

Bayesian network-guided sparse regression with flexible varying effects

Yangfan Ren¹, Christine B. Peterson², Marina Vannucci^{1,*}

¹Department of Statistics, Rice University, Houston, TX 77005, United States, ²Department of Biostatistics, The University of Texas MD Anderson Cancer Center, Houston, TX 77030, United States

*Corresponding author: Marina Vannucci, Department of Statistics, Rice University, 6100 Main Street, Houston, TX 77005, United States (marina@rice.edu).

ABSTRACT

In this paper, we propose Varying Effects Regression with Graph Estimation (VERGE), a novel Bayesian method for feature selection in regression. Our model has key aspects that allow it to leverage the complex structure of data sets arising from genomics or imaging studies. We distinguish between the *predictors*, which are the features utilized in the outcome prediction model, and the subject-level *covariates*, which modulate the effects of the predictors on the outcome. We construct a varying coefficients modeling framework where we infer a network among the predictor variables and utilize this network information to encourage the selection of related predictors. We employ variable selection spike-and-slab priors that enable the selection of both network-linked predictor variables and covariates that modify the predictor effects. We demonstrate through simulation studies that our method outperforms existing alternative methods in terms of both feature selection and predictive accuracy. We illustrate VERGE with an application to characterizing the influence of gut microbiome features on obesity, where we identify a set of microbial taxa and their ecological dependence relations. We allow subject-level covariates, including sex and dietary intake variables to modify the coefficients of the microbiome predictors, providing additional insight into the interplay between these factors.

KEYWORDS: Bayesian variable selection; Gaussian process prior; graphical model; spike-and-slab prior; varying coefficient model.

1 INTRODUCTION

In this paper, we propose a novel Bayesian hierarchical regression model that enables the selection of both network-linked predictor variables and covariates that modify the predictor effects. Here, we distinguish between the *predictors*, which are the features utilized in the outcome prediction model, and the subject-level *covariates*, which modulate the effects of the predictors. Our method is motivated by applications to multivariate data sets arising from genomic and neuroimaging studies, where the observed predictors are linked by metabolic or functional networks. As an illustration of the utility of our method, we consider an application to a microbiome data set examining the interplay between the microbiome, obesity, and subject-level covariates (Wu et al., 2011). In this context, the dependence network among the predictor variables describes ecological relationships between microorganisms inhabiting the same niche (Kurtz et al., 2015). Subject-level variables, including sex and dietary intake variables, modify the influence of the microbiome features by regulating their activity and their effects on the host (Leeming et al., 2021). This effect may be partially driven by the production of dietary metabolites (Sonnenburg and Bäckhed, 2016); however, many aspects of these relations are not completely understood.

Our proposed model builds upon the framework of varying coefficient models (Cleveland and Grosse, 1991; Hastie and Tibshirani, 1993). Varying-coefficient models relax the assump-

tion of linear effects in classical regression by allowing predictor effects to depend on factors that may modify their effects. This class of model has been extended to allow for high-dimensional covariates using spline- and tree-based approaches (Marx, 2009; Bürgin and Ritschard, 2015). Additionally, Bayesian varying-coefficient models have been developed for scenarios with spatial or temporal dependence (Reich et al., 2010; Scheipl et al., 2012). These approaches generally focus on selecting either the main predictors or the modifying covariates. The simultaneous selection of predictors and covariates was first introduced in the pliable lasso (Tibshirani and Friedman, 2020) and extended by Kim et al. (2021) to account for grouping structure among the predictors via a weighted hierarchical penalty. In the Bayesian framework, Ni et al. (2019) proposed a varying-sparsity regression model, which allows for subject-specific predictor selection and coefficient values. However, none of these approaches utilize information on the network among the predictor variables.

In our formulation of the varying coefficient modeling framework, we enable flexibility in the predictor effects by utilizing a Gaussian process prior, which allows the model coefficients to vary smoothly as a function of the observed covariates. Furthermore, to achieve model sparsity, we rely on spike-and-slab priors for the selection of both the predictor and covariate effects. Our prior formulation allows to infer a network among the predictors and to utilize this information to encourage the selection of network-linked predictors. In order to infer the

predictors network, we adopt the prior formulation proposed by Wang (2015), which imposes a mixture of normals on the off-diagonal elements of the precision matrix, along with an efficient blocked Gibbs sampling scheme. This approach directly represents edge selection, with respect to alternative shrinkage priors, such as the Bayesian graphical lasso (Wang, 2012) or the graphical horseshoe prior (Li et al., 2019). For posterior inference, we design a stochastic search Markov chain Monte Carlo (MCMC) that requires careful consideration of how to handle the changing dimensions of the parameter space, cleverly dealing with the multiple layers of selection while ensuring good mixing. We also look into prediction. We call our proposed method VERGE (Varying Effects Regression with Graph Estimation).

We demonstrate through simulation studies that our method outperforms existing alternative methods in terms of both feature selection and predictive accuracy. We illustrate VERGE with an application to characterizing the influence of gut microbiome features on obesity, where we identify a set of microbial taxa and their ecological dependence relations. Our findings highlight bacterial genera with both protective and detrimental effects, and provide insight into how these effects are modulated by dietary intake and biological sex.

Section 2 details our proposed VERGE approach, and Section 3 discusses posterior inference and prediction. In Section 4, we present simulation studies and comparisons. Section 5 contains an application to microbiome data and Section 6 conclusions and discussion.

2 METHODS

2.1 Varying-effects regression model

Let Y_i denote the observed response variable, and $\mathbf{X}_i = (X_{i1}, \dots, X_{ip})$ denote the P -dimensional vector of predictors for subject $i = 1, \dots, n$. We assume a joint distribution for the random variables (Y, \mathbf{X}) . As in Peterson et al. (2016), our joint distribution (Y, \mathbf{X}) can be factorized as $(Y, \mathbf{X}) = f(Y|\mathbf{X}) \cdot f(\mathbf{X})$, where $f(Y|\mathbf{X})$ is a regression model and $f(\mathbf{X})$ is a multivariate normal distribution. Both the response variable Y_i and predictors \mathbf{X}_i are assumed to be centered. Additionally, for each subject i , we also observe a K -dimensional covariate vector $\mathbf{Z}_i = (Z_{i1}, \dots, Z_{iK})$. Our proposed model allows for the effects of the predictors on the outcome to depend on specific covariates via a varying-effects regression model formulation where the coefficients of X_{ij} change based on the value of \mathbf{Z}_i as

$$y_i = \sum_{j=1}^p X_{ij} \mu_j(\mathbf{Z}_i) + \varepsilon_i, \quad \varepsilon_i \sim \mathcal{N}(0, \tau^2), \quad (1)$$

with $\mu_j(\mathbf{Z}_i)$ an unknown function of \mathbf{Z}_i , and ε_i 's iid white noise with variance parameter, τ^2 , on which we assume a standard conjugate inverse gamma prior as $\tau^2 \sim \text{IG}(a_0, b_0)$. Equation 1 represents a full model with no predictor selection; in the next section, we introduce our prior formulation that enables model sparsity.

