

RESEARCH ARTICLE

WILEY

Bayesian varying-effects vector autoregressive models for inference of brain connectivity networks and covariate effects in pediatric traumatic brain injury

Yangfan Ren¹ | Nathan Osborne² | Christine B. Peterson³ | Dana M. DeMaster⁴ | Linda Ewing-Cobbs⁴ | Marina Vannucci¹ 

¹Department of Statistics, Rice University, Houston, Texas, USA

²Intuit, San Diego, California, USA

³Department of Biostatistics, The University of Texas MD Anderson Cancer Center, Houston, Texas, USA

⁴Department of Pediatrics, Children's Learning Institute, University of Texas Health Science Center, Houston, Texas, USA

Correspondence

Marina Vannucci, Department of Statistics, Rice University, Houston, TX, USA.
Email: marina@rice.edu

Funding information

National Science Foundation, Grant/Award Numbers: DMS-2113602, DMS-2113557; National Institutes of Health, Grant/Award Number: R01 NS046308

Abstract

In this article, we develop an analytical approach for estimating brain connectivity networks that accounts for subject heterogeneity. More specifically, we consider a novel extension of a multi-subject Bayesian vector autoregressive model that estimates group-specific directed brain connectivity networks and accounts for the effects of covariates on the network edges. We adopt a flexible approach, allowing for (possibly) nonlinear effects of the covariates on edge strength via a novel Bayesian nonparametric prior that employs a weighted mixture of Gaussian processes. For posterior inference, we achieve computational scalability by implementing a variational Bayes scheme. Our approach enables simultaneous estimation of group-specific networks and selection of relevant covariate effects. We show improved performance over competing two-stage approaches on simulated data. We apply our method on resting-state functional magnetic resonance imaging data from children with a history of traumatic brain injury (TBI) and healthy controls to estimate the effects of age and sex on the group-level connectivities. Our results highlight differences in the distribution of parent nodes. They also suggest alteration in the relation of age, with peak edge strength in children with TBI, and differences in effective connectivity strength between males and females.

KEYWORDS

brain connectivity, fMRI data, Gaussian process, spike-and-slab prior, traumatic brain injury, variational inference

1 | INTRODUCTION

Functional magnetic resonance imaging (fMRI) measures blood-oxygenation-level-dependent (BOLD) contrast, which reflects the

difference in magnetization between oxygenated and deoxygenated blood arising from patterns in cerebral blood flow. Changes in BOLD response are treated as a proxy for changes in neurological activity. This technique remains one of the most popular for measuring brain connectivity, mostly due to its noninvasive nature. In this article, we are interested in effective connectivity, that is, the directed influence

Yangfan Ren and Nathan Osborne contributed equally to this study.

This is an open access article under the terms of the [Creative Commons Attribution-NonCommercial-NoDerivs](https://creativecommons.org/licenses/by-nc-nd/4.0/) License, which permits use and distribution in any medium, provided the original work is properly cited, the use is non-commercial and no modifications or adaptations are made.

© 2024 The Author(s). *Human Brain Mapping* published by Wiley Periodicals LLC.

that one neural system exerts over another, which is often estimated based on resting-state fMRI data (Friston, 2011).

Statistical approaches to modeling effective connectivity among brain regions include dynamic causal modeling (DCM, Friston et al., 2003), structural equation modeling (SEM, McIntosh & Gonzalez-Lima, 1994), Bayesian networks (BNs, Li et al., 2008; Rajapakse & Zhou, 2007), and Granger causality (CG) modeling via vector autoregressive (VAR) models (Granger, 1969; Roebroeck et al., 2005). DCM and SEM are typically used as confirmatory techniques to test predefined hypotheses about neural activity (Friston, 2011). BNs employ directed acyclic graphs, ignoring the high prevalence of reciprocal connections that commonly renders brain connectivity cyclic (Friston, 2011). GC is based on the notion that causes both precede and help predict their effects. This approach infers effective connectivity by estimating coefficients from VAR models. It is important to note that even though such methods allow inference on directed connections between brain regions, causality between fMRI signals does not translate into causality of the corresponding neuronal activity (Wen et al., 2013).

VAR models have been used to estimate whole-brain connectivity networks, where nodes represent brain regions of interest (ROIs) obtained from a parcellation of the brain. Sparsity plays an essential role in the estimation of these models. A common approach is to impose sparsity through an ℓ_1 penalty on the VAR coefficients (Arnold et al., 2007; Valdés-Sosa et al., 2005). For fMRI studies with multiple subjects, Gorrostieta et al. (2012, 2013) proposed mixed-effect VAR models that achieve group-level selection, by identifying significant connections between ROIs that are consistent across a group of subjects. Chiang et al. (2017) considered a supervised setting with subjects belonging to multiple groups and employed spike-and-slab priors to achieve sparsity in the group-level networks. To improve the computational scalability, Kook et al. (2021) developed a variational Bayes (VB) approach for the model of Chiang et al. (2017). See also Wang et al. (2023) for a recently proposed VB method for an autoregressive state-space model. While these VAR-based models have significantly advanced the understanding of brain connectivity, they do not account for subject heterogeneity.

In fMRI studies, it is common to measure subject-level covariates, such as age, sex and behavioral assessment scores, in addition to the imaging data. In recent years, many researchers have focused on the question of how to relate subject-level covariates to the observed imaging data (Guha & Guhaniyogi, 2021; Kundu et al., 2021; Scheffler et al., 2019; Zhao et al., 2021). Generally speaking, external covariates can be incorporated into the estimation of graphs by linking them to either node values (as covariate-adjusted models) or edge strengths (covariate-dependent). However, most of these contributions in graphical modeling have considered the framework of undirected networks. In this setting, graphs can be expressed as regression models that link the mean values of the network nodes to external covariates, allowing the estimation of a conditional network that reflects dependencies among node variables after adjusting for external covariates (Li et al., 2012; Yin & Li, 2011). Time-varying graphs can also be considered a special case of covariate-dependent graphical models, where

the network structure is allowed to change smoothly over time (Kolar et al., 2010; Zhou et al., 2010). More advanced models may allow the external covariates to have nonlinear effects on edge strength. These approaches build on regression settings with varying effects, or varying coefficient models (Cleveland & Grosse, 1991; Hastie & Tibshirani, 1993). By modeling regression coefficients as a function of a covariate, these modeling approaches allow greater flexibility than linear regression.

In this article, we develop an analytical approach for estimating brain connectivity networks that accounts for subject heterogeneity by modeling the effect of covariates on the edge strengths. More specifically, we build upon the VAR framework of Chiang et al. (2017) and Kook et al. (2021) to construct a varying-effect VAR modeling framework that estimates group-specific brain connectivity networks and accounts for the effects of subject-level covariates on the network edges. We name our method Varying-Effects Vector Autoregression (VEVAR). Within this modeling framework, we are able to estimate both a group-level graph structure and edge strengths as (possibly) nonlinear functions of subject-level covariates. We achieve this via Gaussian process priors that capture the edge strengths as smooth functions of covariate values. We model the covariate effects as a sum of univariate Gaussian processes, which allows for edge-specific covariate effect selection. Additionally, we use variable selection spike-and-slab priors to determine the presence or absence of edges. For posterior inference, we address the scalability limitations of the existing models by implementing a VB approach (Blei et al., 2017; Dance & Paige, 2022; Titsias & Lázaro-Gredilla, 2011). We use simulated data to compare our method with two-stage frequentist and Bayesian approaches that first estimate the networks and then select the covariates that explain the edge strengths. Our results show that the proposed VEVAR model does well in terms of both group-level edge selection and covariate effect selection.

Next, we apply our method to resting-state fMRI data from children with a history of traumatic brain injury (TBI) and healthy controls (HCs) to characterize age and sex effects on neural circuitry. TBI is particularly concerning because it can disrupt the typical course of brain development and lead to cascading effects on health-related quality of life. Furthermore, TBI outcomes are characterized by significant heterogeneity. Consequently, statistical approaches that can account for variability related to subject-level covariates can substantially refine the sensitivity and utility of connectivity network modeling. Here, in addition to the estimation of the group-level connectivities, we evaluate whether specific group-level edge connectivity strengths are affected by age and sex. A unique feature of our approach is that, unlike other methods, VEVAR allows the estimation of group-level edges as (possibly) nonlinear functions of these covariates. Our results highlight differences in the distribution of parent nodes. They also suggest alteration in the relation of age, with peak edge strength in children with TBI, and differences in effective connectivity strength between males and females. We provide some discussion corroborating our findings.

The rest of this article is organized as follows. In Section 2, we introduce the model and the prior construction and briefly describe

the proposed variational approach for inference. We also introduce the simulation design and the empirical data. In Section 3, we assess the performance of our proposed method against competing approaches using simulated data and illustrate our method on resting-state functional MRI data collected on children with a history of traumatic injury and HCs. Section 4 provides some interpretation of the results and Section 5 concludes the paper with a discussion on limitations and future directions.

2 | METHODS AND MATERIALS

2.1 | Statistical model

We describe the proposed varying-effects VAR model for simultaneous group- and subject-level network estimation and selection of edge-specific covariate effects. A graphical representation of the model is provided in Figure 1.

2.1.1 | Likelihood

We first introduce the structure of the observed data and the likelihood. In our setting of interest, we observe multivariate time series data on different groups of subjects. Specifically, let $\mathbf{x}_t^{(s)} = [x_{t,1}^{(s)}, \dots, x_{t,R}^{(s)}]$ represent the R dimensional vector of observed values for subject s at time t , with $s = 1, \dots, n$, and $t = 1, \dots, T$. We assume that subjects are classified into G groups with $\boldsymbol{\eta} = [\eta_1, \dots, \eta_n]$ indicating the group membership for each subject, where $\eta_s = g$ if subject s belongs to group $g = 1, \dots, G$. Furthermore, we assume that an additional set of P fixed covariates $\mathbf{m} = [m_1, \dots, m_P]$ are observed for each subject.

