihoods yield explicit derived class labels of association with any (binary, categorical, or continuous) clinical outcome of interest, which can be used to identify, confirm, or explore various pathological states.

Our framework stands at the forefront of developing a regression model with a binary outcome and a multiplex graph predictor. It is noteworthy that our framework assumes a bi-linear effect of graph edges on the binary outcome. While bi-linear edge effects have been demonstrated to effectively capture intricate graph properties (Hoff et al., 2002), an approach that captures complex non-linear associations between edges of the multiplex graph and the outcome could potentially yield deeper scientific insights into the impact of aging on functional connectivity. To address this, we will incorporate a broader class of semi-parametric models to capture such complex non-linear associations in our future work. Our contribution will involve advancing the semi-parametric regression framework by introducing multiplex graph

predictors. This builds upon existing literature (e.g., see Zhou et al., 2024), where a single graph predictor is employed.

Future studies may also take a further step in examining the posterior probability distributions, specifically for brain asymmetry related to differences between the left and right hemispheres. Since the brain networks are the inputs, not the outputs of the BMGS model for classifying ROIs as inactive, borderline, or actively participating in the binary classification (output) of participants into normal and accelerated aging. This can facilitate drawing inference about connectivity pattern different cohorts does not appear feasible. For instance, it may be interesting to explore bilaterally the ROIs and brain networks (Schaefer-17) that are most probabilistically distinct between the two cohorts of normal (MMSE = 30) and aging (MMSE ≤ 30).

Information Sharing Statement

The complete fMRI dataset and the regional brain correlation functional connectivity networks reflecting cognitive control throughout the adult lifespan are available online (<https://osf.io/5zyws/>) Rieck et al. (2021a). As firm supporters of open, reproducible and trustworthy scientific discovery, we have provided data, software tools, protocols, interactive charts, and supplementary materials on these websites: https://socr.umich.edu/docs/uploads/2024/fMRI_Corr_Pilot.html and <https://github.com/jeroda7105/Classification-with-Multi-Layer-Graphs>.

Appendix

Below we present the full conditional posterior distributions for all parameters within the generalized linear model detailed in “[Model and Prior Formulation](#)” section. These distributions are utilized for implementing a Markov Chain Monte Carlo (MCMC) algorithm via Gibbs sampling. The samples derived from the MCMC algorithm yield samples from the complete posterior distributions for the model parameters jointly, facilitating posterior inference as elaborated in “[Posterior Inference](#)” section. Let $\kappa_i = (y_i - 0.5)/\omega_i$, and $\Omega = \text{diag}\left(\frac{1}{\omega_1}, \dots, \frac{1}{\omega_n}\right)$. Assuming X to be an $n \times p$ matrix with its i th row as $\mathbf{x}_i$, and $\mathbf{A}^{(a)}$ to be an $n \times V(V-1)/2$

dimensional matrix with its i th row given by $\mathbf{w}_i^{(\alpha)}$. Consequently, the full conditional distributions for the model parameters are expressed as:

$$\bullet \mu | - \sim N\left(\frac{\mathbf{1}_n^T \Omega^{-1} (\kappa - X\beta_x - \sum_{a=1}^L A^{(a)} \gamma^{(a)})}{\mathbf{1}_n^T \Omega^{-1} \mathbf{1}_n}, \frac{1}{\mathbf{1}_n^T \Omega^{-1} \mathbf{1}_n}\right), \quad \text{where}$$

$\kappa = (\kappa_1, \dots, \kappa_n)^T$ is an n -dimensional vector of continuous outcomes over all samples.

$$\bullet \beta_x | - \sim N(\mu_\beta, \Sigma_\beta), \text{ where}$$

$$\Sigma_\beta = (X^T \Omega^{-1} X + I_p)^{-1}, \quad \mu_\beta = \Sigma_\beta X^T \Omega^{-1} (\kappa - \mu \mathbf{1}_n - \sum_{a=1}^L A^{(a)} \gamma^{(a)}).$$

$$\bullet \xi_v | - \sim \eta_v N(\mu_v, \Sigma_v) + (1 - \eta_v) \Delta(\mathbf{0}), \text{ where}$$

$$\Sigma_v = (Z_v^T \Omega^{-1} Z_v + K^{-1})^{-1}, \quad \mu_v = \Sigma_v Z_v^T \Omega^{-1} \tilde{\kappa}_v.$$