2.2 Priors for variable and graph selection

Our model builds on the Bayesian variable selection approach originally proposed by Kuo and Mallick (1998). We innovate on

this framework in 2 key regards: We utilize network relations to link the probability of predictor selection and allow the non-zero coefficients to vary as a smooth function of the covariate values. To represent the predictor selection, we introduce a set of latent indicator variables, $\boldsymbol{\gamma} = (\gamma_1, \dots, \gamma_p)$, and write $\mu_j(\mathbf{Z}_i) = \gamma_j \beta_j(\mathbf{Z}_i)$, assuming a priori independence between the indicators γ_j and the effects $\beta_j(\mathbf{Z}_i)$. If $\gamma_j = 0$, then X_j has no effect on the response, as in the discrete spike-and-slab prior formulation (Vannucci, 2021). Thus, the varying-effects regression model in Equation 1 can be written as

$$y_i = \sum_{j=1}^p X_{ij} \gamma_j \beta_j(\mathbf{Z}_i) + \varepsilon_i, \quad \varepsilon_i \sim \mathcal{N}(0, \tau^2). \quad (2)$$

In this model formulation, γ_j acts as an indicator for the relevance of the predictor X_j . When $\gamma_j = 1$, it implies that the corresponding $\beta_j(\mathbf{z})$ is included in the model; otherwise, $\beta_j(\mathbf{z})$ is effectively 0. Although γ_j and $\beta_j(\mathbf{z})$ are a priori independent, the MCMC sampling process captures their relationship, with γ_j more likely to be 1 when the Gaussian Process (GP) realization of $\beta_j(\mathbf{z})$ (see Section 2.3) is significantly different from zero. In our setting, this model form allows us to integrate out the coefficients $\beta_j(\mathbf{z})$. This integration makes sampling more efficient than alternative spike-and-slab prior formulations (George and McCulloch, 1997).

Given our focus on both selecting a subset of explanatory predictors and understanding the interconnections among these variables, we rely on the Gaussian graphical model to infer a network among the predictors and assume that \mathbf{X}_i follow a multivariate normal distribution

$$\mathbf{X}_i \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Omega}^{-1}), \quad (3)$$

where $\mathbf{0}$ is a P -vector of 0s, and $\boldsymbol{\Omega}$ is the precision matrix, which can be used to represent the conditional dependencies among predictors. Non-zero off-diagonal entries ω_{ij} in $\boldsymbol{\Omega}$ correspond to conditional dependence relations between the corresponding predictors, while $\omega_{ij} = 0$ indicates that predictors i and j are conditionally independent given the remaining variables. To identify a sparse set of dependence relations, we place a mixture prior on the entries in $\boldsymbol{\Omega}$ as proposed by Wang (2015):

$$p(\boldsymbol{\Omega} | \theta) = \{C(\theta)\}^{-1} \prod_{i < j} \left\{ (1 - \pi) \mathcal{N}(\omega_{ij} | 0, v_0^2) + \pi \mathcal{N}(\omega_{ij} | 0, v_1^2) \right\} \prod_i \text{Exp}(\omega_{ii} | \frac{\lambda}{2}) I_{(\boldsymbol{\Omega} \in M^+)}, \quad (4)$$

where $\theta = \{v_0, v_1, \lambda, \pi\}$ represents the set of prior hyperparameters, $C(\theta)$ is a normalizing constant, and π indicates the prior probability of edge selection. The prior distribution on the off-diagonal elements of $\boldsymbol{\Omega}$ is a mixture of normals, with v_0 and v_1 being set small and large, respectively. This allows a clear separation between selected edges, with values significantly different from zero, and non-selected edges, where ω_{ij} is close to zero. The diagonal elements follow an exponential distribution with parameter $\frac{\lambda}{2}$. The final term $I_{(\boldsymbol{\Omega} \in M^+)}$ expresses the constraint that $\boldsymbol{\Omega}$ belongs to the cone of symmetric positive definite matrices.

The graph can also be represented using a set of binary latent variables $G = \{g_{ij}\}_{i < j} \in \{0, 1\}$, where $g_{ij} = 1$ indicates the

presence of edge (i, j) in the graph G . The model in Equation 4 can then be expressed as the hierarchical model

$$\begin{aligned}
 p(\boldsymbol{\Omega} \mid G, \nu_0, \nu_1, \lambda) &= \{C(G, \nu_0, \nu_1, \lambda)\}^{-1} \\
 &\times \prod_{i < j} \mathcal{N}(\omega_{ij} \mid 0, \nu_{g_{ij}}^2) \prod_i \text{Exp}(\omega_{ii} \mid \frac{\lambda}{2}), \\
 p(G \mid \theta) &= \{C(G, \theta)\}^{-1} C(G, \nu_0, \nu_1, \lambda) \\
 &\times \prod_{i < j} \{\pi^{g_{ij}} (1 - \pi)^{1-g_{ij}}\}. \tag{5}
 \end{aligned}$$

For standardized data, Wang (2015) recommends setting $\pi = \frac{2}{p-1}$ and $\lambda = 1$, noting that edge selection tends to be insensitive to the choice of λ . For ν_0 and ν_1 , Wang (2015) observes that stable MCMC convergence is achieved with $\nu_0 \geq 0.01$ and $\nu_1 \leq 10$. Additional sensitivity analysis results are provided in Wang (2015).

Instead of using the conventional approach of employing an independent Bernoulli prior for the variable selection indicator $\boldsymbol{\gamma}$, we adopt the model proposed by Peterson et al. (2016), which utilizes a Markov random field prior to link the selection of predictors according to their relations in the graph G as

$$p(\boldsymbol{\gamma} \mid G) \propto \exp(a\mathbf{1}'\boldsymbol{\gamma} + b\boldsymbol{\gamma}'G\boldsymbol{\gamma}), \tag{6}$$

where a and b are scalar hyperparameters. This prior connects the variable inclusion to the inference of the dependence network, encouraging the selection of predictors that are connected with other relevant predictors. The parameter $a < 0$ controls the prior probability of selecting a variable without accounting for information in the graph, while the parameter b controls the extent to which a variable's inclusion probability is influenced by the inclusion of connected variables in the graph. As discussed in Li and Zhang (2010), b should be carefully selected, as high values result in very dense models, a phenomenon known as phase transition.

2.3 Priors on predictor effects and covariate selections

To allow for covariates to modulate the strength of the predictor effects, we employ the Gaussian process prior framework proposed by Savitsky et al. (2011). Specifically, for each predictor, the prior distribution of $\boldsymbol{\beta}_j(\mathbf{Z})$ is defined by a Gaussian process regression model

$$\boldsymbol{\beta}_j(\mathbf{Z}) = f_j(\mathbf{Z}) + \boldsymbol{\delta}_j, \tag{7}$$

where \mathbf{Z} is an $n \times K$ matrix, and $f_j(\mathbf{Z})$ is a realization of a Gaussian process $f(\mathbf{Z}) \sim \mathcal{N}(0, \mathbf{C}_j)$. The ‘‘jitter’’ term $\boldsymbol{\delta}_j$ is distributed as $\mathcal{N}(0, \frac{1}{r_j}\mathbf{I}_n)$, and r_j is a precision parameter with prior $\text{Ga}(a_r, b_r)$. We can integrate out $f_j(\mathbf{Z})$ to obtain the marginalized likelihood

$$\boldsymbol{\beta}_j(\mathbf{Z}) \sim \text{GP}(\mathbf{0}, \mathbf{C}_j + \frac{1}{r_j}\mathbf{I}_n). \tag{8}$$

As noted in Neal (1998), the ‘‘jitter’’ term is added to the covariance matrix to maintain the positive definite condition in computation. Given Equation 8, by selecting an appropriate covariance matrix, we establish a non-linear relationship between the predictor effects $\boldsymbol{\beta}_j$ and the covariates \mathbf{Z} . Different covariance

matrices can capture this relationship, as discussed in Rasmussen and Williams (2006). We adopt the single-term exponential covariance structure from Savitsky et al. (2011), for its simplicity and flexibility, which accommodates a wide range of linear and non-linear relationships. However, VERGE is general and is capable of using any valid kernel function, for example, the Matern kernel. The covariance matrix \mathbf{C}_j in our model comprises a constant term and an exponential term,

$$\mathbf{C}_j = \frac{1}{\lambda_{aj}}\mathbf{J}_n + \frac{1}{\lambda_{zj}}\exp(-\mathbf{M}), \tag{9}$$

where \mathbf{J}_n is an $n \times n$ matrix of 1's and \mathbf{M} is a matrix with entries $m_{ii'}$. Here, $m_{ii'}$ is defined as $(\mathbf{Z}_i - \mathbf{Z}_{i'})'\mathbf{P}(\mathbf{Z}_i - \mathbf{Z}_{i'})$, where \mathbf{P} is the diagonal matrix $\text{diag}(-\log(\rho_{j1}, \dots, \rho_{jK}))$, and $\rho_{jk} \in [0, 1]$ is the parameter associated with covariate \mathbf{Z}_k for $k = 1, \dots, K$.