Within each group, we model the subject-level time series data via a VAR model that expresses the observed values at time t as a function of the R variables at the previous L lagged time points

$$\underbrace{\mathbf{x}_t^{(s)}}_{1 \times R} = \underbrace{\mathbf{u}_t^{(s)}}_{1 \times RL} \cdot \underbrace{\mathbf{B}^{(s)}}_{RL \times R} + \underbrace{\mathbf{e}_t^{(s)}}_{1 \times R}, \quad (1)$$

where $\mathbf{u}_t^{(s)} = [\mathbf{x}_{t-1}^{(s)}, \mathbf{x}_{t-2}^{(s)}, \dots, \mathbf{x}_{t-L}^{(s)}]$ is the $1 \times RL$ vector of concatenated lagged measurements, $\mathbf{B}^{(s)}$ is the $RL \times R$ matrix of subject-specific VAR coefficients, and, given $\eta_s = g$, $\mathbf{e}_t^{(s)} \sim \mathcal{N}(0, \Xi^{(g)})$ with $\Xi^{(g)} = \text{diag}(\xi_1^{(g)}, \dots, \xi_R^{(g)})$ represents independent Gaussian noise. To accommodate all T time points, we can write $\mathbf{X}^{(s)} = [\mathbf{x}_2^{(s)}, \dots, \mathbf{x}_T^{(s)}]'$, $\mathbf{U}^{(s)} = [\mathbf{u}_1^{(s)}, \dots, \mathbf{u}_{T-L}^{(s)}]'$, and $\mathbf{E}^{(s)} = [\mathbf{e}_1^{(s)}, \dots, \mathbf{e}_{T-1}^{(s)}]'$. We then use the vec operator, which converts a matrix to a single column vector, and can write model (1) as

$$\underbrace{\mathbf{x}^{(s)}}_{(T-L)R \times 1} = \left(\underbrace{\mathbf{I}}_{R \times R} \otimes \underbrace{\mathbf{U}^{(s)}}_{(T-L) \times RL} \right) \underbrace{\boldsymbol{\beta}^{(s)}}_{(RL)R \times 1} + \underbrace{\mathbf{e}^{(s)}}_{(T-L)R \times 1}, \quad (2)$$

where $\mathbf{x}^{(s)}$, $\boldsymbol{\beta}^{(s)}$, and $\mathbf{e}^{(s)}$ are equivalent to $\text{vec}(\mathbf{X}^{(s)})$, $\text{vec}(\boldsymbol{\beta}^{(s)})$, and $\text{vec}(\mathbf{E}^{(s)})$, respectively, and $\otimes$ represents the Kronecker product (Van Loan, 2000). By writing the model in form (2), we can see that


$$\mathbf{x}^{(s)} | \eta_s = g \sim \mathcal{N} \left((\mathbf{I} \otimes \mathbf{U}^{(s)}) \boldsymbol{\beta}^{(s)}, \Xi^{(g)} \otimes \mathbf{I} \right), \quad (3)$$

which allows us to recognize this as a linear regression problem. The observed values for subject s , given the group membership g , follow a normal distribution with mean $(\mathbf{I} \otimes \mathbf{U}^{(s)}) \boldsymbol{\beta}^{(s)}$ and covariance $\Xi^{(g)} \otimes \mathbf{I}$, where $\mathbf{I}$ denotes the identity matrix and $\Xi^{(g)}$ is an error covariance matrix defined as above. The subject-specific coefficients $\boldsymbol{\beta}^{(s)} = [\beta_1^{(s)}, \dots, \beta_{(RL)R}^{(s)}]$ capture the temporal relationship of variables. This model is illustrated in the right portion of Figure 1.

2.1.2 | Varying-effects selection priors

Chiang et al. (2017) used the VAR setting (1) to model fMRI data observed on multiple groups of subjects. In that setting, inference on the VAR regression coefficients allows the estimation of directed networks, where network nodes are brain regions and an edge represents a directional influence of one region on another (Friston, 1994). The authors adopted a hierarchical structure on the $\boldsymbol{\beta}^{(s)}$ coefficients by

FIGURE 1 Graphical representation of the proposed model: circular nodes indicate parameters, and square nodes represent observed data. Links between nodes represent direct probabilistic dependence. Some hyperparameters are not shown for clarity. Indices refer to subject, $s = 1, \dots, n$, coefficient, $j = 1, \dots, (RL)R$, and group, $g = 1, \dots, G$.



assuming that they are generated from a common group-level coefficient vector $\underline{\beta}^{(g)}$, and imposed spike-and-slab priors on the group-level coefficients. Here, we build upon this supervised hierarchical setting by proposing a novel prior formulation that allows the edge strengths of the group-level graphs to vary as functions of the covariate values. We achieve this by imposing nonparametric spike-and-slab priors at the group level that allow edges to be either zero or nonzero functions of the covariate values. Additionally, we employ univariate spike-and-slab priors to select covariates influencing individual edge strengths, for each subject group. Our novel prior formulation allows for nonlinear effects of covariates on the edge strengths of the inferred networks and simultaneous selection of the relevant effects.

As in Chiang et al. (2017), we assume that the VAR coefficients for each subject, $\beta_j^{(s)}$, are noisy realizations of a group-level network. However, at the group level, we model the edge strengths as functions of the subject-specific covariate values, $\mathbf{m}$. We write

$$\beta_j^{(s)} | \mathbf{m} \sim \mathcal{N}(\beta_j^{(g)}(\mathbf{m}), \sigma_j^{(g)}), \quad (4)$$

for all subjects s such that $\eta_s = g$, with $j = 1, \dots, (RL)R$ and $g = 1, \dots, G$. We place the priors $\sigma_j^{(g)} = \sigma_0^{(g)} \sim IG(a_0^{(g)}, b_0^{(g)})$ for $\delta_j^{(g)} = 0$ and $\sigma_j^{(g)} = \sigma_1^{(g)} \sim IG(a_1^{(g)}, b_1^{(g)})$ for $\delta_j^{(g)} = 1$. We seek to estimate a network for each of the G groups, that is, to find the nonzero group-level edges. We achieve this by proposing a novel discrete nonparametric spike-and-slab prior that imposes sparsity on the networks at the group level while modeling the nonzero edges as smooth functions of the covariate values:

$$\beta_j^{(g)}(\mathbf{m}) = \delta_j^{(g)} f_j^{(g)}(\mathbf{m}) + (1 - \delta_j^{(g)}) \delta_0(\mathbf{m}), \quad (5)$$

with $\delta_j^{(g)} \sim \text{Bernoulli}(\pi_\delta)$, where the functional notation $\beta_j^{(g)}(\mathbf{m})$ represents possible values of the VAR coefficient as a function of the covariate values $\mathbf{m}$, the “spike” δ_0 represents a Dirac delta function, and the notation $\delta_j^{(g)} = 0$ implies that the group-level coefficient for group g is zero for all possible values of $\mathbf{m}$. Spike-and-slab priors typically include a normal distribution as the “slab” portion (George & McCulloch, 1997; Vannucci, 2021). We innovate on this framework, to instead allow the strengths of the selected nonzero edges to depend on covariate values in a nonlinear manner, as described below. A nice result of this modeling choice is that we can obtain a set of directed edges for each group, that is, a form of interpretable functional group networks where the strengths of the included edges vary as smooth functions of the covariate values. Our nonparametric “slab” portion of the mixture prior (5) is defined as follows. We first incorporate the covariates as

$$f_j^{(g)}(\mathbf{m}) = \mu_j^{(g)} + \sum_{p=1}^P w_{j,p}^{(g)} \phi_{j,p}^{(g)}(m_p), \quad (6)$$

with $\mu_j^{(g)} \sim \mathcal{N}(0, \sigma_\mu^2)$ a baseline value reflecting edge strength not driven by covariate effects, $\phi_{j,p}^{(g)}(m_p)$ a univariate function dependent on the p th covariate, and $w_{j,p}^{(g)}$ the corresponding weight. We then

model the functions $\phi_{j,p}^{(g)}(m_p)$ via Gaussian process priors (Williams & Rasmussen, 2006) as

$$\phi_{j,p}^{(g)} \equiv GP(\mu(\cdot), K(\cdot, \cdot)), \quad (7)$$

with $\mu(\cdot) = 0$ and covariance function $K(m_{p,k}, m_{p,k'}) = \text{cov}(f(m_{p,k}), f(m_{p,k'}))$, which allows us to model nonlinear effects of the covariates on edge strength. Finally, we impose a parametric discrete spike-and-slab prior on the coefficients in (6) as

$$w_{j,p}^{(g)} \sim \pi_\phi \mathcal{N}(0, \sigma_w^2) + (1 - \pi_\phi) \delta_0, \quad (8)$$

where π_ϕ is the prior probability that the coefficient is nonzero, that is, the corresponding covariate is determined to be important in estimating the edge strength function, $\beta_j^{(g)}(\mathbf{m})$. For each edge j in group g , formulation (6) employs P nonlinear functions, $\phi_{j,p}^{(g)}(m_p)$, one for each covariate. Thus, using the spike-and-slab prior (8) on the individual $w_{j,p}^{(g)}$'s allows us to select possibly different covariates, with nonlinear effects, for each group-level edge. We complete our model by assuming priors $\xi_r^{(g)} \sim IG(a_\xi, b_\xi)$, for $r = 1, \dots, R$, on the diagonal elements of $\Xi^{(g)}$.

The proposed modeling framework described here represents a novel approach to inference of covariate effects on edge strength within a VAR model that leads to the estimation of sparse networks at the group level and the selection of key covariates that influence individual edge strengths. Figure 2 shows an illustration of our proposed varying-effect nonparametric selection prior. The top row shows group-level networks, which will be estimated using the spike-and-slab prior (5). The lower row shows the underlying group-level function, $\beta_j^{(g)}(\mathbf{m})$, modeled via the nonparametric slab (6) incorporating the covariates, as a function of a single covariate, m_1 , and for a single edge, $j = 1$. Subject-level edge strengths, $\beta_1^{(s)}(m_1)$, for all s such that $\eta_s = g$, are shown as points following the shape of the function $\beta_1^{(g)}(m_1)$. Notice that the edge corresponding to $\beta_1^{(g)}$ is present in the group-level network for $g = 1$ but not for $g = G$.