Here Z_v is an $n \times LH$ matrix whose i th row is given by $(\sum_{k < v} w_{i,(k,v)}^{(1)} \xi_k^{(1)T} \Theta^{(1)} + \sum_{v < k} w_{i,(v,k)}^{(1)} \xi_k^{(1)T} \Theta^{(1)}, \dots, \sum_{k < v} w_{i,(k,v)}^{(L)} \xi_k^{(L)T} \Theta^{(L)} + \sum_{v < k} w_{i,(v,k)}^{(L)} \xi_k^{(L)T} \Theta^{(L)})$ and $\tilde{\kappa}_v$ is an n dimensional vector with its i th entry

$$\tilde{\kappa}_i = \kappa_i - \mu - X\beta_x - \sum_{k < k': k, k' \neq v} \sum_{a=1}^L w_{i,(k,k')}^{(\alpha)} \xi_k^{(\alpha)T} \Theta^{(\alpha)} \xi_{k'}^{(\alpha)}.$$

$$\bullet \eta_v | - \sim \text{Ber}(\tilde{\eta}_v), \text{ where}$$

$$\tilde{\eta}_v = \frac{\delta N(\tilde{\kappa}_v | \mathbf{0}, Z_v K Z_v^T + \Omega)}{\delta N(\tilde{\kappa}_v | \mathbf{0}, Z_v K Z_v^T + \Omega) + (1 - \delta) N(\tilde{\kappa}_v | \mathbf{0}, \Omega)}.$$

$$\bullet \delta | - \sim \text{Beta}(a + \sum_{v=1}^V \eta_v, b + V - \sum_{v=1}^V \eta_v).$$

$$\bullet K | - \sim IW\left(v + \#\{v : \eta_v = 1\}, I_{LH} + \sum_{\{v: \eta_v = 1\}} \mathbf{u}_v \mathbf{u}_v^T\right),$$

where $\mathbf{u}_v = (\mathbf{u}_v^{(1)T}, \dots, \mathbf{u}_v^{(L)T})^T$.

$$\theta_h^{(\alpha)} | - \sim \begin{cases} 1 & w.p. p_{h,1}^{(\alpha)} \\ 0 & w.p. p_{h,2}^{(\alpha)} \\ -1 & w.p. p_{h,3}^{(\alpha)} \end{cases}$$

$$\bullet \text{ where } p_{h,1}^{(\alpha)} = \frac{\pi_{h,1}^{(\alpha)} N(\kappa | \mu \mathbf{1}_n + X\beta_x + A^{(\alpha)} (\gamma^{(\alpha)})^{(\theta_h^{(\alpha)}=1)} + \sum_{a' \neq \alpha} A^{(a')} (\gamma^{(a')}), \Omega)}{S},$$

$$p_{h,2}^{(\alpha)} = \frac{\pi_{h,2}^{(\alpha)} N(\kappa | \mu \mathbf{1}_n + X\beta_x + A^{(\alpha)} (\gamma^{(\alpha)})^{(\theta_h^{(\alpha)}=0)} + \sum_{a' \neq \alpha} A^{(a')} (\gamma^{(a')}), \Omega)}{S},$$

$$p_{h,3}^{(\alpha)} = \frac{\pi_{h,3}^{(\alpha)} N(\kappa | \mu \mathbf{1}_n + X\beta_x + A^{(\alpha)} (\gamma^{(\alpha)})^{(\theta_h^{(\alpha)}=-1)} + \sum_{a' \neq \alpha} A^{(a')} (\gamma^{(a')}), \Omega)}{S},$$

$$\text{where, } S = \sum_{s \in \{0,1,-1\}} \pi_{h,s}^{(\alpha)} N(\kappa | \mu \mathbf{1}_n + X\beta_x + A^{(\alpha)} (\gamma^{(\alpha)})^{(\theta_h^{(\alpha)}=s)} + \sum_{a' \neq \alpha} A^{(a')} (\gamma^{(a')}), \Omega)$$

$$\bullet \omega_i | - \sim PG(1, \mu + x_i^T \beta_x + \sum_{a=1}^L \mathbf{w}_i^{(\alpha)T} \gamma^{(\alpha)}), \text{ for } i = 1, \dots, n.$$