To identify which specific covariates are important in modulating the effect of each predictor, we place spike-and-slab priors on the covariance parameters

$$p(\rho_{jk} \mid \tilde{\gamma}_{jk}) = \tilde{\gamma}_{jk}\mathbf{I}[0 < \rho_{jk} < 1] + (1 - \tilde{\gamma}_{jk})\delta_1(\rho_{jk}), \tag{10}$$

for $j = 1, \dots, P$ and $k = 1, \dots, K$, where δ_1 represents a point mass distribution at 1, which translates to 0 once the log transformation is applied. The indicator variable $\tilde{\gamma}_{jk}$ follows a Bernoulli distribution $\text{Ber}(\alpha_{jk})$. When $\tilde{\gamma}_{jk} = 1$, the magnitude of $\rho_{jk} \in (0, 1)$ controls the smoothness of the function, while $\tilde{\gamma}_{jk} = 0$ indicates the covariate \mathbf{Z}_k having no effect on the j th predictor, with $\rho_{jk} = 1$. Finally, we complete our model by assuming Gamma priors on the scaling parameters of Equation 9 as $\lambda_{aj} \sim \text{Ga}(a_\lambda, b_\lambda)$ and $\lambda_{zj} \sim \text{Ga}(a_z, b_z)$ for $j = 1, \dots, P$. Due to the sensitivity to scaling of prior (9), normalizing the covariates \mathbf{Z} to the unit cube is recommended by Savitsky et al. (2011).

3 POSTERIOR INFERENCE

Given that the posterior is intractable, we utilize MCMC methods to sample parameters from the posterior distribution. To avoid directly sampling the realizations for $\boldsymbol{\beta}_j(\mathbf{Z})$, and to minimize uncertainty, we integrate out $\boldsymbol{\beta}_j(\mathbf{Z})$ for $j = 1, \dots, P$. Since we have spike-and-slab priors on both the predictor coefficients and the covariance parameters for each covariate, our MCMC scheme requires careful consideration of how to handle the changing dimensions of the parameter space. To deal with the multiple layers of selection and ensure good mixing, we incorporate both between-model moves (where we update the predictor or covariate selection) and within-model moves (where we update the parameters while keeping the predictor and covariate selection fixed). The details of the MCMC scheme are provided in the [Supplementary Materials](#).

3.1 Variable and edge selections

For variable selection, as recommended by Barbieri and Berger (2004), we utilize the median probability model, which includes variables with a marginal posterior probability of inclusion (PPI) of at least 0.5. The marginal PPI for each predictor j is determined by calculating the frequency of inclusion in the model

across the post burn-in MCMC samples, represented as $\delta_j = \frac{\sum_{t=1}^N \mathbf{I}(\gamma_j^{(t)}=1)}{N}$, where N is the total number of samples. For each covariate Z_{jk} , the PPI is calculated as the number of iterations where Z_{jk} is included in the model out of the total number of iterations where its corresponding X_j was selected. As an alternative, the expected false discovery rate (FDR) can be used to set a threshold for the PPIs. This is calculated as $\text{FDR}(\kappa) = \frac{\sum_{j=1}^p \mathbf{I}(\delta_j > \kappa)(1-\delta_j)}{\sum_{j=1}^p \mathbf{I}(\delta_j > \kappa)}$, where $\kappa \in (0, 1)$ is the threshold value. Furthermore, for edge selection, following the approach from Wang (2015) and Peterson et al. (2016), we adopt the posterior median graph for the selection of the graph structure. Specifically, we select edges that have a marginal PPI greater than 0.5.

3.2 Prediction

To perform prediction, we follow the method in Rasmussen and Williams (2006) to incorporate the information provided by the training data about the function. Let \mathbf{X}^* and \mathbf{Z}^* be the predictors and covariates for future observations, and $\boldsymbol{\beta}^* = \boldsymbol{\beta}(\mathbf{Z}^*)$ represent the corresponding $n^* \times 1$ latent vector. We then consider the joint distribution

$$\begin{bmatrix} \boldsymbol{\beta} \\ \boldsymbol{\beta}^* \end{bmatrix} \sim \mathcal{N}\left(\begin{bmatrix} \mathbf{0} \\ \mathbf{0} \end{bmatrix}, \begin{bmatrix} \mathbf{C}(\mathbf{Z}, \mathbf{Z}) & \mathbf{C}(\mathbf{Z}, \mathbf{Z}^*) \\ \mathbf{C}(\mathbf{Z}^*, \mathbf{Z}) & \mathbf{C}(\mathbf{Z}^*, \mathbf{Z}^*) \end{bmatrix}\right),$$

where $\mathbf{C}(\mathbf{Z}, \mathbf{Z}^*) := \mathbf{C}(\mathbf{Z}, \mathbf{Z}^*)(\boldsymbol{\Theta})$ denotes the $n \times n^*$ covariance matrix calculated for all pairs of training and test points, $\boldsymbol{\Theta} = \{\boldsymbol{\Theta}_1, \dots, \boldsymbol{\Theta}_p\}$, and $\boldsymbol{\Theta}_j = \{\tilde{\gamma}_j, \rho_j, \lambda_{aj}, \lambda_{zj}, r_j\}$ for $j = 1, \dots, p$. The expectation of the conditional joint predictive distribution for $\boldsymbol{\beta}^* | \boldsymbol{\beta}$ is $\mathbf{C}(\mathbf{Z}^*, \mathbf{Z})\mathbf{C}(\mathbf{Z}, \mathbf{Z})^{-1}\boldsymbol{\beta}$, which we can estimate based on the MCMC samples as

$$\hat{\boldsymbol{\beta}}_j^*(\boldsymbol{\Theta}^{(t)}) := \mathbf{C}(\mathbf{Z}^*, \mathbf{Z})(\boldsymbol{\Theta}^{(t)})\mathbf{C}(\mathbf{Z}, \mathbf{Z})^{-1}(\boldsymbol{\Theta}^{(t)})\hat{\boldsymbol{\beta}}_j, \quad (11)$$

where $\hat{\boldsymbol{\beta}}_j = \frac{\sum_{t=1}^N \mathbf{I}(\gamma_j^{(t)}=1)\tilde{\boldsymbol{\beta}}_j(\mathbf{Z})}{\sum_{t=1}^N \mathbf{I}(\gamma_j^{(t)}=1)}$, when the j th predictor is included in the model. Here, $\tilde{\boldsymbol{\beta}}_j(\mathbf{Z})$ represents the sampled values from each iteration. We can then obtain the estimated response value $\hat{\mathbf{y}}^* = \frac{1}{L} \sum_{t=1}^L (\sum_{j=1}^p \mathbf{I}(\tilde{\gamma}_j > 0.5) \mathbf{X}_j^* \hat{\boldsymbol{\beta}}_j^*(\boldsymbol{\Theta}^{(t)}))$, where $\tilde{\gamma}_j$ represents the marginal PPI for the j th predictor, and L is the total number of samples where all predictors with marginal PPI greater than 0.5 are selected. This involves averaging over the MCMC samples to obtain the final estimate. Note that only covariates that have been selected based on the marginal PPIs are included in the computation of the covariance matrices in Equation 11. Moreover, as suggested by Neal (1998), we rely on the Cholesky decomposition for computing $\mathbf{C}(\mathbf{Z}, \mathbf{Z})^{-1}$ in Equation 11.

4 SIMULATION STUDY

4.1 Simulation setup

In our simulation design, we first construct the graph representing the dependence relations among the predictor variables. We create a sparse network with clusters of correlated features, similar to Li and Li (2008) and Peterson et al. (2016). Predictors are represented as clusters of genes, including a transcription factor and its regulated genes. Our graph consists of $P = 60$ nodes, divided into 12 clusters. Each cluster contains 1 primary node

functioning as a hub, connected to 4 remaining nodes in the cluster, resulting in a network with 48 total edges.

