

Statistica Sinica Preprint No: SS-2019-0202

Title	Joint Bayesian Variable and DAG Selection Consistency for High-dimensional Regression Models with Network-structured Covariates
Manuscript ID	SS-2019-0202
URL	http://www.stat.sinica.edu.tw/statistica/
DOI	10.5705/ss.202019.0202
Complete List of Authors	Xuan Cao and Kyoungjae Lee
Corresponding Author	Kyoungjae Lee
E-mail	leekjstat@gmail.com
Notice: Accepted version subject to English editing.	

JOINT BAYESIAN VARIABLE AND DAG SELECTION CONSISTENCY FOR HIGH-DIMENSIONAL REGRESSION MODELS WITH NETWORK-STRUCTURED COVARIATES

Xuan Cao and Kyoungjae Lee

University of Cincinnati and Inha University

Abstract: We consider the joint sparse estimation of regression coefficients and the covariance matrix for covariates in a high-dimensional regression model, where the predictors are both relevant to a response variable of interest and functionally related to one another via a Gaussian directed acyclic graph (DAG) model. Gaussian DAG models introduce sparsity in the Cholesky factor of the inverse covariance matrix, and the sparsity pattern in turn corresponds to specific conditional independence assumptions on the underlying predictors. A variety of methods have been developed in recent years for Bayesian inference in identifying such network-structured predictors in regression setting, yet crucial sparsity selection properties for these models have not been thoroughly investigated. In this paper, we consider a hierarchical model with spike and slab priors on the regression coefficients and a flexible and general class of DAG-Wishart distributions with multiple shape parameters on the Cholesky factors of the inverse covariance matrix. Under mild regularity assumptions, we establish the joint selection consistency for both the variable and the underlying DAG of the covari-

ates when the dimension of predictors is allowed to grow much larger than the sample size. We demonstrate that our method outperforms existing methods in selecting network-structured predictors in several simulation settings.

Key words and phrases: DAG-Wishart prior, Posterior ratio consistency, Strong selection consistency.

1. Introduction

In modern day statistics, datasets where the number of variables is much larger than the number of samples are more pervasive than they have ever been. One of the major problems is high-dimensional variable selection, where the challenge is to select a subset of predictor variables which significantly affect a given response. The literature on Bayesian variable selection in linear regression is vast and rich. George and McCulloch (1993) propose the stochastic search variable selection which uses the Gaussian distribution with a zero mean and a small but fixed variance as the spike prior, and another Gaussian distribution with a large variance as the slab prior. Ishwaran, Kogalur, and Rao (2005) also use Gaussian spike and slab priors, but with continuous bimodal priors for the variance of the regression coefficient to alleviate the difficulty of choosing specific prior parameters. Narisetty and He (2014) introduce shrinking and diffusing priors as spike

and slab priors, and establish model selection consistency of the approach in a high-dimensional setting.

Another important problem is how to formulate models and develop inferential procedures to understand the complex relationships and multivariate dependencies in these high-dimensional datasets. A covariance matrix is one of the most fundamental objects that quantifies these relationships. A common and effective approach for covariance estimation in sample starved settings is to induce sparsity either in the covariance matrix, its inverse, or the Cholesky factor of the inverse. The sparsity patterns in these matrices can be uniquely encoded in terms of appropriate graphs. Hence the corresponding models are often referred to as covariance graph models (sparsity in Σ), concentration graph models (sparsity in $\Omega = \Sigma^{-1}$), and directed acyclic graph (DAG) models (sparsity in the Cholesky factor of Ω).

In this work, we will work in a high-dimensional regression setting, where the predictors are both relevant to a response variable of interest and functionally related to one another via a Gaussian DAG model. Our goal is to jointly perform variable and DAG selection, and to establish the selection consistency in a high-dimensional regime. The advantage of joint modeling is that we can borrow information from the DAG structure to

improve the performance of variable selection. One popular motivation for this type of problem comes from genomic studies: the mechanism for an effect on an outcome such as a quantitative molecular phenotypes including gene expression, proteomics, or metabolomics data often displays a coordinated change along a pathway, and the impact of one single genotype may not be apparent. In this setting, our proposed method can incorporate and highlight unknown pathways or regulatory networks that impact the response, which can potentially improve the performance of variable selection by borrowing information from the network structure. To uncover these relationships, we develop a Bayesian hierarchical model that favors the inclusion of variables that are not only relevant to the outcome of interest but also linked through a DAG.

When the underlying graph structure is known, several approaches including both frequentist and Bayesian methods have been proposed and studied in the literature to solve the variable selection problem. Li and Li (2008, 2010) study a graph-constrained regularization procedure and its theoretical properties to take into account the neighborhood information of the variables measured on a known graph. Pan et al. (2010) propose a grouped penalty based on the L_γ -norm that smooths the regression coefficients of the predictors over the available network. On the Bayesian

side, Li and Zhang (2010) and Stingo and Vannucci (2010) incorporate a graph structure in the Markov random field (MRF) prior on indicators of variable selection, encouraging the joint selection of predictors with known relationships. Stingo et al. (2011) and Peng et al. (2013) propose the selection of both pathways and genes within them based on prior knowledge on gene-gene interactions or functional relationships.