Author Contributions S.G. conceived the idea, formulated the model, calculated the priors for parameters and full conditionals for MCMC computation. S.G. took the lead in drafting the article. J.R.A. implemented the model and drafted parts of the simulation study and the real

data analysis, generated and added simulation figures and tables. I.D.D. took the lead in drafting the real data analysis section, editing the draft, providing scientific data and valuable scientific insights, generated the real data result figures. All authors reviewed the manuscript.

Data Availability Data is provided within the supplementary information files.

Declarations

Competing Interests The authors declare no competing interests.

References

- Ahmed, T., Raja, H., & Bajwa, W. U. (2020). Tensor regression using low-rank and sparse tucker decompositions. *SIAM Journal on Mathematics of Data Science*, 2(4), 944–966. <https://doi.org/10.1137/19M1299335>
- Airoldi, E. M., Blei, D. M., Fienberg, S. E., & Xing, E. P. (2008). Mixed membership stochastic blockmodels. *Journal of Machine Learning Research*, 9(65), 1981–2014.
- Allaire, J., & Chollet, F. (2023). Keras: R Interface to ‘Keras’. R package version 2.11.1. <https://tensorflow.rstudio.com/>
- Allaire, J., & Tang, Y. (2023). Tensorflow: R Interface to ‘TensorFlow’. R package version 2.11.0.9000. <https://github.com/rstudio/tensorflow>
- Carvalho, C. M., Polson, N. G., & Scott, J. G. (2010). The horseshoe estimator for sparse signals. *Biometrika*, 97(2), 465–480.
- Chen, Y., Chen, D., Wang, Y., Wang, T., & Liang, Y. (2021). CaF-Graph: Context-aware facial multi-graph representation for facial action unit recognition. *Proceedings of the 29th ACM International Conference on Multimedia*. (pp. 1029–1037).
- Chen, F., Wang, Y.-C., Wang, B., & Kuo, C.-C.J. (2020). Graph representation learning: A survey. *APSIPA Transactions on Signal and Information Processing*, 9, 15.
- Churchill, N. W., Raamana, P., Spring, R., & Strother, S. C. (2017). Optimizing fMRI preprocessing pipelines for block-design tasks as a function of age. *NeuroImage*, 154, 240–254.
- Contisciani, M., Power, E. A., & De Bacco, C. (2020). Community detection with node attributes in multilayer networks. *Scientific Reports*. <https://doi.org/10.1038/s41598-020-72626-y>
- Craddock, R. C., Holtzheimer, P. E., Hu, X. P., & Mayberg, H. S. (2009). Disease state prediction from resting state functional connectivity. *Magnetic Resonance in Medicine*, 62(6), 1619–1628. <https://doi.org/10.1002/mrm.22159>
- Dinh, V. C., & Ho, L. S. (2020). Consistent feature selection for analytic deep neural networks. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, & H. Lin (Eds.), *Advances in Neural Information Processing Systems* (Vol. 33, pp. 2420–2431).
- Dinov, I. D., & Velev, M. V. (2021). *Data science: Time complexity, inferential uncertainty, and space-time analytics*. Walter de Gruyter GmbH & Co KG
- Dinov, I. D. (2023). Springer. *Cham*. https://doi.org/10.1007/978-3-031-17483-4_2
- Durante, D., Mukherjee, N., & Steorts, R. C. (2017). Bayesian learning of dynamic multilayer networks. *Journal of Machine Learning Research*, 18(43), 1–29.
- Fan, J., Gong, W., & Zhu, Z. (2019). Generalized high-dimensional trace regression via nuclear norm regularization. *Journal of Econometrics*, 212(1), 177–202. <https://doi.org/10.1016/j.jeconom.2019.04.026>. Big Data in Dynamic Predictive Econometric Modeling.