The predictor variables X_i are sampled from a multivariate normal distribution with mean zero and covariance matrix $\boldsymbol{\Sigma}_G$. The matrix $\boldsymbol{\Sigma}_G$ has unit variances for each predictor. Within each cluster, we set a correlation of 0.7 between the primary node and the 4 subsidiary nodes, and the correlations among the subsidiary nodes are fixed to 0.7². This results in a sparse graph structure with edges limited to within-cluster connections, where each primary node is connected to all 4 subsidiary nodes.

We assume 10% of the predictors are relevant to the outcome, resulting in $P_{\text{true}} = P/10 = 6$. The response variable y_i is generated using the linear model $y_i = \sum_{j=1}^6 X_{ij}\boldsymbol{\beta}_j(\mathbf{Z}_i) + \varepsilon_i$, for $i = 1, \dots, n$, where $\varepsilon_i \sim \mathcal{N}(0, 1)$. We specify $n = 200$ as the number of training samples for parameter estimation, and $n_t = 50$ as the number of test samples for the evaluation of prediction performance. The covariates \mathbf{Z}_k are randomly sampled from $\text{Unif}(-1, 1)$ with $K = 3$. For the predictor effects $\boldsymbol{\beta}_j(\mathbf{Z})$, we consider different generating functions, including constant, linear, and non-linear forms. The true values of the $\boldsymbol{\beta}_j(\mathbf{Z})$'s are defined as follows:

$$\begin{aligned} \boldsymbol{\beta}_1(\mathbf{Z}_1) &= 0.3, \quad \boldsymbol{\beta}_2(\mathbf{Z}_2) = 2\sin(\pi\mathbf{Z}_2), \quad \boldsymbol{\beta}_3(\mathbf{Z}_3) \\ &= 2\mathbf{Z}_3^2 - 1, \quad \boldsymbol{\beta}_4(\mathbf{Z}_4) = -2\mathbf{Z}_4, \\ \boldsymbol{\beta}_5(\mathbf{Z}_5) &= 2\cos(\pi\mathbf{Z}_5), \quad \boldsymbol{\beta}_6(\mathbf{Z}_6) = -2\mathcal{N}(\mathbf{Z}_6 | 0.3, 0.3^2) \\ &\quad - 3\mathcal{N}(\mathbf{Z}_6 | -0.5, 0.3^2), \quad \boldsymbol{\beta}_j(\mathbf{Z}_j) = 0, \end{aligned} \quad (12)$$

for $j = 7, \dots, 60$, where \mathbf{Z}_j , which represents the covariate influencing the j th predictor, is randomly selected from the $K = 3$ covariates. We exclude categorical functions in this setup since not all methods in our simulation studies are designed for handling categorical covariates. However, VERGE is effective with binary and categorical covariates, as shown in the application section. Finally, we center \mathbf{y} and standardize the predictors \mathbf{X} and covariates \mathbf{Z} to ensure stable results when applying the graphical model (Wang, 2015) and the Gaussian process model (Savitsky et al., 2011).

Parameter settings and sensitivity analyses are discussed in the [Supplementary Materials](#).

4.2 Comparative analysis

To characterize the impact of the components of our proposed VERGE model, including the incorporation of graph information and the selection of predictors, we consider 2 reduced forms of our model as comparators. In the first reduced model, we omit the second term in Equation 6, so that no graph information is incorporated under the prior. In this model, the prior on $\boldsymbol{\gamma}$ simplifies to an independent Bernoulli. In the second reduced model, we include a GP prior on $\boldsymbol{\beta}_j(\mathbf{Z})$, but do not perform selection of the primary predictors.

Furthermore, we compare VERGE with 2 established methods: the pliable lasso method (Kim et al., 2021), implemented in the R package `svreg`, and the spline-based Bayesian Hierarchical Varying-Sparsity Regression (BEHAVIOR) model Ni et al. (2019). The pliable lasso was fit with 2 penalty parameters selected using 5-fold cross-validation on the training data. The Bayesian models were run in Matlab Release 2022b. For the

TABLE 1 Simulation results for predictor selection and prediction accuracy for $n = 200, P = 60, K = 3$.

	pLasso	BEHAVIOR	VERGE	No graph	No predictor selection
TPR	0.460 (0.265)	0.853 (0.055)	0.960 (0.073)	0.833 (0.068)	–
FPR	0.157 (0.190)	0.002 (0.006)	0.001 (0.004)	0.002 (0.005)	–
MCC	0.340 (0.194)	0.905 (0.046)	0.974 (0.051)	0.896 (0.056)	–
F1	0.356 (0.163)	0.911 (0.042)	0.975 (0.047)	0.902 (0.052)	–
AUC	0.381 (0.159)	0.966 (0.027)	0.999 (0.001)	0.972 (0.033)	–
PMSE	6.657 (2.195)	1.326 (0.199)	1.278 (0.260)	1.295 (0.276)	7.125 (2.220)

Methods compared include the pliable lasso Tibshirani and Friedman (2020), BEHAVIOR Ni et al. (2019), VERGE, and 2 reduced versions of VERGE: one with no graph selection, and one with no predictor selection. Performance metrics evaluated are the true positive rate (TPR), false positive rate (FPR), F1 score, Matthews correlation coefficient (MCC), the area under the ROC curve (AUC), and mean squared prediction error (PMSE).

TABLE 2 Simulation results for covariate selection for $n = 200, P = 60, K = 3$.

	pLasso	BEHAVIOR	VERGE	No graph	No predictor selection
TPR	0.280 (0.271)	0.992 (0.040)	0.992 (0.040)	0.984 (0.055)	0.800 (0.062)
FPR	0.018 (0.050)	0.001 (0.001)	0.007 (0.002)	0.001 (0.001)	0.011 (0.012)
MCC	–	0.986 (0.041)	0.985 (0.035)	0.988 (0.034)	0.758 (0.043)
F1	0.263 (0.241)	0.985 (0.042)	0.985 (0.036)	0.988 (0.035)	0.764 (0.042)
AUC	0.075 (0.084)	0.964 (0.018)	0.996 (0.022)	0.992 (0.029)	0.995 (0.007)

Methods compared include the pliable lasso Tibshirani and Friedman (2020), BEHAVIOR Ni et al. (2019), VERGE, and 2 reduced versions of VERGE: one with no graph selection, and one with no predictor selection. Performance metrics evaluated are the true positive rate (TPR), false positive rate (FPR), F1 score, Matthews correlation coefficient (MCC), and the area under the ROC curve (AUC). Note: some MCC values for the pliable lasso are omitted due to zero denominators.

BEHAVIOR model, 150 000 iterations for burn-in and 150 000 for inference were needed to reach convergence, using the Matlab code provided by Ni et al. (2019). Our VERGE model required 60 000 iterations, with the first 30 000 for burn-in. The acceptance rates for the Level 1 moves were 5% for the “Add” proposal, 45% for the “Delete” proposal, and 25% for the “Keep” proposal. Despite the lower acceptance rate for the “Add” move, the rates for “Keep” and “Delete” moves facilitate updates in the GP kernel and improve mixing. The acceptance rates for updating covariate inclusion in the within-model move ranged from 40% to 60%. We evaluated the Pearson correlation for the PPIs of both predictor and covariate selections from 2 independent chains and found good indications of convergence.