However, when the underlying graph is unknown and needs to be selected, comparatively fewer methods have been proposed. Dobra (2009) estimate a network among relevant predictors by first performing a stochastic search in the regression setting to identify possible subsets of predictors, then applying a Bayesian model averaging method to estimate a dependency network. Liu et al. (2014) develop a Bayesian method for regularized regression, which provides inference on the inter-relationship between variables by explicitly modeling through a graph Laplacian matrix. Peterson et al. (2016) simultaneously infer a sparse network among the predictors and perform variable selection using this network as guidance by incorporating it into a prior favoring selection of connected variables based on a Gaussian graphical model among the predictors, which provides a sparse and interpretable representation of the conditional dependencies found in the data. In a slightly different context, Chekouo et al. (2015) and Chekouo

et al. (2016) relate two sets of covariates via a DAG to integrate multiple genomic platforms and select the most relevant features. Given the ordering of variables, they use a mixture of a non-local prior (Johnson and Rossell, 2012) and a point mass at zero to infer the DAG structure.

Despite the developments in Bayesian methods for joint variable and graph selection, a rigorous investigation of the high-dimensional consistency properties of these methods has not been undertaken to the best of our knowledge. Hence, our goal was to investigate if joint selection consistency results could be established in the high-dimensional regression setting with network-structured predictors. This is a challenging goal, particularly because of the interaction between the regression coefficients and the graph in the posterior analysis, and the massive parameter space to be explored for both the coefficients and the graph.

In this paper, we consider a hierarchical multivariate regression model with DAG-Wishart priors on the covariance matrix for the predictors, spike and slab priors on regression coefficients, independent Bernoulli priors for each edge in the DAG, and a MRF prior linking the variable indicators to the graph structure. Under high-dimensional settings, we establish posterior ratio consistency, following the nomenclature in Cao et al. (2019c) and Narisetty and He (2014), for both the variable and the DAG with given

DAG and variable, respectively (Theorems 1 and 2). In Theorems 3 and 4, we also establish the posterior ratio consistency and the strong selection consistency for any pair of the DAG and variable. In particular, the strong selection consistency implies that under the true model, the posterior probability of the true variable indicator and the true graph converges in probability to 1 as $n \rightarrow \infty$. Finally, through simulation studies, we demonstrate that the models studied in this paper can outperform existing state-of-the-art methods including both penalized likelihood and Bayesian approaches in several settings.

The rest of paper is organized as follows. Section 2 provides background material regarding Gaussian DAG model and the DAG-Wishart distribution. In Section 3, we introduce our hierarchical Bayesian model. Model selection consistency results are stated in Section 4 with proofs provided in Supplementary material. In Section 5, we conduct simulation experiments to illustrate the performance of the proposed method. Benefits of our Bayesian method for identifying network-structured predictors are demonstrated vis-a-vis existing Bayesian and penalized likelihood approaches. We end our paper with a discussion session in Section 6.

2. Preliminaries

In this section, we provide the necessary background material from graph theory, Gaussian DAG models, and DAG-Wishart distributions.

2.1 Gaussian DAG models

Throughout this paper, a directed acyclic graph (DAG) $\mathcal{D} = (V, E)$ consists of a vertex set $V = \{1, \dots, p\}$ and an edge set E such that there is no directed path starting and ending at the same vertex. As in Ben-David et al. (2016) and Cao et al. (2019c), we will assume a parent ordering, where that all the edges are directed from larger vertices to smaller vertices. Thus, the ordering of variables is assumed to be known throughout the paper. The set of parents of i , denoted by $pa_i(\mathcal{D})$, is the collection of all vertices which are larger than i and share an edge with i . A Gaussian DAG model over a given DAG \mathcal{D} , denoted by $\mathcal{N}_{\mathcal{D}}$, consists of all multivariate Gaussian distributions which obey the directed Markov property with respect to a DAG \mathcal{D} . In particular, if $x = (x_1, \dots, x_p)^T \sim N_p(0, \Sigma)$ and $N_p(0, \Sigma) \in \mathcal{N}_{\mathcal{D}}$, then $x_i \perp x_{\{i+1, \dots, p\} \setminus pa_i(\mathcal{D})} | x_{pa_i(\mathcal{D})}$ for each i .

Any positive definite matrix Ω can be uniquely decomposed as $\Omega = LD^{-1}L^T$, where L is a lower triangular matrix with unit diagonal entries, and D is a diagonal matrix with positive diagonal entries. This decom-

2.1 Gaussian DAG models

position is known as the modified Cholesky decomposition of Ω (see for example Pourahmadi (2007)). It is well-known that if $\Omega = LD^{-1}L^T$ is the modified Cholesky decomposition of Ω , then $N_p(0, \Omega^{-1}) \in \mathcal{N}_{\mathcal{D}}$ if and only if $L_{ij} = 0$ whenever $i \notin pa_j(\mathcal{D})$. In other words, the structure of the DAG \mathcal{D} is reflected in the Cholesky factor L of the inverse covariance matrix.