The GP construction in (7) requires selecting a kernel function. Here, we used a squared exponential (SE) of the form $K(x, x') = \sigma^2 \exp\left(-\frac{(x-x')^2}{2l^2}\right)$. We found a length scale fixed at $l = 0.5$ to be sufficiently flexible for our application. In addition, we found this kernel to work for both continuous and categorical covariates, as shown by our results in the simulation study presented below. Our methodology is general and can accommodate any valid kernel function. Moreover, a common kernel does not need to be used across all covariates or across all edges. Additional flexibility could be gained by placing a prior on any of the hyperparameters of the kernel.

2.1.3 | Variational algorithm for scalable inference

One challenge in network estimation is the large number of parameters that need to be estimated. This is particularly challenging in Bayesian estimation, as Bayesian methods frequently rely on Markov




FIGURE 2 An example diagram of the proposed varying-effects selection prior, with an illustration of group-level edges, subject-level coefficients, and group-level coefficient functions. The top row shows group-level networks. The lower row shows the underlying group level function, $\beta_j^{(g)}(\mathbf{m})$, as a function of a single covariate, m_1 , for a single edge, $j = 1$. Subject-level edge strengths, $\beta_1^{(g)}(m_1)$, are shown as points following the shape of the function $\beta_1^{(g)}(m_1)$.

chain Monte Carlo methods (Gelman et al., 2013) to sample from the posterior. In our model formulation, complexity and computational costs are compounded by the consideration of the covariates and the large number of inclusion indicators we need to estimate in the non-parametric spike-and-slab construction.

To allow the model to scale up to large data sizes, in particular those encountered in our application setting, we implement a variational approximation method for posterior inference which dramatically improves computational time. Variational schemes for linear models that use spike-and-slab priors have been previously described (Carbonetto & Stephens, 2012; Kook et al., 2021; Titsias & Lázaro-Gredilla, 2011). Variational inference (VI) aims at finding an approximation of the posterior by using optimization methods. It works by specifying a family of approximate distributions, Q , which are densities over model parameters and latent variables that depend on free parameters Θ , and then seeks to find the values of Θ that minimize the Kullback–Leibler (KL) divergence between the approximate distribution and the true posterior. Let $\mathbf{Z}$ indicate the set of model

parameters and latent variables. As discussed in Blei et al. (2017), minimizing the KL divergence is equivalent to maximizing the Evidence Lower Bound (ELBO), defined as

$$\text{ELBO} = \mathbb{E}_{\Theta}[\log p(\mathbf{Z}, \mathbf{Y})] - \mathbb{E}_{\Theta}[\log q(\mathbf{Z})], \quad (9)$$

with $p(\mathbf{Z}, \mathbf{Y})$ the joint distribution of $\mathbf{Z}$ and the data, and $q(\mathbf{Z})$ the approximate distribution. The complexity of the optimization procedure is determined by the complexity of the variational distribution. A common family of approximate distributions is the mean field approximation, which assumes that the approximate distribution factorizes over some partition of the parameters and latent variables as

$$q(\mathbf{Z}) = \prod_k q(Z_k). \quad (10)$$

The exact parametric form of each $q(Z_k)$ is selected based on whether Z_k is continuous or discrete, and may exploit conditional

conjugate distributions to allow for simpler derivation of the optimization steps.

In our framework, we adopt a mean field approach and introduce a family of approximate distributions $Q(\Theta)$. We follow the work of Titsias and Lázaro-Gredilla (2011) by reparameterizing the spike-and-slab prior in (8) through the introduction of new variables $\tilde{w}_{j,p}^{(g)} \sim \mathcal{N}(0, \sigma_w^2)$ and $s_{j,p} \sim \pi_{\phi}^{s_{j,p}} (1 - \pi_{\phi})^{1-s_{j,p}}$. This allows us to rewrite $w_{j,p}^{(g)}$ as $w_{j,p}^{(g)} = \tilde{w}_{j,p}^{(g)} s_{j,p}^{(g)}$. We then find $q(\tilde{w}_{j,p}^{(g)} | s_{j,p}^{(g)})$ and $q(s_{j,p}^{(g)})$ and model $\{\tilde{w}_{j,p}^{(g)}, s_{j,p}^{(g)}\}$ jointly, according to the mean field approach. The remaining variational distributions are proposed as:

$$\begin{aligned} q(\mu_j^{(g)}) &\sim \mathcal{N}(u_j^{(g)}, v_j^{(g)}) \\ q(\delta_j^{(g)}) &\sim \text{Bernoulli}(\gamma_{(\delta)j}^{(g)}) \\ q(\underline{\beta}^{(s)} | \underline{\beta}^{(g)}(\mathbf{m}^{(s)}), \sigma_j) &\sim \mathcal{N}(\mathbf{u}_\beta^{(s)}, \mathbf{S}^{(s)}) \\ q(\xi_r^{(g)} | a_\xi^{(g)}, b_\xi^{(g)}) &\sim \text{IG}(z_{1r}^{(g)}, z_{2r}^{(g)}) \\ q(\sigma_0^{(g)} | a_0^{(g)}, b_0^{(g)}) &\sim \text{IG}(a_0^{(g)}, b_0^{(g)}) \\ q(\sigma_1^{(g)} | a_1^{(g)}, b_1^{(g)}) &\sim \text{IG}(a_1^{(g)}, b_1^{(g)}) \\ q(\phi_{j,p}^{(g)}(\mathbf{m})) &\sim \mathcal{N}(\tilde{\mu}_{j,g}^\phi, K_{j,g}^\phi) \\ q(\tilde{w}_{j,p}^{(g)} | s_{j,p}^{(g)}) &\sim \mathcal{N}(\omega_{j,p}^{(g)}, \tilde{\sigma}_{j,p}^{(g)}) \\ q(s_{j,p}^{(g)}) &\sim \text{Bernoulli}(\gamma_{(\phi)j,p}^{(g)}), \end{aligned} \quad (11)$$

where $\underline{\beta}^{(g)}(\mathbf{m}) = [\beta_1^{(g)}(\mathbf{m}), \dots, \beta_{(RL)R}^{(g)}(\mathbf{m})]$ is the collection of realizations of the group-level function at subject s 's covariate measurements, and $K_{j,g}^\phi$ is the covariance matrix corresponding to the kernel $K(\mathbf{m}_p^{(g)}, \mathbf{m}_p^{(g)})$ where $\mathbf{m}_p^{(g)} = [\mathbf{m}_p^{(s=1)}, \dots, \mathbf{m}_p^{(s=N)}]$ for $\eta_s = g$. Using a mean field family, and the reparameterizations discussed above, we write the full approximate distribution for all variational parameters as a product over the approximate distributions:

$$Q(\Theta) = \prod_{g=1}^G \left[\prod_{j=1}^{R \times R \times L} \left\{ q(\mu_j^{(g)}) q(\delta_j^{(g)}) \prod_{p=1}^P q(\phi_{j,p}^{(g)}) q(\tilde{w}_{j,p}^{(g)} | s_{j,p}^{(g)}) q(s_{j,p}^{(g)}) \right\} \right. \\ \left. \times \prod_{r=1}^R q(\xi_r^{(g)}) \right] \prod_{s=1}^N q(\beta_j^{(s)}) q(\sigma_1^{(s)}) q(\sigma_0^{(s)})$$

where Θ represents the parameters to be optimized. As common with VI, these are updated via coordinate ascent VI. This algorithm is repeated until the ELBO has converged or has changed by some minimum threshold. Updates of the parameters and the parameter blocks are outlined in the Supplementary Material.

Posterior inference based the VI algorithm results in estimates of the group-level networks, together with estimates of the edge strengths at the group- and subject-level, and simultaneous selection of the covariates that are relevant to the edge strengths. After convergence of the VI algorithm, we classify edge j as being present in group g if the estimated $\gamma_{(\delta)j}^{(g)}$ is greater than 0.5, and select covariate p as affecting the j th edge strength in group g if the estimated $\gamma_{(\phi)j,p}^{(g)}$ is greater than 0.5. We infer the nonzero group-level edge strength functions by the values of $u_j^{(g)} + \sum_{p=1}^P \gamma_{(\phi)j,p}^{(g)} \omega_{j,p}^{(g)} \tilde{\mu}_{j,g}^\phi$ in (11) and the

subject-level edge strengths by the values of $u_\beta^{(s)}$ in (11). We note that variational approaches are only suitable for point estimation and do not allow the assessment of uncertainty about the estimates. Additionally, in situations with correlated covariates, performance can be sensitive to initializations and to the order that variables are updated, possibly resulting in poor selection performance. Ray and Szabó (2022) proposed a prioritized updating scheme, where the importance of the variable is determined using a preliminary estimator. Here, we obtained the preliminary estimator as the average of the output from several cold starts, where each covariate is afforded the opportunity to be the first and also the last covariate proposed.

2.2 | Simulation experiment

We use simulated data to test the performance of the proposed model and compare results to competing approaches.

In our simulation, we considered two groups of subjects with sample sizes 30 and 60, respectively, to assess the robustness of the model to an unbalanced sample size setting. For each subject, we generated the time series data $\mathbf{x}_t^{(s)}$ to represent realistic fMRI signal across ROIs over time. These data were drawn from a multivariate normal distribution $\mathcal{N}(\mathbf{u}_t^{(s)} \mathbf{B}^{(s)}, \sigma^2 \mathbf{I}_R)$, for $t = 2, \dots, 200$, with noise variance $\sigma^2 = 0.5$. We selected $R = 100$ ROIs to mirror the level of granularity commonly observed in brain network studies. This choice ensures that our simulation represents the complex connectivity between different brain regions. The initial time point $\mathbf{x}_1^{(s)}$ was randomly drawn from $\mathcal{N}(\mathbf{0}, 0.25 \mathbf{I}_R)$. We simulated the VAR coefficients $\mathbf{B}^{(s)}$, representing the connectivity strength and direction between different ROIs, in a manner that allowed us to test a range of possible functions and covariate associations. The time lag value was set to $L = 1$, ensuring the simulation of direct, immediate connections typical in brain networks. We checked to make sure that the process generated only stationary time series for each subject, and repeated the data generating process if not. We considered six random subject-level covariates, $\mathbf{m} = (m_1, \dots, m_6)$, with the first five drawn from a uniform $(-1, 1)$, and the sixth one from a Bernoulli (0.5). These covariates affect the edge strengths within the brain connectivity network, inducing individual differences in neural connections. Next, we generated each element of the subject level vector of coefficients $\mathbf{B}^{(s)}$ from a $\mathcal{N}(f(\mathbf{m}), 0.08)$, with $f(\cdot)$ the generating function of covariates on the edge strengths defined by setting each element $B_{jj'}^{(s)}$, with $j = 1, \dots, RL$ and $j' = 1, \dots, R$, as follows:

$$\begin{aligned} \text{if } |j - j'| = 0, f(\cdot) &= 0.15 \text{ (i.e., constant)} \\ \text{if } |j - j'| = 1, f(\cdot) &= m_p \cdot 0.25 \\ \text{if } |j - j'| = 2, f(\cdot) &= (1.0 + m_p)^{0.40} \\ \text{if } |j - j'| = 3, f(\cdot) &= (m_p)^2 \\ \text{if } |j - j'| = 4, f(\cdot) &= 0.20 \cdot \text{sine}(\pi m_p) \\ \text{if } |j - j'| = 5, f(\cdot) &= 0.20 \cdot m_6 \\ \text{if } |j - j'| = 6, f(\cdot) &= 0.3m_{p_1} - 0.3m_{p_2} \\ \text{if } |j - j'| = 7, f(\cdot) &= (1.0 + m_p)^{0.70}, \end{aligned} \quad (12)$$

where m_p is randomly selected from the first five covariates, and m_{p_1} and m_{p_2} refer to two randomly selected covariates from the first five covariates. Note that m_p, m_{p_1} , and m_{p_2} are randomly selected separately for each $B_{jj'}^{(g)}$. Additionally, $f(\cdot)$ was rescaled to be within $[-0.2, 0.2]$ in the case where $|j - j'|$ is equivalent to one of $[2, 3, 4]$, and within $[-0.15, 0.35]$ when $|j - j'| = 7$, to explore scenarios with both a nonzero mean and covariate effect. Additionally, to more easily sample stationary time series, each function was multiplied by -1 with probability 0.5. Finally, to allow differences between the two sample groups, while remaining mostly similar, each $f^{(g)}(\cdot)$ outlined above was set to 0 with probability 0.2, resulting in some group-level edges being present in one group but not the other. By not restricting our simulation to predefined patterns, we can capture the diverse and complex network structure observed in the brain. Furthermore, the magnitude and variance of the edge strengths in our simulation are designed to replicate the estimated values in previous fMRI studies (Kook et al., 2021).

2.3 | Observational study on traumatic brain injury

We analyzed data from an fMRI study on children with a history of TBI following a vehicle collision and HCs. Subjects were recruited from the Emergency Department or Level 1 Pediatric Trauma Center at the Children's Memorial Hermann Hospital, University of Texas Health Science Center at Houston (UTHealth), between September 2011 and August 2015 (Ewing-Cobbs et al., 2019; Kook et al., 2021; Watson et al., 2019).

2.3.1 | Participants

Participants were included in the study if they met the following criteria: (1) injured in a vehicle accident between 8 and 15 years of age; (2) proficiency in English or Spanish; (3) residing within a 125 mile catchment radius; (4) no prior history of major neuropsychiatric disorder (intellectual deficiency or low-functioning autism spectrum disorder [ASD]) that would complicate assessment of the impact of injury on brain outcomes; (5) no metabolic, endocrine, or systemic health problems (e.g., hypertension); (6) no prior medically-attended TBI; and (7) no habitual use of steroids, tobacco, or alcohol. The latter four criteria were assessed during screening using a brief parent interview. A quality control evaluation of all scans resulted in 70 TBI samples being selected for analysis. Additionally, 50 subjects with no history of trauma were measured as part of a HC group. Subjects were excluded for excessive motion or scanner error (e.g., operator error, crack in scanner head coil). Written informed consent was obtained from each child's guardian and written assent was obtained from all children in accordance with Institutional Review Board guidelines. The demographic characteristics of the participants in the study are shown in Table 1.

TABLE 1 Demographic information for participants. Age is summarized as mean (range), and sex is summarized as number (percentage).

	TBI (n = 70)	HC (n = 50)
Age (years)	12.54 (8.08–16.0)	12.13 (8.08–16.0)
Sex (M, %)	45 (64%)	30 (60%)

2.3.2 | Data preprocessing

The fMRI data were preprocessed using SPM12 from the Wellcome Trust Center for Neuroimaging (<http://www.fil.ion.ucl.ac.uk/spm/>). The preprocessing steps involved correcting motion through realignment, correcting slice timing, separating gray matter, white matter, and CSF, registering the data to the subject's T1-weighted MPAGE image and a standard space using the ICBM space template, and smoothing with an 8 mm full-width half maximum Gaussian kernel. The data were then screened using the Artifact Detection Tools toolbox (Whitfield-Gabrieli et al., 2011) to eliminate volumes with excessive motion. Participants with more than 15% of their volumes affected by motion outliers were excluded from the analysis. Finally, a 3D parcellation was performed using the MarsBaR toolbox in SPM 12 and the automated anatomical labeling brain atlas, resulting in 90 ROIs excluding regions associated with the cerebellum.

2.3.3 | Previous findings

In Kook et al. (2021) and Vaughn et al. (2022), Bayesian vector-autoregressive models were employed to identify unique effective connectivity patterns within the default mode network in children with TBI relative to HCs. The connectivity patterns differed according to the severity of TBI and showed specific directional relations with symptom profiles. Fewer post-concussion and anxiety symptoms were associated with stronger regional effective orbitofrontal to posterior cingulate cortex connectivity for mild TBI, whereas weaker connectivity was associated with better outcomes for more severe TBI (Vaughn et al., 2022). These findings were similar to prior reports in adults with TBI showing different relations of frontal lobe connectivity with outcomes after mild TBI versus more severe TBI (Wu et al., 2015; Zhou et al., 2012).

3 | RESULTS

3.1 | Simulation results

We fit our proposed model by setting parameters $\pi_\phi = \pi_\delta = 0.1$ for the selection priors on covariate effects and group-level edges, respectively. We also set $\sigma_w = 1$ as the variance of the slab in (8) and $\sigma_\mu = 1$ as the variance of the baseline in (6). For the priors on the variance of the subject-level edge strengths, that is, $\sigma_0^{(g)} \sim \text{IG}(a_0^{(g)}, b_0^{(g)})$

TABLE 2 Simulation study: Results of competing methods and VEAR for group-level edge selection and covariate effect selection, evaluated using true positive rate (TPR), false positive rate (FPR), Matthew's correlation coefficient (MCC), F1 score, and accuracy (Acc). GC indicates Granger causality, and VEAR-S1 refers to the VEAR model without covariate effects. GC-LASSO employs LASSO to select covariates affecting edge strengths estimated by GC, while GC-plsmselect utilizes the method described in Ghosal and Kormaksson (2019). VEAR-S1-SS applies the spike-and-slab approach from Carbonetto and Stephens (2012) for selecting covariates impacting edge strengths in VEAR-S1. Bold values denote the best performance for each metric.

	Group 1					Group 2				
	TPR	FPR	MCC	F1	ACC	TPR	FPR	MCC	F1	ACC
<i>Edge selection</i>										
GC	.373	.060	.339	.407	.874	.456	.064	.400	.467	.881
VEVAR-S1	.560	.003	.711	.707	.947	.633	.018	.692	.716	.942
VEVAR	.935	.029	.852	.863	.962	.999	.039	.859	.868	.965
<i>Covariate effect selection</i>										
GC-LASSO	.275	.030	.190	.204	.950	.276	.022	.220	.236	.964
GC-plsmselect	.364	.127	.099	.098	.863	.359	.105	.115	.112	.884
VEVAR-S1-SS	.693	.020	.532	.527	.974	.780	.025	.546	.528	.971
VEVAR	.879	.000	.936	.935	.998	.998	.000	.999	.999	.999

and $\sigma_1^{(g)} \sim IG(a_1^{(g)}, b_1^{(g)})$, we imposed vague priors by setting $a_0 = a_1 = 2$ and $b_0 = b_1 = 1$. Similarly, for the prior on the variance of the time series data $\mathbf{x}$, that is, $\xi_r^{(g)} \sim IG(a_\xi, b_\xi)$, we set $a_\xi = 2, b_\xi = 1$. For the kernel of the Gaussian process, we set length-scale $l = 0.5$ and output variance $\sigma^2 = 1$. A sensitivity analysis of the parameter choices is provided in Section 3.1.2. We executed our simulations on a MacBook Pro with an M1 chip. The running time ranged from 12 to 16 h, according to the convergence speed of the VI algorithm for the different replicated datasets.

3.1.1 | Comparative performance

Table 2 provides results for group-level edge selection and for covariate effect selection, averaged over 25 replicated datasets, in terms of true positive rate (TPR), false positive rate (FPR), Matthew's correlation coefficient (MCC), F1 score, and accuracy (Acc). Specifically, the metrics are defined as

$$\begin{aligned}
 \text{TPR} &= \frac{\text{TP}}{\text{TP} + \text{FN}}, \\
 \text{FPR} &= \frac{\text{FP}}{\text{FP} + \text{TN}}, \\
 \text{MCC} &= \frac{\text{TP} \times \text{TN} - \text{FP} \times \text{FN}}{(\text{TP} + \text{FP})(\text{TP} + \text{FN})(\text{TN} + \text{FP})(\text{TN} + \text{FN})}, \\
 \text{F1} &= \frac{2\text{TP}}{2\text{TP} + \text{FP} + \text{FN}}, \\
 \text{Acc} &= \frac{\text{TP} + \text{TN}}{\text{TP} + \text{TN} + \text{FP} + \text{FN}},
 \end{aligned}$$

where TP, TN, FP, and FN denote the number of true positives, true negatives, false positives and false negatives, respectively. In the same table, we report results from alternative approaches. Since we are not aware of other methods that achieve simultaneous group-level edge selection and covariate effect selection, we considered

two-stage approaches that, at the first stage, estimate the networks and at the second stage select the covariates that explain the edge strengths. For one approach, at the first stage, that is, group-level edge selection, we considered a traditional GC model. GC (Granger, 1969) is a statistical hypothesis test to assess whether one time series can predict another time series. Here, it estimates subject-level VAR coefficients via ordinary least squares and then performs group-level inference through one-sample t tests. Nonzero group-level edges were identified by thresholding p -values with false discovery rate control at the .05 level. At the second stage, that is, covariate effect selection, we regressed the subject-level strengths estimated via the GC method on the covariates and performed selection using LASSO (Tibshirani, 1996) (GC-LASSO), which identifies linear covariate effects, and plsmselect (Ghosal & Kormaksson, 2019) (GC-plsmselect), which fits generalized additive models with flexible penalties, allowing for linear and smooth covariate effects. We also considered a second two-stage approach, obtained by considering at the first stage model (6) with the GP component removed and only an intercept term included. The resulting model, which we call VEAR-S1, estimates subject- and group-level networks without accounting for covariate effects. At the second stage, we regressed the estimated subject-level coefficients on the covariates and performed variable selection via the model of Carbonetto and Stephens (2012) (VEVAR-S1-SS), which uses mixture priors on the regression coefficients to perform variable selection. Given the relatively low FPRs for edge selection using GC and VEAR-S1, we performed selection at the second stage by using all estimated edges rather than only the selected ones.