We assess the performance of the models using the following metrics: the true positive rate (TPR), false positive rate (FPR), F_1 score, Matthews correlation coefficient (MCC), and area under the ROC curve (AUC), for both variable and covariate selection. Additionally, we employ the mean squared prediction error (PMSE) as a measure to evaluate the predictive performance of the models. Specifically, the metrics are defined as $TPR = \frac{TP}{TP+FN}$, $FPR = \frac{FP}{FP+TN}$, $F_1 = \frac{2TP}{2TP+FP+FN}$, $MCC = \frac{TP \times TN - FP \times FN}{(TP+FP)(TP+FN)(TN+FP)(TN+FN)}$ and $PMSE = \frac{1}{n_t} \sum_{i=1}^{n_t} (\hat{y}_i - y_{test,i})^2$, where TP, TN, FP, and FN denote the true positives, true negatives, false positives, and false negatives, respectively, and n_t is the sample size for the test data. For the pliable lasso, which has 2 tuning parameters that regulate the penalties for predictor and covariate selections, the AUC calculation requires varying one of these parameters while selecting the second parameter through a 5-fold cross-validation at each level of the first parameter. AUCs for VERGE and BEHAVIOR are calculated by changing the thresholds for the PPIs.

We first considered a simulated setting with $n = 200, P = 60, K = 3$. Results for predictor selection are shown in Table 1,

and those for covariate selection are shown in Table 2, across 25 simulated data sets. In comparison to the full VERGE model, the predictor selection for the reduced model with no graph information exhibited a lower TPR. Upon examining each predictor, this decrease is attributed to the extremely low TPR (0.080) for the predictor with a small constant effect. Without graph information, the reduced model fails to identify this effect. However, the covariate selection and PMSE remain similar. The reduced model variant with no selection of the primary predictors achieved acceptable performance in covariate selection, demonstrating its ability to identify most of the covariate effects. However, it led to a high PMSE, likely due to overfitting caused by the absence of predictor selection, as all predictors were considered in the model.

With respect to competing methods, results indicate that VERGE performs well in both predictor and covariate selection, notably outperforming existing methods in predictor selection due to the integration of graph information. This helps in identifying connected predictors with slightly weaker effects. Additional simulations were conducted with more covariates ($K = 6$) and a smaller sample size ($n = 100$), with detailed results reported in [Supplementary Materials](#). Overall, VERGE achieves TPRs that are either 1 or close to 1 across all scenarios, with FPRs below 0.001. In terms of covariate selection, VERGE is comparable with BEHAVIOR. However, VERGE experiences a slight reduction in performance with an increase in covariates or a decrease in sample size, while BEHAVIOR shows a significant decrease in smaller sample scenarios but maintains strong performance with increased covariates. The pliable lasso, limited to identifying only linear covariate effects, exhibits lower accuracy across all scenarios. Additionally, its penalty parameters, optimized for prediction error through cross-validation, are not geared toward predictor selection, as discussed in Meinshausen and Bühlmann (2006). In predictive performance, our method

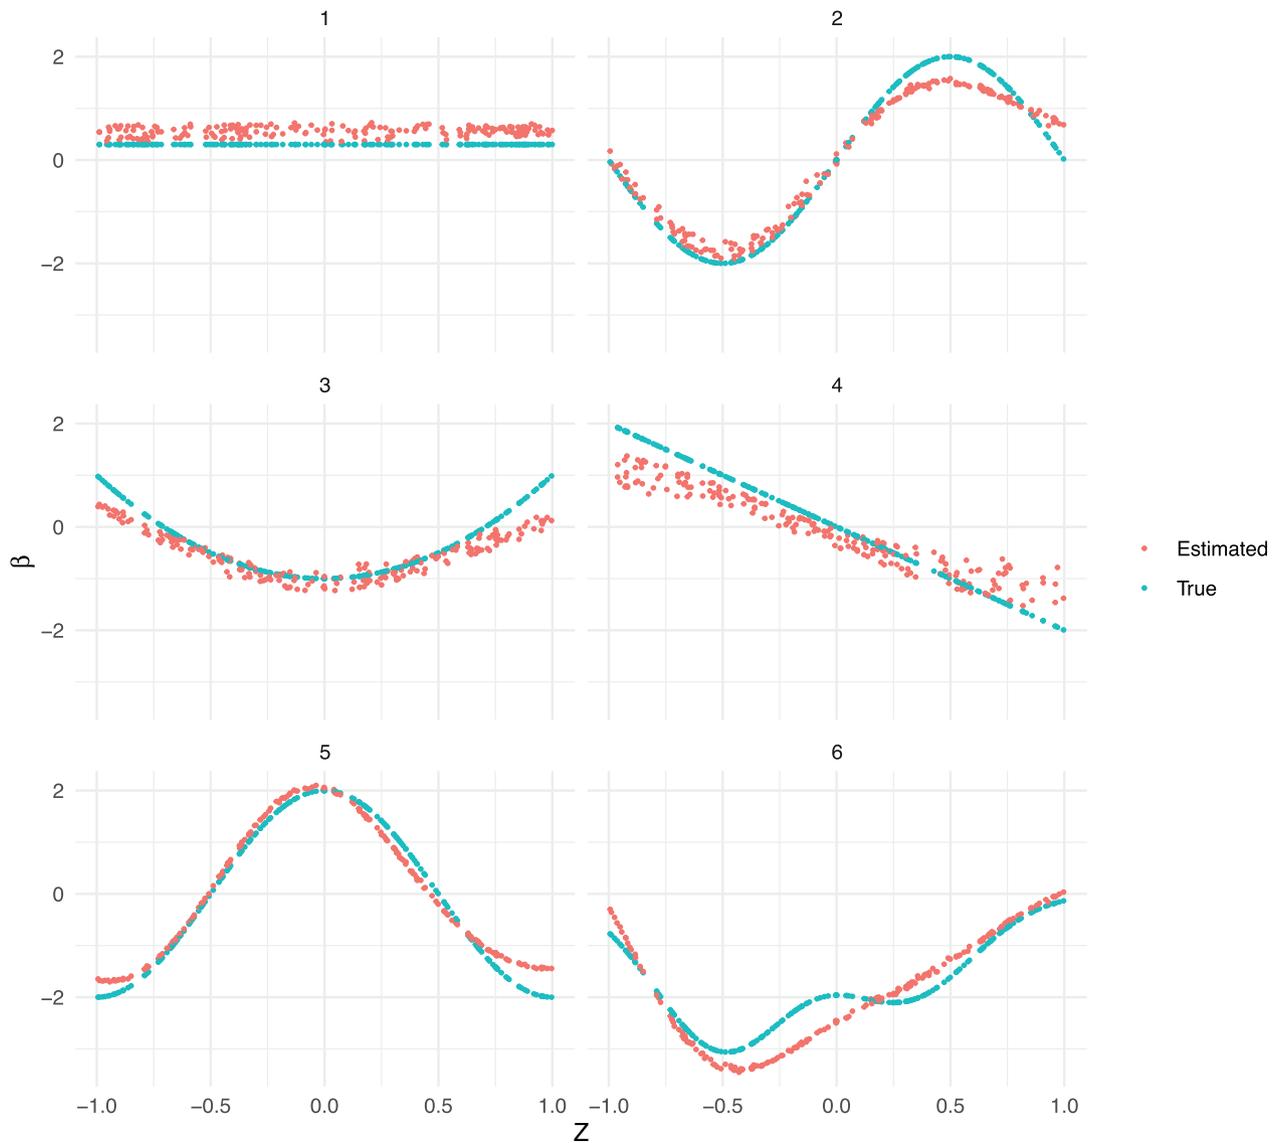


FIGURE 1 True and estimated coefficients for simulation study by generating functions.

has the lowest PMSE in all scenarios, whereas the pliable lasso has the highest.

Figure 1 displays true and estimated coefficients for each underlying function listed in Equation 12, from 1 simulated dataset. The estimated coefficients were calculated by averaging the sampled β 's from the third step of the MCMC algorithm, as detailed in the [Supplementary Materials](#). VERGE effectively recovers all function types, including constant, linear, and non-linear, closely matching the true values. We also evaluated the average TPRs for predictor selection across all methods by generating function. Focusing on the base scenario ($n = 200$, $P = 60$, $K = 3$), results in [Web Table 3](#) show that our method performs well across all functions. BEHAVIOR excels with covariate-dependent effects but struggles with small constant coefficients. The pliable lasso only performs well with functions that do not cross 0. VERGE's ability to detect small constant coefficients likely benefits from the graph structure guiding the predictor selection.