Given a DAG \mathcal{D} on p vertices, denote $\mathcal{L}_{\mathcal{D}}$ as the set of lower triangular matrices with unit diagonals and $L_{ij} = 0$ if $i \notin pa_j(\mathcal{D})$, and let \mathcal{D}_+^p be the set of strictly positive diagonal matrices in $\mathbb{R}^{p \times p}$. We refer to $\Theta_{\mathcal{D}} = \mathcal{D}_+^p \times \mathcal{L}_{\mathcal{D}}$ as the Cholesky space corresponding to \mathcal{D} , and $(D, L) \in \Theta_{\mathcal{D}}$ as the Cholesky parameter corresponding to \mathcal{D} . In fact, the relationship between the DAG and the Cholesky parameter implies that $\mathcal{N}_{\mathcal{D}} = \{N_p(0, (L^T)^{-1}DL^{-1}) : (D, L) \in \Theta_{\mathcal{D}}\}$.

The skeleton of \mathcal{D} , denoted by $\mathcal{D}^u = (V, E^u)$, can be obtained by replacing all the directed edges of \mathcal{D} by undirected ones. We define the adjacency matrix of \mathcal{D} to be a $(0,1)$ -matrix such that the elements of the matrix indicate whether pairs of vertices are adjacent or not in \mathcal{D} , i.e., 1 representing adjacent, 0 representing not adjacent.

2.2 DAG-Wishart Distribution

In this section, we revisit the multiple shape parameter DAG-Wishart distributions introduced in Ben-David et al. (2016). Given a directed graph $\mathcal{D} = (V, E)$ with $V = \{1, \dots, p\}$ and a $p \times p$ matrix A , denote the column vectors $A_{\mathcal{D},i}^> = (A_{ij})_{j \in pa_i(\mathcal{D})}^T$ and $A_{\mathcal{D},i}^{\geq} = (A_{ii}, (A_{\mathcal{D},i}^>)^T)^T$. Also,

$$A_{\mathcal{D}}^{\geq i} = \begin{bmatrix} A_{ii} & (A_{\mathcal{D},i}^>)^T \\ A_{\mathcal{D},i}^> & A_{\mathcal{D}}^{\geq i} \end{bmatrix},$$

where $A_{\mathcal{D}}^{\geq i} = (A_{kj})_{k,j \in pa_i(\mathcal{D})}$. In particular, we have $A_{\mathcal{D},p}^{\geq} = A_{\mathcal{D}}^{\geq p} = A_{pp}$. Let $\nu_i(\mathcal{D}) = |pa_i(\mathcal{D})| = |\{j : j > i, (j, i) \in E(\mathcal{D})\}|$.

The DAG-Wishart distributions in Ben-David et al. (2016) corresponding to a DAG \mathcal{D} are defined on the Cholesky space $\Theta_{\mathcal{D}}$. Given a $p \times p$ positive definite matrix U and a p -dimensional vector $\boldsymbol{\alpha}(\mathcal{D}) = (\alpha_1(\mathcal{D}), \dots, \alpha_p(\mathcal{D}))$ with $\min_{1 \leq i \leq p} \{\alpha_i(\mathcal{D}) - \nu_i(\mathcal{D})\} > 2$, the probability density of the DAG-Wishart distribution is given by

$$\pi_{U, \boldsymbol{\alpha}(\mathcal{D})}^{\Theta_{\mathcal{D}}}(D, L) = \frac{1}{z_{\mathcal{D}}(U, \boldsymbol{\alpha}(\mathcal{D}))} \exp\left\{-\frac{1}{2} \text{tr}((LD^{-1}L^T)U)\right\} \prod_{i=1}^p D_{ii}^{-\frac{\alpha_i(\mathcal{D})}{2}} \mathbb{I}((D, L) \in \Theta_{\mathcal{D}}), \quad (2.1)$$

where

$$z_{\mathcal{D}}(U, \boldsymbol{\alpha}(\mathcal{D})) = \prod_{i=1}^p \frac{\Gamma\left(\frac{\alpha_i(\mathcal{D})}{2} - \frac{\nu_i(\mathcal{D})}{2} - 1\right) 2^{\frac{\alpha_i(\mathcal{D})}{2} - 1} (\sqrt{\pi})^{\nu_i(\mathcal{D})} \det(U_{\mathcal{D}}^{\geq i})^{\frac{\alpha_i(\mathcal{D})}{2} - \frac{\nu_i(\mathcal{D})}{2} - \frac{3}{2}}}{\det(U_{\mathcal{D}}^{\geq i})^{\frac{\alpha_i(\mathcal{D})}{2} - \frac{\nu_i(\mathcal{D})}{2} - 1}}$$

and $\mathbb{I}(\cdot)$ stands for the indicator function. The above density has the same form as the classical Wishart density, but is defined on the lower dimensional space $\Theta_{\mathcal{D}}$ and has p shape parameters $\{\alpha_i(\mathcal{D})\}_{i=1}^p$ which can be used for differential shrinkage of variables in high-dimensional settings.