Results from Table 2 show that our proposed VEAR model does well in terms of all performance measures, for both group-level edge selection and covariate effect selection. For edge selection, VEAR does considerably better than the other methods in all metrics, with only slightly higher FPRs compared to VEAR-S1. This is because the




FIGURE 3 Simulation study: True and estimated edge strengths and group-level functions. Gray points represent simulated subject-level edge strengths and gray lines represent the true group-level functions. Blue points represent estimated subject-level edge strengths and blue lines indicate the estimated group-level functions.

TABLE 3 Simulation study: Results of the proposed VEAR method, showing true positive rates (TPRs) of group-level edge selection for each of the underlying functions in Formula (12), which were used to generate the VAR coefficients. The results, averaged over all replicated datasets, indicate that all functions were recovered relatively well for both linear and nonlinear functions.

TPR by generating function							
$ i - j =$	1	2	3	4	5	6	7
Group 1	.944	.909	.807	.873	.999	.989	1.00
Group 2	1.00	1.00	.999	1.00	1.00	1.00	1.00

inclusion of covariate effects introduces more variability while improving the edge selection. For covariate effect selection, we see again that VEAR performs the best across all metrics.

Unlike competing methods, VEAR allows the estimation of group-level edges as (possibly) nonlinear functions of the covariate values. Figure 3 provides true and estimated subject-level edge strengths and group-level functions, for one of the simulated datasets, and Table 3 shows TPRs for each of the underlying functions in (12) that we used to generate the VAR coefficients, averaged over all replicated datasets. All functions were recovered relatively well for both linear and nonlinear functions and the estimated edge strengths were close to the true simulated values. For example, the average MSEs across selected edges generated by function 1 were 0.0034 and 0.0027 for the two groups, respectively, and those for function 3 were 0.0054 and 0.0043, respectively. As expected, it was more difficult to recover complex functions for the group with a smaller sample size, though the performance was generally good. It is also worth noting that the categorical covariate effects were selected correctly at a high rate, as seen in column 5.

3.1.2 | Sensitivity analysis

We investigated the sensitivity to the key parameter choices that determine the sparsity of the model, that is π_δ and π_ϕ . We also considered sensitivity to the choice of σ^2 , the variance of the GP kernel. We used the same simulation setting as outlined above, but with smaller networks ($R = 10$) for computational convenience. Parameters were varied one at a time keeping the others fixed at the values specified in the simulation as described above, and the reported results were averaged across 25 replicated datasets. Results from this sensitivity analysis are shown in Table 4, for both network edge selection and covariate effect selection. The parameter π_δ can be interpreted as the prior probability of a group-level VAR coefficient being nonzero. As π_δ increases, so does the number of selected edges. As expected, this parameter had no influence on the selection of the covariate effects. Parameter π_ϕ represents the prior probability of a covariate being selected as relevant for a certain edge strength. As π_ϕ increases so does the number of selected covariate effects. This in turn causes a slight increase in the number of selected edges. Despite this increase, the performance results remain pretty stable. The parameter σ^2 represents the variance of the GP kernel. As shown in Table 4, the numbers of selected edges and selected covariate effects increase as σ^2 increases, but this trend could be the opposite for data with different scales. A general trend that can be noticed in all results is that the inference is more robust to parameter selection for groups with larger sample sizes.

3.2 | Results from observational study on TBI

We are interested in estimating group-level connectivity networks and evaluating whether specific edge strengths are affected by

TABLE 4 Simulation study: Sensitivity analysis to assess model robustness by varying key parameters and observing their effects on performance metrics. The metrics evaluated are true positive rate (TPR), false positive rate (FPR), Matthew's correlation coefficient (MCC), F1 score, and accuracy (Acc). We focus on the parameters π_δ and π_ϕ , which determine the sparsity of the model, and σ^2 , the variance of the GP kernel. This analysis helps identify how the model's performance depends on specific input variables.

	Network edge selection									
	Group 1					Group 2				
	TPR	FPR	MCC	F1	Acc	TPR	FPR	MCC	F1	Acc
Varying π_δ										
.1	.733	.000	.636	.845	.798	.998	.010	.988	.997	.996
.5	.892	.460	.450	.881	.814	1.00	.626	.550	.907	.846
.9	1.00	1.00	.000	.871	.772	1.00	1.00	.000	.862	.758
Varying π_ϕ										
.1	.733	.000	.636	.845	.798	.998	.010	.988	.997	.996
.5	.821	.001	.739	.901	.868	1.00	.018	.988	.997	.995
.9	.931	.000	.876	.964	.948	.999	.016	.987	.996	.996
Varying σ^2										
.1	.487	.000	.429	.653	.609	.928	.018	.857	.959	.940
.5	.651	.000	.787	.786	.957	.998	.008	.976	.978	.994
1	.733	.000	.636	.845	.798	.998	.010	.988	.997	.996
	Covariate effect selection									
	Group 1					Group 2				
	TPR	FPR	MCC	F1	Acc	TPR	FPR	MCC	F1	Acc
Varying π_δ										
.1	.535	.000	.707	.692	.943	.979	.002	.981	.983	.996
.5	.525	.000	.701	.687	.941	.981	.001	.984	.985	.996
.9	.522	.000	.697	.683	.940	.981	.002	.982	0.984	0.996
Varying π_ϕ										
.1	.535	.000	.707	.692	.943	.979	.002	.981	.983	.996
.5	.706	.000	.824	.827	.965	.999	.003	.987	.988	.997
.9	.882	.000	.931	.936	.985	.999	.004	.985	.987	.997
Varying σ^2										
.1	.160	.000	.370	.271	.895	.864	.003	.906	.915	.980
.5	.420	.000	.602	.575	.873	.923	.002	.976	.978	.994
1	.535	.000	.707	.692	.943	.979	.002	.981	.983	.996

covariates of age and sex. For the application of the proposed model, we rescaled the continuous covariate (age) so that the min and max values were -1 and 1 . Putting covariates on the same scale is commonly done with GP priors, as it allows the use of the same length scale parameter in the GP kernel function across all covariates. We fitted the model to the TBI and HC data using the same parameter settings as outlined in the simulation study, with minor adjustments. Due to the different scale of the fMRI data, we set $\pi_\phi = 0.9$, $\sigma_\mu^2 = 0.01$, and adjusted the variance of the GP kernel function to be 0.5 to better identify the effects of covariates. This prior specification allows us to achieve the expected level of sparsity in brain networks and covariate space and to identify only the most significant connections, therefore retaining interpretability and focusing on the

most relevant brain regions and covariates. We are cautious to avoid over-fitting, which could lead to a model too finely tuned to the data. Therefore, the selection of the model is a balance between achieving a sparse, interpretable network and maintaining general applicability. When fitting the model, we use the ELBO to determine the convergence of the VB algorithm.

Estimated group-level networks for the two groups, TBI and HC, are shown in Figure 4, with shared edges in red. The number of selected edges is 419 for the HC group and 405 for the TBI group, with 182 overlapping edges between the two groups and with the HC group having more parent nodes than the TBI group in the left frontal lobe regions and right temporal gyrus regions. We conducted t tests to compare the estimated subject-level edge strengths between the




FIGURE 4 Application: Estimated connectograms for HC and TBI groups. Arcs indicate group-level edges (self-connections are not drawn) and node sizes are representative of the number of connected edges, including incoming, outgoing and self-connections. Edges in red are shared between the two groups. The number of selected edges is 419 for the HC group and 405 for the TBI group, with 182 overlapping edges between the 2 groups. The HC group has more parent nodes than the TBI group in the left frontal lobe regions and right temporal gyrus regions.




FIGURE 5 Application: (Top) Group-level edges affected by age. An arc between ROIs indicates that the group-level edge is selected and influenced by the covariate “age,” with more edges exhibiting age dependence in the HC group. Self-connections not shown. (Bottom) Subject-level edge strength estimates and the resulting estimated function of the covariate “age” values. Only four of the most interesting estimated effect functions are shown. Notably, the effective connectivity between the left and right supplementary motor areas in the HC group displays a nonlinear, inverted U-shaped response to increasing age, initially decreasing then increasing, with a pivot at around 11 years old.




FIGURE 6 Application: (Top) Group-level edges affected by sex. An arc between ROIs indicates that the group-level edge is selected and influenced by the covariate “sex,” with edges in red shared between the two groups and more sex-dependent edges in the HC group. Self-connections not shown. (Bottom) Subject-level edge strength estimates and the resulting estimated function of the covariate “sex” values. Only four of the most interesting estimated effect functions are shown, indicating how edge strengths vary with sex. For example, there is a stronger connectivity in the TBI group between the left triangular part of the inferior frontal gyrus and the left inferior temporal gyrus for males compared to females.

HC and TBI groups. Our analysis revealed that among the overlapping edges, 154 out of 182 showed significant differences between the two groups. Furthermore, for edges selected in one group but not in the other one, over 95% of these edges exhibited statistically significant differences in edge strength.