- Fraley, C., Raftery, A. E., Murphy, T. B., & Scrucca, L. (2012). *mclust version 4 for R: normal mixture modeling for model-based clustering, classification, and density estimation*. Citeseer: Technical report.
- Frank, O., & Strauss, D. (1986). Markov graphs. *Journal of the American Statistical Association*, 81(395), 832–842. <https://doi.org/10.1080/01621459.1986.10478342>
- Friedman, J., Tibshirani, R., & Hastie, T. (2010). Regularization paths for generalized linear models via coordinate descent. *Journal of Statistical Software*, 33(1), 1–22.
- Gollini, I., & Murphy, T. B. (2016). Joint modeling of multiple network views. *Journal of Computational and Graphical Statistics*, 25(1), 246–265.
- Guhaniyogi, R., Qamar, S., & Dunson, D. B. (2017). Bayesian tensor regression. *Journal of Machine Learning Research*, 18(79), 1–31.
- Guha, S., & Guhaniyogi, R. (2021). Bayesian generalized sparse symmetric tensor-on-vector regression. *Technometrics*, 63(2), 160–170.
- Guha, S., & Guhaniyogi, R. (2024). Covariate-dependent clustering of undirected networks with brain-imaging data. *Technometrics*, 1–23.
- Guha, S., & Rodriguez, A. (2021). Bayesian regression with undirected network predictors with an application to brain connectome data. *Journal of the American Statistical Association*, 116(534), 581–593. <https://doi.org/10.1080/01621459.2020.1772079>
- Guha, S., & Rodriguez, A. (2023). High-dimensional Bayesian network classification with network global-local shrinkage priors. *Bayesian Analysis*. <https://doi.org/10.1214/23-BA1378>
- Han, Q., Xu, K. S., & Airolidi, E. M. (2015). Consistent estimation of dynamic and multi-layer block models. *Proceedings of the 32nd International Conference on International Conference on Machine Learning* (Vol. 37).
- He, L., Chen, K., Xu, W., Zhou, J., & Wang, F. (2018). Boosted sparse and low-rank tensor regression. *Advances in Neural Information Processing Systems*, 31.
- Heaney, M. T. (2014). Multiplex networks and interest group influence reputation: An exponential random graph model. *Social Networks*, 36, 66–81. <https://doi.org/10.1016/j.socnet.2012.11.003>. Special Issue on Political Networks.
- Heidari, N., & Iosifidis, A. (2021). Progressive spatio-temporal bilinear network with Monte Carlo dropout for landmark-based facial expression recognition with uncertainty estimation. *2021 IEEE 23rd International Workshop on Multimedia Signal Processing (MMSP)* (pp. 1–6). IEEE.
- Hoff, P. D. (2015). Multilinear tensor regression for longitudinal relational data. *The Annals of Applied Statistics*. <https://doi.org/10.1214/15-aos839>
- Hoff, P. D., Raftery, A. E., & Handcock, M. S. (2002). Latent space approaches to social network analysis. *Journal of the American Statistical Association*, 97(460), 1090–1098. <https://doi.org/10.1198/016214502388618906>
- Holland, P. W., & Leinhardt, S. (1981). An exponential family of probability distributions for directed graphs. *Journal of the American Statistical Association*, 76(373), 33–50. <https://doi.org/10.1080/01621459.1981.10477598>
- Ishwaran, H., & Rao, J. S. (2005). Spike and slab variable selection: Frequentist and Bayesian strategies. *The Annals of Statistics*, 33(2), 730–773. <https://doi.org/10.1214/009053604000001147>
- Li, Y., Zhang, L., Lan, X., & Jiang, D. (2023). Towards adaptable graph representation learning: An adaptive multi-graph contrastive transformer. *Proceedings of the 31st ACM International Conference on Multimedia* (pp. 6063–6071).
- Liu, Z., & Zhou, J. (2022). *Introduction to graph neural networks*. Springer Nature.
- Lubben, N., Ensink, E., Coetzee, G. A., & Labrie, V. (2021). The enigma and implications of brain hemispheric asymmetry in neurodegenerative diseases. *Brain Communications*, 3(3), 211.