The class of densities $\pi_{U, \alpha(\mathcal{D})}^{\Theta_{\mathcal{D}}}$ form a conjugate family of priors for the Gaussian DAG model $\mathcal{N}(\mathcal{D})$. In particular, if the prior on $(D, L) \in \Theta_{\mathcal{D}}$ is $\pi_{U, \alpha(\mathcal{D})}^{\Theta_{\mathcal{D}}}$ and $X_1, \dots, X_n \mid D, L, \mathcal{D} \stackrel{i.i.d.}{\sim} N_p(0, (L^T)^{-1}DL^{-1})$, then the resulting posterior distribution of (D, L) is $\pi_{\tilde{U}, \tilde{\alpha}(\mathcal{D})}^{\Theta_{\mathcal{D}}}$, where $S = \frac{1}{n} \sum_{i=1}^n X_i X_i^T$, $\tilde{U} = U + nS$, and $\tilde{\alpha}(\mathcal{D}) = (n + \alpha_1(\mathcal{D}), \dots, n + \alpha_p(\mathcal{D}))$.

3. Model Specification

In this section, we specify our hierarchical model to facilitate the purpose of joint variable and DAG selection for regression models with network-structured predictors. We start by considering the standard Gaussian linear regression model with p coefficients and by introducing some required notations. Similar to Peterson et al. (2016) and Li and Li (2008), consider both the response $Y = (y_1, \dots, y_n) \in \mathbb{R}^{n \times 1}$ and the predictors $X = (X_1, \dots, X_n)^T \in \mathbb{R}^{n \times p}$ to be random variables. In particular, $Y \sim N_n(X\beta, \sigma^2 I_n)$, and the predictors are assumed to obey a multivariate Gaussian distribution, i.e., $X_i \stackrel{i.i.d.}{\sim} N_p(0, (LD^{-1}L^T)^{-1})$ for $i = 1, 2, \dots, n$, where

$\beta \in \mathbb{R}^{p \times 1}$ is a vector of regression coefficients and (L, D) represents the Cholesky parameter corresponding to a DAG \mathcal{D} . Let symmetric matrix $G = (G_{ij})_{1 \leq i, j \leq p}$ represent the adjacency matrix corresponding to DAG \mathcal{D} where $G_{ij} = G_{ji} = 1$ if and only if there is an edge between vertex i and vertex j , and $G_{ij} = G_{ji} = 0$ otherwise. Our goal is both (i) the variable selection, i.e., to correctly identify all the non-zero regression coefficients, and (ii) network estimation, i.e., to precisely recover the sparsity pattern in \mathcal{D} .

For variable selection, we denote a variable indicator $\gamma = \{\gamma_1, \dots, \gamma_p\}$, where $\gamma_j \in \{0, 1\}$ for $1 \leq j \leq p$, and $|\gamma| = \sum_{j=1}^p \gamma_j$. Let $\beta_\gamma = (\beta_j)_{\{j: \gamma_j=1\}}^T \in \mathbb{R}^{|\gamma| \times 1}$ be the vector formed by the active components in β corresponding to a model γ . For any $n \times p$ matrix A , let A_k represent the submatrix formed from the columns of A corresponding to model k . In particular, Let X_γ denote the design matrix formed from the columns of X corresponding to model γ . For the network estimation, the class of DAG-Wishart distribution in Section 2.2 can be used for joint variable and DAG selection through the

following hierarchical model.

$$Y|X_\gamma, \beta_\gamma \sim N_n(X_\gamma\beta_\gamma, \sigma^2 I_n), \quad (3.1)$$

$$X_i|(L, D), \mathcal{D} \stackrel{i.i.d.}{\sim} N_p(0, (LD^{-1}L^T)^{-1}), \quad \text{for } i = 1, 2, \dots, n, \quad (3.2)$$

$$(L, D)|\mathcal{D} \sim \pi_{U, \alpha(\mathcal{D})}^\Theta(D, L), \quad (3.3)$$

$$\beta_\gamma|\gamma \sim N_{|\gamma|}(0, \tau^2 \sigma^2 I_\gamma), \quad (3.4)$$

$$\pi(\mathcal{D}) \propto \prod_{j=1}^{p-1} q^{\nu_j(\mathcal{D})} (1-q)^{p-j-\nu_j(\mathcal{D})} \mathbb{I}\{\max_{1 \leq j \leq p-1} \nu_j(\mathcal{D}) < R\}, \quad (3.5)$$

$$\pi(\gamma|\mathcal{D}) \propto \exp(-a1^T\gamma + b\gamma^T G\gamma) \mathbb{I}\{|\gamma| < R\}. \quad (3.6)$$