5 CASE STUDY

The gut microbiome, consisting of trillions of bacteria, plays a critical role in extracting energy from the diet and influences human health outcomes, including obesity (Turnbaugh et al., 2006). However, the mechanisms behind this link are not fully understood. In this case study, we consider how the effect of gut microorganisms on body mass index (BMI) may be modulated by covariates such as sex and dietary intake, using the COMBO dataset originally described by Wu et al. (2011). The data were obtained from a cross-sectional study of 98 healthy volunteers, where stool samples were analyzed using 16S rRNA gene segments via 454/Roche pyrosequencing. Additionally, diet and demographic information, including age, sex, and BMI, were collected. This dataset was previously analyzed by Lin et al. (2014) and Zhang et al. (2021), who focused on microbiome feature selection without considering potential covariate effects.

TABLE 3 Selected genera in the gut microbiome data and their corresponding covariates.

Phylum	Family	Genus	Selected covariates	Averaged coefficient
Bacteroidetes	<i>Bacteroidaceae</i>	<i>Bacteroides</i>	tfat, aofib	0.88 (0.30)
Firmicutes	<i>Erysipelotrichaceae</i>	<i>Catenibacterium</i>	tfat, aofib	0.59 (0.15)
Firmicutes	Family XIII	AD3011 group	sex, aofib	0.84 (0.16)
Firmicutes	<i>Lachnospiraceae</i>	<i>Anaerostipes</i>	aofib	0.82 (0.17)
Firmicutes	<i>Lachnospiraceae</i>	<i>Lachnoclostridium</i>	–	–0.88 (0.05)
Firmicutes	<i>Lachnospiraceae</i>	NK4A136 group	–	–0.62 (0.03)
Firmicutes	<i>Lachnospiraceae</i>	UCG-004	aofib	0.58 (0.11)
Firmicutes	<i>Ruminococcaceae</i>	<i>Ruminococcus 2</i>	–	0.85 (0.09)
Firmicutes	<i>Ruminococcaceae</i>	UCG-002	sex, aofib	–1.03 (0.14)
Firmicutes	<i>Ruminococcaceae</i>	Unclassified	sex, tfat	–0.45 (0.11)
Firmicutes	<i>Veillonellaceae</i>	<i>Megasphaera</i>	sex, tfat, aofib	0.94 (0.14)

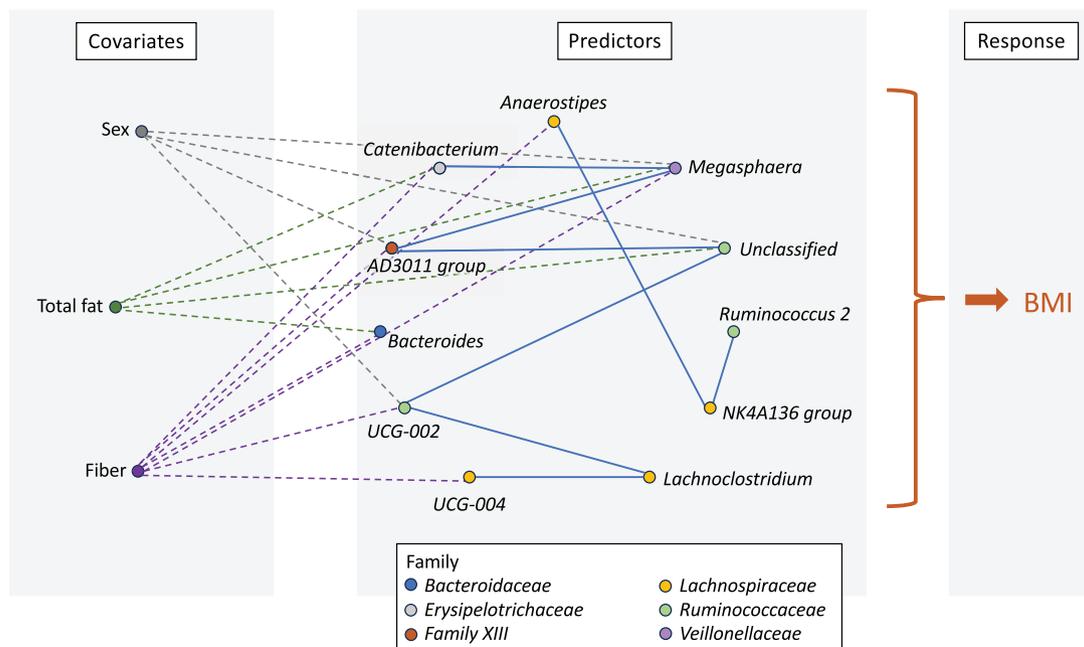


FIGURE 2 Connections among selected predictors (genus-level microbiome features) and their corresponding covariates.

In our analysis, we utilized the reprocessed data from Zhang et al. (2021), which used the updated SILVA rRNA database to assign the sequences to taxonomy. We obtained 156 genera from 1763 Operational Taxonomic Units (OTUs) using the R package `phyloseq`, focusing on 69 taxa with an average abundance $\geq 0.1\%$. Given the compositional nature of the data, the features were transformed using a centered log-ratio transform (Aitchison, 1982) prior to downstream analysis. Since we are interested in predicting obesity as a health outcome, we used BMI as our response variable. We selected sex, total fat (tfat), and total fiber measured using the AOAC method (aofib, McCleary et al., 2010) as the candidate covariates that could modify the effects of the microbiome features on BMI. The response variable was centered, and both the predictors and covariates were normalized in the analysis.

We fit the model to the data using parameter settings similar to those in our simulation study, but with slight modifications. To account for the relatively weaker signal compared to simulated

data, we increased the prior probability of predictor inclusion parameter a to $\log(0.22)$. We also chose smaller values for the hyperparameters ν_0 and ν_1 for graph estimation, setting them at 0.1 and 5, respectively. The MCMC simulations included a burn-in of 100 000 iterations, followed by 100 000 iterations for analysis. Similar to the simulation study, we selected predictors and edges with marginal PPIs above 0.5. For covariate selection, we used a cutoff of 0.5, rather than controlling FDR, due to the weaker signal in real data.

The VERGE model identified 11 genera influencing BMI, as shown in Table 3. It lists the genera, their covariates, and average coefficients across 98 subjects. Four genera exhibited protective effects with a negative association with BMI, including 2 from the family *Lachnospiraceae* and 2 from the family *Ruminococcaceae*. Previous studies have noted the depletion of these families in obese adults, suggesting they digest dietary fiber into short chain fatty acids, potentially modulated by diet (Peters et al., 2018; Vacca et al., 2020). Specifically, both *Ruminococcaceae*