Examining the results further, networks displayed in Figures 5 and 6 show the selected edge strengths that are affected by age and sex, respectively. From these plots, we see that both age and sex have an influence on the edge strength function for both HC and TBI groups, with more edges in the HC group exhibiting a covariate dependence than in the TBI group. While fewer edges were evident for the TBI group compared to HC, in the TBI group age, and to a lesser extent sex, selectively influenced edges of the left putamen. At a more granular level, our inference reveals how edge strengths change as a function of each covariate. This is illustrated in the bottom parts of Figures 5 and 6, for a few of the most interesting estimated patterns of dependence. For example, in the TBI group, we observe stronger effective connectivity between the left triangular part of the inferior frontal gyrus and the left inferior temporal gyrus for males compared to females. A table with all selected edges and estimated covariate effect functions can be found in the Supplementary Material, together with plots of the estimated covariate effects for common edges in two groups. Broadly, we see that the model

TABLE 5 Application: Results of competing methods and VEVAR for numbers of edges that are influenced by covariates. Note that these results include self-connections within the networks.

Covariate effect selection	HC		TBI	
	Age	Sex	Age	Sex
GC-LASSO	7	23	12	7
GC-plsmselect	55	95	273	105
VEVAR-S1-SS	2	51	42	17
VEVAR	26	96	9	75

does well in capturing different functional shapes as well as capturing functions of categorical and continuous data.

We also compared our findings with three other methods: GC-LASSO, GC-plsmselect, and VEVAR-S1-SS. GC selected 1324 edges for the HC group and 1645 edges for the TBI group, with 101 and 171 overlapping edges with our methods, respectively. VEVAR-S1 selected 398 edges for the HC group and 362 edges for the TBI group, with 215 and 233 overlapping edges with our methods, respectively. We note that, as there is no ground truth from real data, the number of selected edges simply illustrates the relative sparsity of

TABLE 6 Application: Sensitivity analysis. Total numbers of selected edges and number of selected edges that are influenced by age and sex.

Sensitivity analysis						
	HC			TBI		
	Total	Age	Sex	Total	Age	Sex
Varying π_δ						
.1	419	26	96	405	9	75
.5	458	37	88	434	13	64
.9	527	26	92	507	10	65
Varying π_ϕ						
.1	353	24	89	282	13	53
.5	391	29	89	381	13	58
.9	419	26	96	405	9	75
Varying σ^2						
.1	324	13	128	440	14	146
.5	419	26	96	405	9	75
	243	28	47	179	7	36

the inferred networks. However, the primary focus of our method lies in its ability to incorporate significant covariates directly into the estimation of edge strengths, thereby enhancing our understanding of individual variability in brain connectivity. Results for covariate effect selections are shown in Table 5. Unlike GC, which provides a single estimation for the whole group, our approach allows for a nuanced analysis that can differentiate and track changes across various demographic and developmental stages. Furthermore, GC-LASSO and VEVAR-S1-SS are limited to identifying only linear effects of covariates, while VEVAR can handle nonlinear effects. Although GC-plsmselect is capable of handling nonlinear effects, the simulation study results showed that GC-plsmselect had a significantly high FPR, which may explain the larger number of selected edges with covariate effects seen in Table 5. Finally, it is also worth mentioning that the two-stage methods GC-LASSO, GC-plsmselect and VEVAR-S1-SS may select edges with covariate effects that were not selected in the first stage, which may not be relevant to the study.

3.2.1 | Sensitivity analysis

Results from the sensitivity analysis for the application conducted on the parameters π_δ , π_ϕ , and σ^2 are shown in Table 6. Consistent with the findings from the simulation study, an increase in the value of π_δ was observed to correspond to an increase in the number of selected edges. This parameter had no impact on the selection of covariate effects. An increase in the value of π_ϕ led to a rise in the number of selected effects for the binary covariate (sex), while having a limited effect on continuous covariate (age). This increase resulted in an overall increase in the number of selected edges, as we also observed in the simulation study. For the TBI group, our observations revealed a

negative relationship between σ^2 and the number of selected edges associated with covariates. On the other hand, for the HC group with a smaller sample size, no clear trends were observed for the total number of selected edges and the number of selected edges associated with age. This highlights the need for careful selection of the σ^2 parameter as it can be impacted by both the scales of covariates and the estimated edge strengths. While there are variations in the covariates selected under different parameter settings, the majority of the covariate effects that are identified tend to overlap, and the relationships between effective connectivity and the covariates remain relatively consistent.

4 | DISCUSSION

We have developed VEVAR, a novel extension of a Bayesian VAR model, and examined its ability to characterize group differences in effective network connectivity while simultaneously evaluating the potentially nonlinear impact of covariates on network edges. Using simulated data sets, VEVAR outperformed other competing approaches in terms of group-level edge selection as well as in covariate effect selection. We then applied VEVAR to assess its utility in characterizing subject and group-level connectivity network edge strengths using resting state fMRI data from a clinical sample of children with TBI and HC. Further, we estimated the effects of age and sex covariates on the group-level edge strengths. The groups differed in the distribution of parent nodes, which were fewer in children with TBI compared to controls. Age effects were largely nonlinear and influenced edges predominantly in healthy children, suggesting alteration in the relation of age with peak edge strength in children with TBI. Group-level edges were also affected by sex; effective connectivity strength was higher in more edges in males than in females. VEVAR has major potential to generate refined analyses of network-level data while clarifying potentially nonlinear relations with diverse covariates.

Unlike traditional methods that assume constant edge strengths across the group, VEVAR offers a flexible estimation of network edges that vary with covariate values, which is particularly advantageous in scenarios where edge strengths are potentially related to covariates in a nonlinear fashion. For group-level edge selection, VEVAR can model the edge strength as a dynamic function influenced by covariates, rather than as a static feature. In cases where competing methods might overlook an edge due to averaging effects (where the mean edge strength across subjects is zero), VEVAR retains the ability to detect these edges because it accounts for the covariate-driven variability in edge strength. In the simulation study, we compared performances to two-stage approaches that estimate the networks at the first stage and then select the covariates that explain the edge strengths at the second stage. We found that VEVAR does well in terms of all performance measures considered, for both group-level edge selection and covariate effect selection. For covariate effect selection, we saw again that VEVAR performs the best across all metrics, demonstrating superior accuracy and sensitivity in identifying the

true effects of covariates on network connectivity, likely due to its capability to model both linear and nonlinear relationships between covariates and edge strengths. One may raise a concern about comparing VEGAR to other methods for covariate selection, as it was shown that other models did not do well in selecting group-level networks. However, upon examining the subject level edge strength estimates, both GC and VEGAR-S1 gave estimates that were close to the underlying values used to simulate the data, minimizing our concern about this causing the difference in performance.

When applying VEGAR to data from the resting state fMRI study, our findings revealed that in the HC group the left frontal lobe regions and right temporal gyrus regions have more parent nodes than the corresponding nodes in the TBI group. These findings are consistent with those from other analytical approaches and reflect a consensus on the vulnerability of structural and functional connectivity in the frontal and temporal regions following TBI (Botchway et al., 2022; Johnson et al., 2011; Vaughn et al., 2022; Ware et al., 2022; Watson et al., 2019). Furthermore, the TBI group showed higher edge strength in the connectivity of the right temporal pole with the right-sided striatal and limbic structures, including the putamen, amygdala and hippocampus. This may reflect the reduced top-down regulation of limbic structures noted after TBI (Ewing-Cobbs et al., 2019). Regional hyperconnectivity has been reported in children and adults across the spectrum of TBI severity and chronicity (Caeyenberghs et al., 2017). Hyperconnectivity, especially in the default mode network, has been linked to better task performance in several studies (Grossner et al., 2019; Lancaster et al., 2019; Palacios et al., 2017; Stephens et al., 2018; Venkatesan & Hillary, 2019). Although hyperconnectivity is often viewed as a compensatory mechanism, the potential long-term metabolic costs have not been established (Hillary & Grafman, 2017).

Both age and sex significantly impacted edge strength in HC and TBI groups. In particular, we identified a higher count of edges in the HC group. Edges in the HC group that were influenced by age primarily connected left inferior temporal to bilateral frontal and temporal-parietal regions. These anterior and posterior association areas are among the later-developing brain regions. Reduced connectivity in these edges following TBI may reflect vulnerability of these connections among regions that develop rapidly during early to late adolescence. We also identified a distinct influence of these demographic factors, particularly age, on the left putamen edges in the TBI group. DTI tractography has previously identified the putamen as a structural hub in children with a history of TBI but not in an age matched HC group (Caeyenberghs et al., 2012). Moreover, frontal-striatal (including putamen and caudate) network disruption is evident following pediatric TBI (Watson et al., 2019). Of the few studies examining functional connectivity changes after pediatric TBI, only one evaluated the influence of age and sex on connectivity metrics. In a sample of children with mild TBI, Onicas et al. (2024) found that by 3–6 months after injury, global clustering was reduced in superior parietal and occipital regions in adolescents as well as in females experiencing persistent symptoms relative to orthopedic injury. Although our sample contained a broad spectrum of TBI severity, our findings converge to

indicate vulnerability of adolescents and females to disruption of specific connectivity metrics. Current findings elucidate the functional implication of striatal injury and characterize how patient characteristics, in this case, age or sex, might strengthen or attenuate the impact of injury on striatal brain networks.

The effective connectivity between the left supplementary motor area and right supplementary motor area for the HC group have shown an interesting nonlinear, inverted U shape as age increases, first decreasing and then increasing. The trends of the estimated effective connectivities with respect to age change at around 11 years old for the HC group while the trends change at around 13 years old for the TBI group. Importantly, our approach does not constrain associations between covariates and edge strength to a nonlinear function but instead allows for discovery of the pattern of influence of the covariate to be detected. Developmental profiles showing inverted U-shapes are consistently reported in brain morphometry via regional volume or thickness of cortical gray matter and are broadly interpreted to reflect initial overly abundant synaptic production and the subsequent synaptic pruning (Giedd et al., 1999, 2006; Gogtay et al., 2004). Although functional profiles corresponding to an inverted U during childhood and adolescence are established in the literature, it is to a lesser degree relative to the morphometric literature. Notably, the pattern of inverted U in brain functioning, even in development, is most often associated with an increase in expertise or skill (often in language/reading domains), such that learning initially triggers an increase (the rise in the inverted U) with the falling end of the U associated with experience, and ultimately expertise, see Perkins and Jiang (2019).