for some constants $\sigma, \tau, a > 0$, $b \geq 0$, $0 < q < 1$ and a positive integer $0 \leq R \leq p$. Here we assume that σ in (3.1) is a known constant for simplicity. However, it can be extended to unknown σ case by imposing an inverse-gamma prior, which will be shown in Corollary 1. Note that in (3.4), we are essentially imposing a spike and slab prior on the regression coefficients, where τ^2 indicates the variance of the slab part. See Narisetty and He (2014), Yang et al. (2016) and the references therein. Prior (3.5) corresponds to an Erdos-Renyi type of prior over the space of DAGs. In particular, similar to Cao et al. (2019c), define $e_{ji} = \mathbb{I}\{(j, i) \in E(\mathcal{D})\}$, $1 \leq j < i \leq p$ to be the edge indicator. Let e_{ji} , $1 \leq i < j < p$ be independent identically distributed Bernoulli(q) random variables. Recall $\nu_j(\mathcal{D}) = |pa_j(\mathcal{D})|$ is the cardinality of the parent set of vertex j . It follows

that $\pi(\mathcal{D}) = \prod_{(j,i):1 \leq j < i \leq p} q^{e_{ji}} (1 - q)^{1 - e_{ji}} = \prod_{j=1}^{p-1} q^{\nu_j(\mathcal{D})} (1 - q)^{p - j - \nu_j(\mathcal{D})}$.

Note that to obtain our desired asymptotic consistency results, appropriate conditions for these hyperparameters τ, R, a, b as well as the edge probability q will be introduced in Section 4.

Remark 1. In (3.6), given a DAG \mathcal{D} , we are imposing a Markov random field (MRF) prior on the variable indicator γ that favors the inclusion of variables linked to other variables in the associated DAG. MRF priors have also been used in the variable selection setting in Peterson et al. (2016); Li and Zhang (2010) and Stingo and Vannucci (2010). In particular, as indicated in Peterson et al. (2016), the parameter a in (3.6) controls the variable inclusion probability, with larger values of a corresponding to sparser models, while b essentially determines how strongly the inclusion probability of a variable is affected by the inclusion of its neighbors in the DAG.

The hierarchical model in (3.1)-(3.6) can be used to estimate a pair of variable and DAG as follows. By (2.1) and Bayes' rule, the following lemma gives the (marginal) joint posterior probabilities with proof provided in the Supplementary material.

Lemma 1. *Under the hierarchical model in (3.1)-(3.6), the (marginal) joint*

variable and DAG posterior is given by,

$$\begin{aligned} & \pi(\gamma, \mathcal{D} | Y, X) \\ & \propto \pi(\gamma | \mathcal{D}) \pi(\mathcal{D}) \frac{z_{\mathcal{D}}(U + X^T X, n + \alpha(\mathcal{D}))}{z_{\mathcal{D}}(U, \alpha(\mathcal{D}))} \\ & \quad \times \det(\tau^2 X_{\gamma}^T X_{\gamma} + I_{|\gamma|})^{-\frac{1}{2}} \exp\left\{-\frac{1}{2\sigma^2} \left(Y^T (I_n + \tau^2 X_{\gamma} X_{\gamma}^T)^{-1} Y\right)\right\}, \end{aligned} \quad (3.7)$$

where $z_{\mathcal{D}}(\cdot, \cdot)$ is the normalized constant in the DAG-Wishart distribution.

Hence, after integrating out β_{γ} , we have the joint posterior available in closed form (up to the multiplicative constant $\pi(X, Y)$). In particular, these posterior probabilities can be used to select a pair of variable and DAG by computing the posterior mode defined by

$$(\hat{\gamma}, \hat{\mathcal{D}}) = \underset{(\gamma, \mathcal{D})}{\operatorname{argmax}} \pi(\gamma, \mathcal{D} | Y, X). \quad (3.8)$$

4. Joint Selection Consistency

In this section we will explore the high-dimensional asymptotic properties of the Bayesian joint variable and DAG selection approach specified in Section 3. For this purpose, we will work in a setting where the number of regression coefficients $p = p_n$ increases with the sample size n . The true data generating mechanism is given by