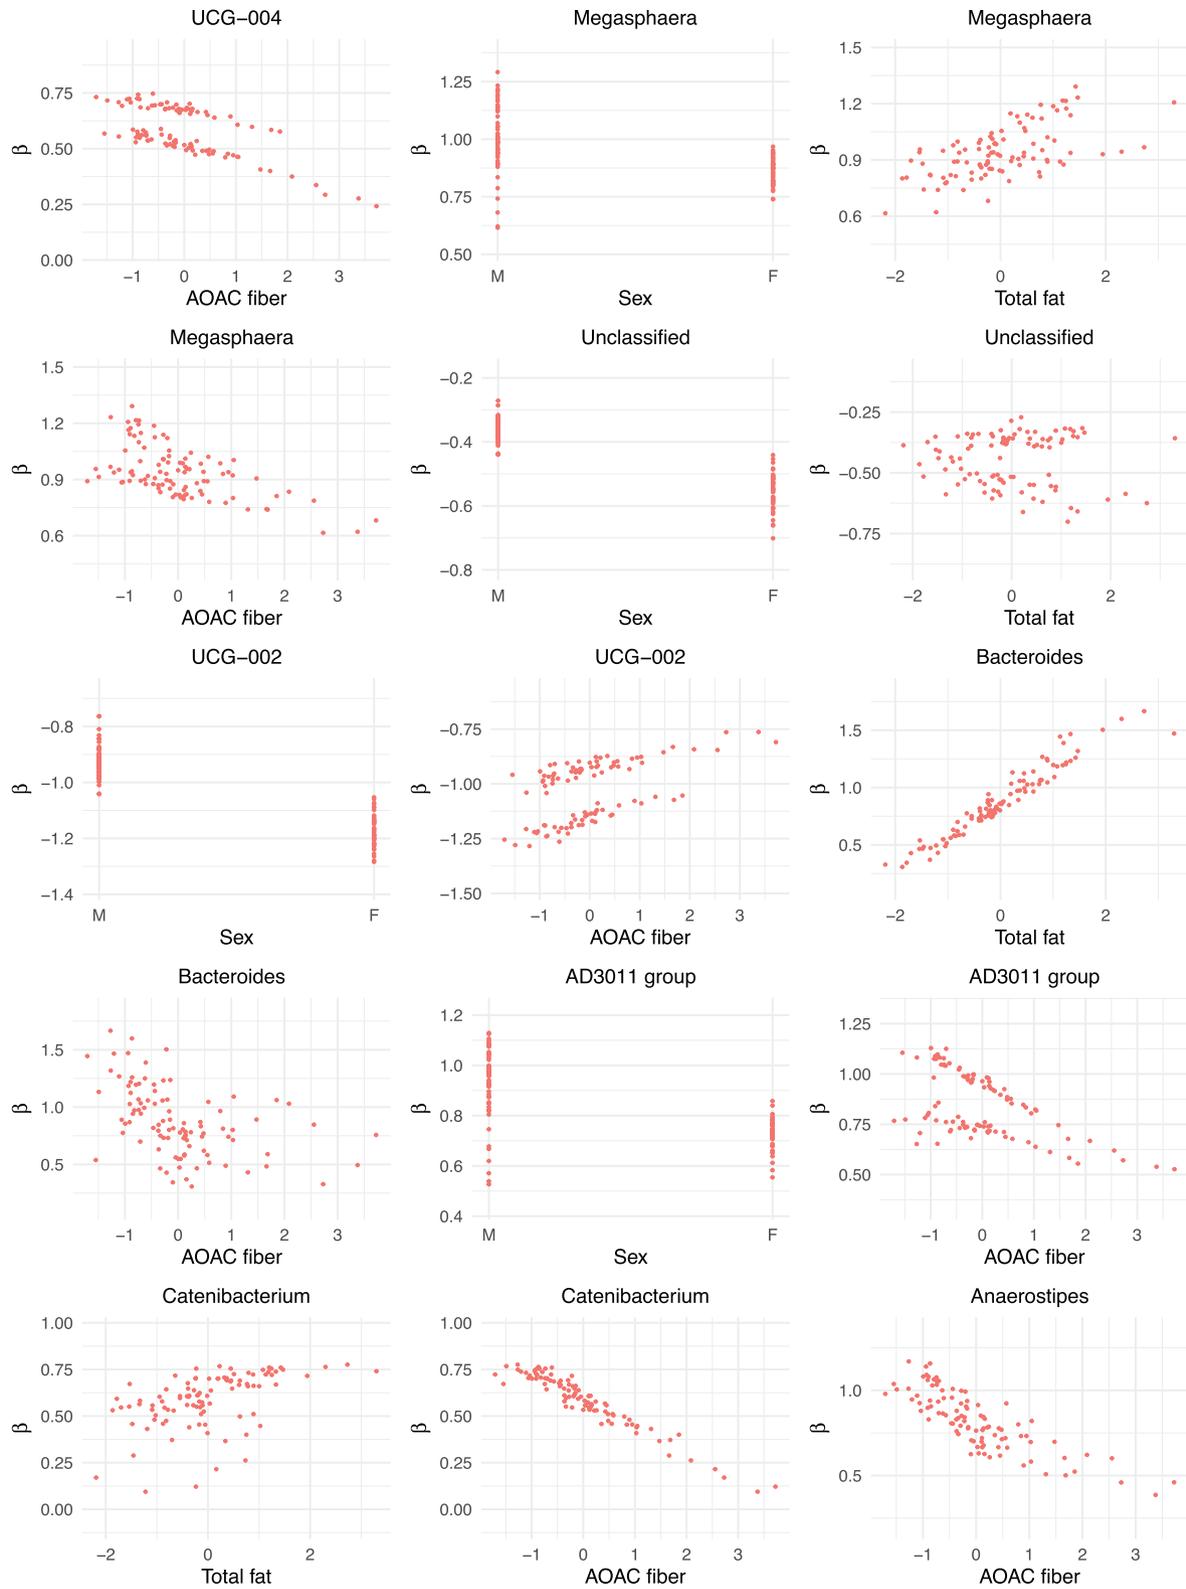


FIGURE 3 Estimated coefficients for selected gut microbiome and their corresponding covariates.

genera showed dietary-dependent effects. Conversely, genera like *Catenibacterium* and *Megasphaera*, linked to increased BMI in prior research (Pinart et al., 2021), were also identified.

In addition to identifying microbiome features, we also uncover their interrelationships and the covariates that modulate their effects. Figure 2 shows the connections among selected predictors and covariates for each genus. The inferred microbiome network consists of 108 edges, including 8 among the selected predictors. Several of the connections link closely related genera, consistent with previous findings that microbiome networks often show an assortative structure, where taxa close in the taxonomic tree are commonly linked in co-occurrence networks (Ha et al., 2020).

The effects of the selected covariates on the estimated predictor coefficients are illustrated in Figure 3, highlighting several interesting relationships. First, increased dietary fiber intake generally attenuates the impact of microbiome features on BMI, underscoring fiber's protective role (Den Besten et al., 2013). Second, coefficient estimates for male versus female subjects show notable separation in many plots, both directly illustrating sex effects and indirectly through a striated pattern. These variations align with studies indicating sex-dependent relationships between microbiome composition and body fat (Min et al., 2019), including sex-specific effects of *Ruminococcaceae*.

We report on comparisons to alternative methods in the [Supplementary Materials](#).

6 CONCLUSION

In this study, we proposed a novel regression framework, Bayesian VERGE, that enables the selection of both network-linked predictor variables and covariates that modify the predictor effects, while learning network connections among these predictors. Our approach employs a Gaussian process prior, allowing for a flexible and varying effect of predictors on the outcome variable, dependent on specific covariates, and spike-and-slab priors to achieve sparsity at both the predictor and covariate levels. Simulated data demonstrate that our method can match or exceed existing methods in terms of predictor and covariate selection and prediction accuracy, particularly in identifying predictors with subtle constant effects through the use of graph information.

We applied our method to identify microbiome effects on obesity, influenced by sex and dietary factors. The selected genera and network relationships align with previously reported associations. Our analysis reveals additional nuanced relationships, particularly highlighting dietary fiber's protective role on health and how microbiome features impact BMI differently based on sex. These associations suggest correlations rather than causality due to the biological system's complexity. To confirm causal links, prospective randomized studies of microbiome and dietary interventions are necessary (Durack and Lynch, 2019).

While our current framework focuses on continuous responses, VERGE is adaptable to other types, such as binary responses or survival outcomes. Additionally, it could be expanded into a longitudinal framework by incorporating time as a varying effect. Furthermore, while the current approach uses the same

set of covariates for all predictors, it could be modified to accommodate different covariates for various predictors based on prior knowledge.

ACKNOWLEDGMENTS

The authors thank Liangliang Zhang for sharing the reprocessed COMBO data.

SUPPLEMENTARY MATERIALS

Supplementary material is available at [Biometrics](#) online.

MCMC Algorithm, Tables, and Figures referenced in Sections 3-5 are available with this paper at the Biometrics website on Oxford Academic. Matlab code to recreate results and example data are posted online with this article and also available on GitHub at <https://github.com/YangfanR/VERGE/tree/main>.