- Mandal, P. K., Mahajan, R., & Dinov, I. D. (2012). Structural brain atlases: design, rationale, and applications in normal and pathological cohorts. *Journal of Alzheimer's Disease*, 31(s3), 169–188.
- Mega, M. S., Cummings, J. L., Fiorello, T., & Gornbein, J. (1996). The spectrum of behavioral changes in Alzheimer's disease. *Neurology*, 46(1), 130–135.
- Moon, S. W., Zhao, L., Matloff, W., Hobel, S., Berger, R., Kwon, D., Kim, J., Toga, A. W., & Dinov, I. D. (2023). Disease neuroimaging initiative: Brain structure and allelic associations in Alzheimer's disease. *CNS Neuroscience & Therapeutics*, 29(4), 1034–1048. <https://doi.org/10.1111/cns.14073>. <https://onlinelibrary.wiley.com/doi/pdf/10.1111/cns.14073>
- Nair, V., & Hinton, G. E. (2010). Rectified linear units improve restricted Boltzmann machines. *ICML'10* (pp. 807–814). Madison, WI, USA: Omnipress.
- Narayan, S. (1997). The generalized sigmoid activation function: Competitive supervised learning. *Information Sciences*, 99(1), 69–82. [https://doi.org/10.1016/S0020-0255\(96\)00200-9](https://doi.org/10.1016/S0020-0255(96)00200-9)
- Ngoc, Q. T., Lee, S., & Song, B. C. (2020). Facial landmark-based emotion recognition via directed graph neural network. *Electronics*, 9(5), 764.
- Nowicki, K., & Snijders, T. A. B. (2001). Estimation and prediction for stochastic blockstructures. *Journal of the American Statistical Association*, 96(455), 1077–1087. <https://doi.org/10.1198/016214501753208735>
- Park, T., & Casella, G. (2008). The Bayesian lasso. *Journal of the American Statistical Association*, 103(482), 681–686. <https://doi.org/10.1198/016214508000000337>
- Polson, N. G., & Ročková, V. (2018). Posterior concentration for sparse deep learning. *Proceedings of the 32nd International Conference on Neural Information Processing Systems. NIPS'18* (pp. 938–949). Red Hook, NY, USA: Curran Associates Inc.
- Polson, N. G., Scott, J. G., & Windle, J. (2013). Bayesian inference for logistic models using Pólya-Gamma latent variables. *Journal of the American Statistical Association*, 108(504), 1339–1349. <https://doi.org/10.1080/01621459.2013.829001>
- Richiardi, J., Eryilmaz, H., Schwartz, S., Vuilleumier, P., & Van De Ville, D. (2011). Decoding brain states from fMRI connectivity graphs. *NeuroImage*, 56(2), 616–626. <https://doi.org/10.1016/j.neuroimage.2010.05.081>. Multivariate Decoding and Brain Reading.
- Rieck, J. R., Baracchini, G., Nichol, D., Abdi, H., & Grady, C. L. (2021a). Dataset of functional connectivity during cognitive control for an adult lifespan sample. *Data in Brief*, 39, 107573. <https://doi.org/10.1016/j.dib.2021.107573>
- Rieck, J. R., Baracchini, G., Nichol, D., Abdi, H., & Grady, C. L. (2021b). Reconfiguration and dedifferentiation of functional networks during cognitive control across the adult lifespan. *Neurobiology of Aging*, 106, 80–94.
- Schaefer, A., Kong, R., Gordon, E. M., Laumann, T. O., Zuo, X.-N., Holmes, A. J., Eickhoff, S. B., & Yeo, B. T. (2018). Local-global parcellation of the human cerebral cortex from intrinsic functional connectivity MRI. *Cerebral cortex*, 28(9), 3095–3114.
- Scott, J. G., & Berger, J. O. (2010). Bayes and empirical-Bayes multiplicity adjustment in the variable-selection problem. *The Annals of Statistics*, 2587–2619.
- Snijders, T. A. B., Lomi, A., & Torló, V. J. (2013). A model for the multiplex dynamics of two-mode and one-mode networks, with an application to employment preference, friendship, and advice. *Social Networks*, 35(2), 265–276. <https://doi.org/10.1016/j.socnet.2012.05.005>
- Spencer, D., Guhaniyogi, R., Shinohara, R., & Prado, R. (2022). Bayesian tensor regression using the Tucker decomposition for sparse spatial modeling. arXiv preprint [arXiv:2203.04733](https://arxiv.org/abs/2203.04733)
- Thompson, P., Moussai, J., Zohoori, S., Goldkorn, A., Khan, A., Mega, M., Small, G., Cummings, J., & Toga, A. (1998). Cortical