$$Y = X\beta_0^n + \epsilon_n,$$

where $Y = (Y_1, \dots, Y_n) \in \mathbb{R}^n$, $X = (X_1, \dots, X_n)^T \in \mathbb{R}^{n \times p_n}$, $X_i \stackrel{i.i.d.}{\sim} N_{p_n}(0, \Sigma_0^n)$ and $\epsilon_n \sim N_n(0, \sigma_0^2 I_n)$. Here β_0^n is the true p_n -dimensional vector of regression coefficients, and Σ_0^n is the true covariance matrix. As in the usual context of variable selection, we assume that the true vector of regression coefficients is sparse, i.e., all the entries of β_0^n are zero except those corresponding to the active entries in the true variable indicator γ_0^n (Castillo et al., 2015; Yang et al., 2016; Narisetty and He, 2014). Denote $\rho_{1n} = \min_{j \in \gamma_0^n} |\beta_{0j}^n|$ and $\rho_{2n} = \max_{j \in \gamma_0^n} |\beta_{0j}^n|$ as the minimum and maximum magnitude of non-zero entries in β_0^n , respectively. We assume that the true quantities $|\gamma_0^n|$, ρ_{1n} and ρ_{2n} vary with n . Let $\Omega_0^n = (\Sigma_0^n)^{-1} = L_0^n (D_0^n)^{-1} (L_0^n)^T$, where (D_0^n, L_0^n) denotes the modified Cholesky parameter of Ω_0^n . Let \mathcal{D}_0^n be the true underlying DAG with structure corresponding to the sparsity pattern in L_0^n , i.e., $L_0^n \in \mathcal{L}_{\mathcal{D}_0^n}$, and let G_0^n be the adjacency matrix for \mathcal{D}_0^n . Denote d_n as the maximum number of non-zero entries in any column of L_0^n , and $s_n = \min_{1 \leq j \leq p_n, i \in pa_j(\mathcal{D}_0^n)} |(L_0^n)_{ij}|$ as the minimum magnitude of non-zero off-diagonal entry in L_0^n . Let \bar{P} denote the probability measure corresponding to the true model presented above. In order to establish the desirable consistency results, we need the following mild assumptions. Each assumption is followed by an interpretation/discussion.

Assumption 1. *There exists $0 < \epsilon_0 \leq 1$ such that $\epsilon_0 \leq \text{eig}_1(\Omega_0^n) \leq \text{eig}_{p_n}(\Omega_0^n) \leq \epsilon_0^{-1}$ for every $n \geq 1$, where $\text{eig}_1(\Omega_0^n)$ and $\text{eig}_{p_n}(\Omega_0^n)$ are the minimum and maximum eigenvalues of Ω_0^n , respectively.*

This is a standard assumption for high dimensional covariance asymptotic consistency, both in the frequentist and Bayesian paradigms. See for example Bickel and Levina (2008); El Karoui (2008); Banerjee and Ghosal (2014); Xiang et al. (2015) and Banerjee and Ghosal (2015). Cao et al. (2019c) relax this assumption by allowing the lower and upper bounds on the eigenvalues to depend on p_n and n .

Assumption 2. *For the true DAG, $d_n \sqrt{\log p_n/n} \rightarrow 0$ and $d_n \log p_n/(s_n^2 n) \rightarrow 0$. For the true regression coefficient, $|\gamma_0^n| \sqrt{\log p_n/n} \rightarrow 0$, $\log n \log p_n/(n \rho_{1n}^2) \rightarrow 0$ and $\rho_{2n}/\sqrt{\log p_n} \rightarrow 0$ as $n \rightarrow \infty$.*

This assumption resembles the dimension assumption in Cao et al. (2019a), and is a much weaker assumption for high dimensional covariance asymptotic than for example, Xiang et al. (2015); Banerjee and Ghosal (2014, 2015) and Cao et al. (2019c). Here we essentially allow the dimension of our covariance matrix to grow slower than $\exp(n/d_n^2)$. Recall that s_n is the smallest (in absolute value) non-zero off-diagonal entry in L_0^n , so the second condition in Assumption 2 can also be interpreted as the lower bound for the signal size. This assumption also known as the “beta-min” condition

provides a lower bound for the signal size that is needed for establishing consistency. This type of condition has been used for the exact support recovery of the high-dimensional linear regression models as well as Gaussian DAG models. See for example Yang et al. (2016); Khare et al. (2017); Lee et al. (2018) and Cao et al. (2019c). Assumption 2 also allows the complexity of γ_0^n as well as the non-zero entries of β_0^n to grow with n while stay uniformly bounded by a function of n and p_n . In addition, the assumption on ρ_{1_n} can be viewed as the beta-min condition in the regression context.

Assumption 3. *The hyperparameters in model (3.4) and the MRF prior (3.6) satisfy $\tau^2 \sim \sqrt{\log p_n}$, $a \sim \alpha_1 \log p_n$, and $bn^2/\{(\log n)^2 \log p_n\} \rightarrow 0$ as $n \rightarrow \infty$, where for any positive sequences a_n and b_n , $a_n \sim b_n$ implies that there exist positive constants c and C such that $c \leq \min(a_n/b_n, b_n/a_n) \leq \max(a_n/b_n, b_n/a_n) \leq C$.*

Recall that the parameter a in (3.6) controls the variable inclusion probability, and b reflects that how strongly the inclusion probability of a variable is affected by the inclusion of its neighbors in the DAG. In Section 4.3, we investigate the behavior of the posterior probability evaluated at the true model under $b > 0$ and $b = 0$. In the Bayesian variable selection literature, similar priors corresponding to $a = C \log p_n$ for some constant $C > 0$ and $b = 0$ have been commonly used to obtain selection consistency (Narisetty

and He, 2014; Castillo et al., 2015; Yang et al., 2016). The assumption on the the variance of the slab prior, τ^2 , is required to approach infinity is also stated here to ensure desired model selection consistency.