FUNDING

This research was supported in part by grants NSF/DMS-2113602 and NSF/DMS-2113557.

CONFLICT OF INTEREST

None declared.

DATA AVAILABILITY

The dataset analyzed in this paper was originally described in Wu et al. (2011). The processed version of the data analyzed in this paper is available as part of the Supplementary Materials.

REFERENCES

- Aitchison, J. (1982). The statistical analysis of compositional data. *Journal of the Royal Statistical Society: Series B (Methodological)*, 44, 139–160.
- Barbieri, M. M. and Berger, J. O. (2004). Optimal predictive model selection. *Annals of Statistics*, 32, 870–897.
- Bürgin, R. and Ritschard, G. (2015). Tree-based varying coefficient regression for longitudinal ordinal responses. *Computational Statistics & Data Analysis*, 86, 65–80.
- Cleveland, W. S. and Grosse, E. (1991). Computational methods for local regression. *Statistics and Computing*, 1, 47–62.
- Den Besten, G., Van Eunen, K., Groen, A. K., Venema, K., Reijngoud, D.-J. and Bakker, B. M. (2013). The role of short-chain fatty acids in the interplay between diet, gut microbiota, and host energy metabolism. *Journal of Lipid Research*, 54, 2325–2340.
- Durack, J. and Lynch, S. V. (2019). The gut microbiome: relationships with disease and opportunities for therapy. *Journal of Experimental Medicine*, 216, 20–40.
- George, E. I. and McCulloch, R. E. (1997). Approaches for Bayesian variable selection. *Statistica Sinica*, 7, 339–373.
- Ha, M. J., Kim, J., Galloway-Peña, J., Do, K.-A. and Peterson, C. B. (2020). Compositional zero-inflated network estimation for microbiome data. *BMC Bioinformatics*, 21, 1–20.
- Hastie, T. and Tibshirani, R. (1993). Varying-coefficient models. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 55, 757–779.
- Kim, R., Mueller, S. and Garcia, T. P. (2021). svreg: structural varying-coefficient regression to differentiate how regional brain atrophy

- affects motor impairment for Huntington disease severity groups. *Biometrical Journal*, 63, 1254–1271.
- Kuo, L. and Mallick, B. (1998). Variable selection for regression models. *Sankhyā: The Indian Journal of Statistics, Series B*, 60, 65–81.
- Kurtz, Z. D., Müller, C. L., Miraldi, E. R., Littman, D. R., Blaser, M. J. and Bonneau, R. A. (2015). Sparse and compositionally robust inference of microbial ecological networks. *PLoS Computational Biology*, 11, e1004226.
- Leeming, E. R., Louca, P., Gibson, R., Menni, C., Spector, T. D. and Le Roy, C. I. (2021). The complexities of the diet-microbiome relationship: advances and perspectives. *Genome Medicine*, 13, 1–14.
- Li, C. and Li, H. (2008). Network-constrained regularization and variable selection for analysis of genomic data. *Bioinformatics*, 24, 1175–1182.
- Li, F. and Zhang, N. R. (2010). Bayesian variable selection in structured high-dimensional covariate spaces with applications in genomics. *Journal of the American Statistical Association*, 105, 1202–1214.
- Li, Y., Craig, B. A. and Bhadra, A. (2019). The graphical horseshoe estimator for inverse covariance matrices. *Journal of Computational and Graphical Statistics*, 28, 747–757.
- Lin, W., Shi, P., Feng, R. and Li, H. (2014). Variable selection in regression with compositional covariates. *Biometrika*, 101, 785–797.
- Marx, B. D. (2009). P-spline varying coefficient models for complex data. In: *Statistical Modelling and Regression Structures*, 19–43, Springer.
- McCleary, B. V., De Vries, J. W., Rader, J. I., Cohen, G., Prosky, L., Muggford, D. C. et al. (2010). Determination of total dietary fiber (CODEX definition) by enzymatic-gravimetric method and liquid chromatography: collaborative study. *Journal of AOAC International*, 93, 221–233.
- Meinshausen, N. and Bühlmann, P. (2006). High-dimensional graphs and variable selection with the lasso. *Annals of Statistics*, 34, 1436–1462.
- Min, Y., Ma, X., Sankaran, K., Ru, Y., Chen, L., Baiocchi, M. et al. (2019). Sex-specific association between gut microbiome and fat distribution. *Nature Communications*, 10, 2408.
- Neal, R. (1998). Regression and classification using Gaussian process priors. *Bayesian Statistics*, 6, 475.
- Ni, Y., Stingo, F. C., Ha, M. J., Akbani, R. and Baladandayuthapani, V. (2019). Bayesian hierarchical varying-sparsity regression models with application to cancer proteogenomics. *Journal of the American Statistical Association*, 114, 48–60.
- Peters, B. A., Shapiro, J. A., Church, T. R., Miller, G. et al. (2018). A taxonomic signature of obesity in a large study of American adults. *Scientific Reports*, 8, 9749.
- Peterson, C. B., Stingo, F. C. and Vannucci, M. (2016). Joint Bayesian variable and graph selection for regression models with network-structured predictors. *Statistics in Medicine*, 35, 1017–1031.
- Pinart, M., Dötsch, A., Schlicht, K., Laudes, M., Bouwman, J., Forslund, S. K. et al. (2021). Gut microbiome composition in obese and non-obese persons: a systematic review and meta-analysis. *Nutrients*, 14, 12.
- Rasmussen, C. E. and Williams, C. K. (2006). *Gaussian Processes for Machine Learning*. Cambridge, MA: MIT Press.
- Reich, B. J., Fuentes, M., Herring, A. H. and Evenson, K. R. (2010). Bayesian variable selection for multivariate spatially varying coefficient regression. *Biometrics*, 66, 772–782.
- Savitsky, T., Vannucci, M. and Sha, N. (2011). Variable selection for non-parametric Gaussian process priors: models and computational strategies. *Statistical Science*, 26, 130.
- Scheipl, F., Fahrmeir, L. and Kneib, T. (2012). Spike-and-slab priors for function selection in structured additive regression models. *Journal of the American Statistical Association*, 107, 1518–1532.
- Sonnenburg, J. L. and Bäckhed, F. (2016). Diet-microbiota interactions as moderators of human metabolism. *Nature*, 535, 56–64.
- Tibshirani, R. and Friedman, J. (2020). A pliable lasso. *Journal of Computational and Graphical Statistics*, 29, 215–225.
- Turnbaugh, P., Ley, R., Mahowald, M., Magrini, V., Gordon, J. I. et al. (2006). An obesity-associated gut microbiome with increased capacity for energy harvest. *Nature*, 444, 1027–1031.
- Vacca, M., Celano, G., Calabrese, F. M., Portincasa, P., Gobbetti, M. and De Angelis, M. (2020). The controversial role of human gut Lachnospiraceae. *Microorganisms*, 8, 573.
- Vannucci, M. (2021). Discrete spike-and-slab priors: models and computational aspects. *Handbook of Bayesian Variable Selection*. Vannucci, M. and Tadesse, M. G., Editors, 3–24. Chapman and Hall/CRC.
- Wang, H. (2012). Bayesian graphical lasso models and efficient posterior computation. *Bayesian Analysis*, 4, 867–886.
- Wang, H. (2015). Scaling it up: stochastic search structure learning in graphical models. *Bayesian Analysis*, 10, 351–377.
- Wu, G. D., Chen, J., Hoffmann, C., Bittinger, K., Chen, Y.-Y., Keilbaugh, S. A. et al. (2011). Linking long-term dietary patterns with gut microbial enterotypes. *Science*, 334, 105–108.
- Zhang, L., Shi, Y., Jenq, R. R., Do, K.-A. and Peterson, C. B. (2021). Bayesian compositional regression with structured priors for microbiome feature selection. *Biometrics*, 77, 824–838.