Despite its importance as a biological variable, few imaging studies have examined how sex may affect connectivity metrics. In our results, many group-level edges were influenced by sex. Except for bilateral temporal pole regions, very few edges were shared in females and males. Relations also differed by group. In the TBI group, females had higher effective connectivity than males in both temporal poles. Males had a unique pattern wherein edges originating in multiple bilateral structures converged in the left putamen. Effective connectivity in males in the HC group was higher in multiple bilateral regions impacting bilateral mid- and inferior temporal regions while females had higher connectivity in edge strengths affecting left temporal regions. We note that there could be other confounding factors, such as head size, influencing the observed sex differences in brain connectivity. Given that we do not have head size data, we cannot conclusively link the correlations directly to sex. Therefore, our results should be viewed as indicative of a correlation between sex and specific brain connectivity patterns, rather than conclusive evidence of sex-based differences. Additional studies should examine the relation of connectivity metrics and covariate effects with cognitive and behavioral outcomes to more fully elucidate specific brain-behavior relations.

There are several aspects of the proposed model that allow the practitioner added flexibility. For example, as previously discussed, the GP kernel can differ across datasets, edges, or covariates. In both the simulation and case study application, the covariates used were

subject-specific measurements, constant across edges and nodes. The model, however, could easily be adjusted to accommodate covariates specific to a subset of edges or groups. Alternatively, covariates whose values vary at each node or edge could be considered. Finally, while the current model formulation relies on the group subject memberships to be known, unsupervised settings could also be considered via nonparametric priors such as the Dirichlet process. This would lead to broader applications and potentially further insights into the data.

5 | CONCLUSION

In this article, we have proposed VEVAR, an analytical approach for estimating brain connectivity networks that accounts for subject heterogeneity. We have employed a hierarchical Bayesian VEVAR model to identify connectivity networks for different groups of subjects, allowing for dependence of the connection strength on covariate values. We have designed a sparse prior to identify key connections within each group and subsets of edges with strengths affected by the covariates. Our novel nonparametric spike-and-slab prior includes a “slab” portion as a function mapping possible covariate values to coefficient values. Additionally, we have assumed a weighted mixture of Gaussian process priors on this function, to allow modeling of (possibly) nonlinear effects and selection of relevant covariates. We have estimated the proposed model using VI, which has allowed for application to large-scale data. The model has been shown to perform better than competing two-stage approaches on simulated data, in both network discovery and covariate selection. We have applied our method to resting-state fMRI data on children with a history of TBI and HCs to estimate group-level connectivity networks and evaluate whether specific edge strengths are affected by age and sex. The identified differences in effective connectivity network patterns between TBI and HC groups, along with the influence of age and sex, have provided valuable insights into the complexity of brain functioning after TBI. Furthermore, TBI as a case study application, where effects of age are evaluated within effective connectivity networks, highlights the promise of the proposed approach for evaluating developmental and aging effects in a breadth of disorders impacting the brain. For future studies, extensions to longitudinal data may provide valuable insights into connectivity changes over an extended period after TBI.

Broadly, VEVAR contributes to the current understanding of brain mechanisms by identifying functionally connected brain regions, enabling researchers to characterize information flow within and between brain networks. Moreover, VEVAR determines the directionality of functional connections between brain regions, offering insights into how specific cognitive functions like memory, attention, language, and decision-making are temporally implemented within overlapping brain networks. In the context of development, VEVAR advances understanding by tracking the development of neural networks from infancy through adolescence and into adulthood. By observing changes in connectivity patterns over time,

researchers can identify sensitive developmental periods and map the integration of different brain regions. VEVAR in brain connectivity research has been implemented to provide crucial insights into altered neural circuits associated with injuries such as TBI (Kook et al., 2021; Vaughn et al., 2022). This method is poised for similar use in neurodevelopmental disorders like ASD (Hanson et al., 2013; Rolls et al., 2020), dyslexia (Di Pietro et al., 2023), or attention deficit hyperactivity disorder (Kumar et al., 2021), where effective connectivity findings suggest that differences in the connectivity patterns in individuals with developmental disorders compared to typically developing individuals. Moreover, VEVAR facilitates the comparison of connectivity profiles between typically developing individuals and those affected by these disorders while controlling for nuisance variables. This enables analyses that may identify aberrant developmental trajectories and potential biomarkers for early diagnosis and intervention.

ACKNOWLEDGMENTS

Y.R., M.V., and C.B.P. were supported in part by grants DMS-2113602 and DMS-2113557 from the National Science Foundation. D.M.D. and L.E.C. were funded in part by National Institutes of Health R01 NS046308. The content is solely the responsibility of the authors and does not necessarily represent the official views of the granting institutes.

CONFLICT OF INTEREST STATEMENT

The authors declare no competing interests.

DATA AVAILABILITY STATEMENT

The processed imaging data analyzed in this article can be downloaded from the MATLAB BVAR package at https://github.com/marinavannucci/BVAR_connect. The subject-level covariate information that supports the findings of this study is available from the corresponding author upon reasonable request. Python code implementing the VEVAR model described in this article can be downloaded from GitHub at <https://github.com/YangfanR/VEVAR-fmri>.

ORCID

Marina Vannucci  <https://orcid.org/0000-0002-7360-5321>

REFERENCES

- Arnold, A., Liu, Y., & Abe, N. (2007). Temporal causal modeling with graphical Granger methods. In *Proceedings of the 13th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 66–75.
- Blei, D. M., Kucukelbir, A., & McAuliffe, J. D. (2017). Variational inference: A review for statisticians. *Journal of the American Statistical Association*, 112(518), 859–877.
- Botchway, E., Kooper, C. C., Pouwels, P. J., Bruining, H., Engelen, M., Oosterlaan, J., & Königs, M. (2022). Resting-state network organisation in children with traumatic brain injury. *Cortex*, 154, 89–104.
- Caeyenberghs, K., Leemans, A., De Decker, C., Heitger, M., Drijkoningen, D., Vander Linden, C., Snaert, S., & Swinnen, S. P. (2012). Brain connectivity and postural control in young traumatic brain injury patients: A diffusion MRI based network analysis. *NeuroImage: Clinical*, 1(1), 106–115.