Assumption 4. Let $q_n = O(p_n^{-\alpha_1})$ for some constant $\alpha_1 > 0$ and R_n in model (3.5) and (3.6) satisfy $R_n \sim n/\log n$ and $bR_n^2/\log p_n \rightarrow 0$ as $n \rightarrow \infty$.

This assumption provides the rate at which the edge probability q_n needs to approach zero. It also states that the prior on the space of the $2^{\binom{p_n}{2}}$ possible models, places zero mass on unrealistically large models. Note that q_n is of slower rate approaching zero compared to the one in Cao et al. (2019c), which helps avoiding the potential computation limitation such as simulation results always favor the most sparse model. This assumption also states that the MRF prior on the space of the 2^{p_n} possible models, places zero mass on unrealistically large models (see similar assumptions in Shin et al. (2018); Narisetty and He (2014) in the context of regression).

Assumption 5. For every $n \geq 1$, the hyperparameters for the DAG-Wishart prior $\pi_{U_n, \alpha}^{\Theta_{\mathcal{D}_n}}$ satisfy (i) $2 < \alpha_i(\mathcal{D}_n) - \nu_i(\mathcal{D}_n) < c$ for every \mathcal{D}_n and $1 \leq i \leq q_n$, and (ii) $0 < \delta_1 \leq \text{eig}_1(U_n) \leq \text{eig}_p(U_n) \leq \delta_2 < \infty$. Here c, δ_1 and δ_2 are constants not depending on n .

This assumption provides mild restrictions on the hyperparameters for the

4.1 Posterior ratio consistency of γ and \mathcal{D}

DAG-Wishart distribution. The assumption $2 < \alpha_i(\mathcal{D}) - \nu_i(\mathcal{D})$ establishes prior propriety. The assumption $\alpha_i(\mathcal{D}) - \nu_i(\mathcal{D}) < c$ implies that the shape parameter $\alpha_i(\mathcal{D})$ can only differ from $\nu_i(\mathcal{D})$ (number of parents of i in \mathcal{D}) by a constant which does not vary with n . Additionally, the eigenvalues of the scale matrix U_n are assumed to be uniformly bounded in n .

For the rest of this paper, $p_n, \Omega_0^n, \Sigma_0^n, L_0^n, D_0^n, \mathcal{D}_0^n, \mathcal{D}^n, d_n, q_n, \beta_n, \gamma_n, \tau_n, A_n$ will be denoted as $p, \Omega_0, \Sigma_0, L_0, D_0, \mathcal{D}_0, \mathcal{D}, d, q, \beta, \gamma, \tau, A$ as needed for notational convenience and ease of exposition. We now state and prove the main joint variable and DAG selection consistency results.

4.1 Posterior ratio consistency of γ and \mathcal{D}

In this section, we show that our method guarantees the posterior ratio consistency of γ and \mathcal{D} . Although Peterson et al. (2016) consider a similar network-structured regression model, theoretical properties of Bayesian models such as posterior ratio consistency and joint selection consistency have not been established yet up to our knowledge. We first establish the posterior ratio consistency with respect to \mathcal{D} under the true variable indicator γ_0 . Theorem 1 says that the true DAG will be the mode of the posterior distribution with probability tending to 1 as $n \rightarrow \infty$ under fixed γ_0 .

4.1 Posterior ratio consistency of γ and \mathcal{D}

Theorem 1. *Under Assumptions 1, 2, 4 and 5,*

$$\max_{\mathcal{D} \neq \mathcal{D}_0} \frac{\pi(\gamma_0, \mathcal{D} | Y, X)}{\pi(\gamma_0, \mathcal{D}_0 | Y, X)} \xrightarrow{\bar{P}} 0, \quad \text{as } n \rightarrow \infty.$$

Remark 2. We would like to point out that the posterior ratio consistency for DAG is achieved under a given parent ordering, where that all the edges are directed from larger vertices to smaller vertices. For several applications in genetics and environmental sciences, a location or time based ordering of variables is naturally available. For temporal data, a natural ordering of variables is provided by the time at which they are observed. In quantitative molecular applications, the variables can be genes or SNPs located on a chromosome, and their spatial location provides a natural ordering. See Huang et al. (2006); Shojaie and Michailidis (2010); Yu and Bien (2016); Khare et al. (2017) and references therein.

The next theorem establishes the posterior ratio consistency with respect to γ under DAG \mathcal{D} . This notion of consistency implies that the true variable indicator γ_0 will be the mode of the posterior distribution with probability tending to 1 as $n \rightarrow \infty$ under fixed \mathcal{D} .