- variability and asymmetry in normal aging and Alzheimer's disease. *Cerebral Cortex* (New York, NY: 1991), 8(6), 492–509.
- Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society Series B (Methodological)*, 58(1), 267–288.
- Wang, M., El-Fiqi, H., Hu, J., & Abbass, H. A. (2019). Convolutional neural networks using dynamic functional connectivity for EEG-based person identification in diverse human states. *IEEE Transactions on Information Forensics and Security*, 14(12), 3259–3272.
- Whitfield-Gabrieli, S., & Nieto-Castanon, A. (2012). Conn: A functional connectivity toolbox for correlated and anticorrelated brain networks. *Brain connectivity*, 2(3), 125–141.
- Xu, S., Zhen, Y., & Wang, J. (2023). Covariate-assisted community detection in multi-layer networks. *Journal of Business & Economic Statistics*, 41(3), 915–926. <https://doi.org/10.1080/07350015.2022.2085726>
- Yang, Y., Ye, C., & Ma, T. (2023). A deep connectome learning network using graph convolution for connectome-disease association study. *Neural Networks*, 164, 91–104.
- Yeo, B. T., Krienen, F. M., Sepulcre, J., Sabuncu, M. R., Lashkari, D., Hollinshead, M., Roffman, J. L., Smoller, J. W., Zöllei, L., Polimeni, J. R., et al. (2011). The organization of the human cerebral cortex estimated by intrinsic functional connectivity. *Journal of Neurophysiology*.
- Zhang, X., Xu, G., & Zhu, J. (2022). Joint latent space models for network data with high-dimensional node variables. *Biometrika*, 109(3), 707–720. <https://doi.org/10.1093/biomet/asab063>
- Zhao, R., Liu, T., Huang, Z., Lun, D. P., & Lam, K.-M. (2022). Spatial-temporal graphs plus transformers for geometry-guided facial expression recognition. *IEEE Transactions on Affective Computing*.
- Zhou, H., Li, L., & Zhu, H. (2013). Tensor regression with applications in neuroimaging data analysis. *Journal of the American Statistical Association*, 108(502), 540–552. <https://doi.org/10.1080/01621459.2013.776499>. PMID: 24791032.
- Zhou, W., Qu, A., Cooper, K. W., Fortin, N., & Shahbaba, B. (2023). A model-agnostic graph neural network for integrating local and global information. Preprint at <http://arxiv.org/abs/2309.13459>
- Zhou, Y., Wong, R. K., & He, K. (2024). Broadcasted nonparametric tensor regression. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 027.
- Zhou, J., Cui, G., Hu, S., Zhang, Z., Yang, C., Liu, Z., Wang, L., Li, C., & Sun, M. (2020). Graph neural networks: a review of methods and applications. *AI open*, 1, 57–81.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Springer Nature or its licensor (e.g. a society or other partner) holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.