- Caeyenberghs, K., Verhelst, H., Clemente, A., & Wilson, P. H. (2017). Mapping the functional connectome in traumatic brain injury: What can graph metrics tell us? *NeuroImage*, 160, 113–123.
- Carbonetto, P., & Stephens, M. (2012). Scalable variational inference for Bayesian variable selection in regression, and its accuracy in genetic association studies. *Bayesian Analysis*, 7(1), 73–108.
- Chiang, S., Guindani, M., Yeh, H. J., Haneef, Z., Stern, J. M., & Vannucci, M. (2017). Bayesian vector autoregressive model for multi-subject effective connectivity inference using multi-modal neuroimaging data. *Human Brain Mapping*, 38(3), 1311–1332.
- Cleveland, W. S., & Grosse, E. (1991). Computational methods for local regression. *Statistics and Computing*, 1(1), 47–62.
- Dance, H., & Paige, B. (2022). Fast and scalable spike and slab variable selection in high-dimensional Gaussian processes. Proceedings of the 25th International Conference on Artificial Intelligence and Statistics, PMLR 151:7976–8002.
- Di Pietro, S. V., Willinger, D., Frei, N., Lutz, C., Coraj, S., Schneider, C., Stämpfli, P., & Brem, S. (2023). Disentangling influences of dyslexia, development, and reading experience on effective brain connectivity in children. *NeuroImage*, 268, 119869.
- Ewing-Cobbs, L., DeMaster, D., Watson, C. G., Prasad, M. R., Cox, C. S., Kramer, L. A., Fischer, J. T., Duque, G., & Swank, P. R. (2019). Post-traumatic stress symptoms after pediatric injury: Relation to pre-frontal limbic circuitry. *Journal of Neurotrauma*, 36(11), 1738–1751.
- Friston, K. J. (1994). Functional and effective connectivity in neuroimaging: A synthesis. *Human Brain Mapping*, 2(1–2), 56–78.
- Friston, K. J. (2011). Functional and effective connectivity: A review. *Brain Connectivity*, 1(1), 13–36.
- Friston, K. J., Harrison, L., & Penny, W. (2003). Dynamic causal modelling. *NeuroImage*, 19(4), 1273–1302.
- Gelman, A., Carlin, J. B., Stern, H. S., Dunson, D. B., Vehtari, A., & Rubin, D. B. (2013). *Bayesian data analysis*. CRC Press.
- George, E. I., & McCulloch, R. E. (1997). Approaches for Bayesian variable selection. *Statistica Sinica*, 7(2), 339–373.
- Ghosal, I., & Kormaksson, M. (2019). plsmselect: Linear and smooth predictor modelling with penalisation and variable selection. R package version 0.2.0.
- Giedd, J. N., Blumenthal, J., Jeffries, N. O., Castellanos, F. X., Liu, H., Zijdenbos, A., Paus, T., Evans, A. C., & Rapoport, J. L. (1999). Brain development during childhood and adolescence: A longitudinal MRI study. *Nature Neuroscience*, 2(10), 861–863.
- Giedd, J. N., Clasen, L. S., Lenroot, R., Greenstein, D., Wallace, G. L., Ordaz, S., Molloy, E. A., Blumenthal, J. D., Tossell, J. W., Stayer, C., Samango-Sprouse, C. A., Shen, D., Davatzikos, C., Merke, D., & Chrousos, G. P. (2006). Puberty-related influences on brain development. *Molecular and Cellular Endocrinology*, 254, 154–162.
- Gogtay, N., Giedd, J. N., Lusk, L., Hayashi, K. M., Greenstein, D., Vaituzis, A. C., Nugent, T. F., III, Herman, D. H., Clasen, L. S., Toga, A. W., Rapoport, J. L., & Thompson, P. M. (2004). Dynamic mapping of human cortical development during childhood through early adulthood. *Proceedings of the National Academy of Sciences*, 101(21), 8174–8179.
- Gorrostieta, C., Fiecas, M., Ombao, H., Burke, E., & Cramer, S. (2013). Hierarchical vector auto-regressive models and their applications to multi-subject effective connectivity. *Frontiers in Computational Neuroscience*, 7, 159.
- Gorrostieta, C., Ombao, H., Bédard, P., & Sanes, J. N. (2012). Investigating brain connectivity using mixed effects vector autoregressive models. *NeuroImage*, 59(4), 3347–3355.
- Granger, C. W. (1969). Investigating causal relations by econometric models and cross-spectral methods. *Econometrica: Journal of the Econometric Society*, 37, 424–438.
- Grossner, E. C., Bernier, R. A., Brenner, E. K., Chiou, K. S., Hong, J., & Hillary, F. G. (2019). Enhanced default mode connectivity predicts metacognitive accuracy in traumatic brain injury. *Neuropsychology*, 33(7), 922–933.
- Guha, S., & Guhaniyogi, R. (2021). Bayesian supervised clustering of undirected networks with cluster specific inference on significant nodes and edges related to predictors. Technical report.
- Hanson, C., Hanson, S. J., Ramsey, J., & Glymour, C. (2013). Atypical effective connectivity of social brain networks in individuals with autism. *Brain Connectivity*, 3(6), 578–589.
- Hastie, T., & Tibshirani, R. (1993). Varying-coefficient models. *Journal of the Royal Statistical Society: Series B (Methodological)*, 55(4), 757–779.
- Hillary, F. G., & Grafman, J. H. (2017). Injured brains and adaptive networks: The benefits and costs of hyperconnectivity. *Trends in Cognitive Sciences*, 21(5), 385–401.
- Johnson, C. P., Juraneck, J., Kramer, L. A., Prasad, M. R., Swank, P. R., & Ewing-Cobbs, L. (2011). Predicting behavioral deficits in pediatric traumatic brain injury through uncinate fasciculus integrity. *Journal of the International Neuropsychological Society*, 17(4), 663–673.
- Kolar, M., Song, L., Ahmed, A., Xing, E. P. (2010). Estimating time-varying networks. *Annals of Applied Statistics*, 4(1), 94–123.
- Kook, J. H., Vaughn, K. A., DeMaster, D. M., Ewing-Cobbs, L., & Vannucci, M. (2021). BVAR-connect: A variational Bayes approach to multi-subject vector autoregressive models for inference on brain connectivity networks. *Neuroinformatics*, 19, 39–56.
- Kumar, U., Arya, A., & Agarwal, V. (2021). Neural network connectivity in ADHD children: An independent component and functional connectivity analysis of resting state fMRI data. *Brain Imaging and Behavior*, 15, 157–165.
- Kundu, S., Ming, J., Nocera, J., & McGregor, K. M. (2021). Integrative learning for population of dynamic networks with covariates. *NeuroImage*, 236, 118181.
- Lancaster, K., Venkatesan, U. M., Lengsfelder, J., & Genova, H. M. (2019). Default mode network connectivity predicts emotion recognition and social integration after traumatic brain injury. *Frontiers in Neurology*, 10, 471442.
- Li, B., Chun, H., & Zhao, H. (2012). Sparse estimation of conditional graphical models with application to gene networks. *Journal of the American Statistical Association*, 107(497), 152–167.
- Li, J., Wang, Z. J., Palmer, S. J., & McKeown, M. J. (2008). Dynamic Bayesian network modeling of fMRI: A comparison of group-analysis methods. *NeuroImage*, 41(2), 398–407.
- McIntosh, A., & Gonzalez-Lima, F. (1994). Structural equation modeling and its application to network analysis in functional brain imaging. *Human Brain Mapping*, 2(1–2), 2–22.
- Onicas, A. I., Deighton, S., Yeates, K. O., Bray, S., Graff, K., Abdeen, N., Beauchamp, M. H., Beaulieu, C., Bjornson, B., Craig, W., Dehaes, M., Deschenes, S., Doan, Q., Freedman, S. B., Goodyear, B. G., Gravel, J., Lebel, C., Ledoux, A. A., Zemek, R., & Ware, A. L. (2024). Longitudinal functional connectome in pediatric concussion: An advancing concussion assessment in pediatrics study. *Journal of Neurotrauma*, 41(5–6), 587–603.
- Palacios, E. M., Yuh, E. L., Chang, Y.-S., Yue, J. K., Schnyer, D. M., Okonkwo, D. O., Valadka, A. B., Gordon, W. A., Maas, A. I., Vassar, M., et al. (2017). Resting-state functional connectivity alterations associated with six-month outcomes in mild traumatic brain injury. *Journal of Neurotrauma*, 34(8), 1546–1557.
- Perkins, K., & Jiang, X. (2019). Neuroimaging and reading comprehension. *Journal of Interdisciplinary Studies in Education*, 8(2), 74–94.
- Rajapakse, J. C., & Zhou, J. (2007). Learning effective brain connectivity with dynamic Bayesian networks. *NeuroImage*, 37(3), 749–760.
- Ray, K., & Szabó, B. (2021). Variational Bayes for high-dimensional linear regression with sparse priors. *Journal of the American Statistical Association*, 117(539), 1–12.
- Roebroek, A., Formisano, E., & Goebel, R. (2005). Mapping directed influence over the brain using granger causality and fMRI. *NeuroImage*, 25(1), 230–242.

- Rolls, E. T., Zhou, Y., Cheng, W., Gilson, M., Deco, G., & Feng, J. (2020). Effective connectivity in autism. *Autism Research*, 13(1), 32–44.
- Scheffler, A. W., Telesca, D., Sugar, C. A., Jeste, S., Dickinson, A., DiStefano, C., & Entürk, D. (2019). Covariate-adjusted region-referenced generalized functional linear model for EEG data. *Statistics in Medicine*, 38(30), 5587–5602.
- Stephens, J. A., Salorio, C. F., Barber, A. D., Risen, S. R., Mostofsky, S. H., & Suskauer, S. J. (2018). Preliminary findings of altered functional connectivity of the default mode network linked to functional outcomes one year after pediatric traumatic brain injury. *Developmental Neuropsychology*, 21(7), 423–430.
- Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society. Series B (Methodological)*, 58(1), 267–288.
- Titsias, M. K., & Lázaro-Gredilla, M. (2011). Spike and slab variational inference for multi-task and multiple kernel learning. In *Advances in neural information processing systems* (pp. 2339–2347). The MIT Press.
- Valdés-Sosa, P. A., Sánchez-Bornot, J. M., Lage-Castellanos, A., Vega-Hernández, M., Bosch-Bayard, J., Melie-García, L., & Canales-Rodríguez, E. (2005). Estimating brain functional connectivity with sparse multivariate autoregression. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 360(1457), 969–981.
- Van Loan, C. F. (2000). The ubiquitous Kronecker product. *Journal of Computational and Applied Mathematics*, 123(1–2), 85–100.
- Vannucci, M. (2021). Discrete spike-and-slab priors: Models and computational aspects. In *Handbook of Bayesian variable selection* (pp. 3–24). Chapman and Hall/CRC.
- Vaughn, K., DeMaster, D., Kook, J., Vannucci, M., & Ewing-Cobbs, L. (2022). Effective connectivity in the default mode network after pediatric traumatic brain injury. *European Journal of Neuroscience*, 55(1), 318–336.
- Venkatesan, U. M., & Hillary, F. G. (2019). Functional connectivity within lateral posterior parietal cortex in moderate to severe traumatic brain injury. *Neuropsychology*, 33(6), 893–910.
- Wang, Y., Yan, G., Wang, X., Li, S., Peng, L., Tudorascu, D., & Zhang, T. (2023). A variational Bayesian approach to identifying whole-brain directed networks with fMRI data. *Annals of Applied Statistics*, 17(1), 518–538.
- Ware, A. L., Yeates, K. O., Geeraert, B., Long, X., Beauchamp, M. H., Craig, W., Doan, Q., Freedman, S. B., Goodyear, B. G., Zemek, R., Lebel, C., & Pediatric Emergency Research Canada A-CAP Study Team. (2022). Structural connectome differences in pediatric mild traumatic brain and orthopedic injury. *Human Brain Mapping*, 43(3), 1032–1046.
- Watson, C. G., DeMaster, D., & Ewing-Cobbs, L. (2019). Graph theory analysis of DTI tractography in children with traumatic injury. *NeuroImage: Clinical*, 21, 101673.
- Wen, X., Rangarajan, G., & Ding, M. (2013). Is Granger causality a viable technique for analyzing fMRI data? *PLoS One*, 8(7), e67428.
- Whitfield-Gabrieli, S., Nieto-Castanon, A., & Ghosh, S. (2011). Artifact detection tools (ART). Cambridge, MA. Release Version, 7(19):11.
- Williams, C. K., & Rasmussen, C. E. (2006). *Gaussian processes for machine learning*. MIT Press.
- Wu, X., Zou, Q., Hu, J., Tang, W., Mao, Y., Gao, L., Zhu, J., Jin, Y., Wu, X., Lu, L., Zhang, Y., Zhang, Y., Dai, Z., Gao, J. H., Weng, X., Zhou, L., Northoff, G., Giacino, J. T., He, Y., & Yang, Y. (2015). Intrinsic functional connectivity patterns predict consciousness level and recovery outcome in acquired brain injury. *Journal of Neuroscience*, 35(37), 12932–12946.
- Yin, J., & Li, H. (2011). A sparse conditional Gaussian graphical model for analysis of genetical genomics data. *The Annals of Applied Statistics*, 5(4), 2630–2650.
- Zhao, Y., Wang, B., Mostofsky, S. H., Caffo, B. S., & Luo, X. (2021). Covariate assisted principal regression for covariance matrix outcomes. *Biostatistics*, 22(3), 629–645.
- Zhou, S., Lafferty, J., & Wasserman, L. (2010). Time varying undirected graphs. *Machine Learning*, 80(2), 295–319.
- Zhou, Y., Milham, M. P., Lui, Y. W., Miles, L., Reaume, J., Sodickson, D. K., Grossman, R. I., & Ge, Y. (2012). Default-mode network disruption in mild traumatic brain injury. *Radiology*, 265(3), 882–892.

SUPPORTING INFORMATION

Additional supporting information can be found online in the Supporting Information section at the end of this article.

How to cite this article: Ren, Y., Osborne, N., Peterson, C. B., DeMaster, D. M., Ewing-Cobbs, L., & Vannucci, M. (2024). Bayesian varying-effects vector autoregressive models for inference of brain connectivity networks and covariate effects in pediatric traumatic brain injury. *Human Brain Mapping*, 45(10), e26763. <https://doi.org/10.1002/hbm.26763>