Theorem 2. *Under Assumptions 1-5, the following holds:*

$$\max_{(\gamma, \mathcal{D}) \neq (\gamma_0, \mathcal{D}_0)} \frac{\pi(\gamma, \mathcal{D} | Y, X)}{\pi(\gamma_0, \mathcal{D}_0 | Y, X)} \xrightarrow{\bar{P}} 0, \quad \text{as } n \rightarrow \infty.$$

4.2 Strong selection consistency of γ and \mathcal{D}

Remark 3. Based on a reviewer's comment, by carefully examining the proof of Theorem 2, we find out that even under a DAG with mis-specified ordering, the consistency result for γ under fixed \mathcal{D} will still hold. We also investigate the performance of the proposed method under mis-specified ordering in Section 5. The results suggest that our method recovers the true variable indicator γ_0 well even in the mis-specified case.

From Theorem 1, Theorem 2 and the fact that

$$\frac{\pi(\gamma, \mathcal{D}|Y, X)}{\pi(\gamma_0, \mathcal{D}_0|Y, X)} = \frac{\pi(\gamma_0, \mathcal{D}|Y, X)}{\pi(\gamma_0, \mathcal{D}_0|Y, X)} \times \frac{\pi(\gamma, \mathcal{D}|Y, X)}{\pi(\gamma_0, \mathcal{D}|Y, X)},$$

we can obtain the joint posterior ratio consistency with respect to both γ and \mathcal{D} . It implies that the true variable indicator and DAG, $(\gamma_0, \mathcal{D}_0)$, will be the mode of the posterior distribution with probability tending to 1.

Theorem 3. *Under Assumptions 1-5, the following holds:*

$$\max_{(\gamma, \mathcal{D}) \neq (\gamma_0, \mathcal{D}_0)} \frac{\pi(\gamma, \mathcal{D}|Y, X)}{\pi(\gamma_0, \mathcal{D}_0|Y, X)} \xrightarrow{\bar{P}} 0 \quad \text{as } n \rightarrow \infty,$$

which implies that

$$\bar{P}((\hat{\gamma}, \hat{\mathcal{D}}) = (\gamma_0, \mathcal{D}_0)) \rightarrow 1, \quad \text{as } n \rightarrow \infty.$$

4.2 Strong selection consistency of γ and \mathcal{D}

In this section, we establish the joint strong selection consistency with respect to both γ and \mathcal{D} . Theorem 4 shows that the posterior probability

4.2 Strong selection consistency of γ and \mathcal{D}

assigned to the true variable indicator γ_0 and the true underlying graph \mathcal{D}_0 grows to 1 as $n \rightarrow \infty$. We call this property the joint strong selection consistency. Note that the result given in Theorem 3 does not guarantee the joint strong selection consistency.

Theorem 4. *Under Assumptions 1-5, if we further assume $\alpha_1 > 2\kappa$ for some constant $\kappa > 1$, the following holds:*

$$\pi(\gamma_0, \mathcal{D}_0 | Y, X) \xrightarrow{\bar{P}} 1 \quad \text{as } n \rightarrow \infty.$$

We would like to point out that the condition on α_1 , which controls the rate of independent Bernoulli probability specified in Assumption 4, is only needed for strong selection consistency (Theorem 4). Similar restrictions on the hyperparameters have been considered for establishing consistency properties in the regression setup (Yang et al., 2016; Lee et al., 2018; Cao et al., 2019b). The model selection consistency for the posterior mode in Theorem 3 does not require any restriction on α_1 .

All the aforementioned theorems are based on known σ^2 , which tends to be not flexible enough, as in real applications, the underlying true variance often remains unavailable. Therefore, we introduce the following corollary for a fully Bayesian hierarchical approach, where an appropriate inverse-gamma prior is imposed on σ^2 . It turns out that even with the unknown

4.3 Behavior of the posterior probability when $b = 0$

σ^2 , strong model selection consistency still holds under the same conditions given in Theorem 4.

Corollary 1. *Suppose σ^2 is unknown and a proper inverse-gamma density with some positive constant parameters (a_0, b_0) is placed on σ^2 . Under Assumptions 1-5, and $\alpha_1 > 2\kappa$ for some constant $\kappa > 1$, the following holds:*

$$\pi(\gamma_0, \mathcal{D}_0 | Y, X) \xrightarrow{\bar{P}} 1 \quad \text{as } n \rightarrow \infty.$$

4.3 Behavior of the posterior probability when $b = 0$

In this section, we aim to examine the behavior of the posterior probability for $(\gamma_0, \mathcal{D}_0)$ corresponding to two different scenarios when the MRF prior parameter $b > 0$ and $b = 0$ respectively. The goal is to show that under certain assumption on the connection between the sparsity patterns in γ_0 and \mathcal{D}_0 , by borrowing the graph information through the MRF prior, the posterior probability assigned to $(\gamma_0, \mathcal{D}_0)$ will increase. In particular, we introduce the following condition with respect to the true sparsity patterns encoded in both the variable indicator and the graph.

Condition 1. The true adjacency matrix G_0 and the true variable indicator γ_0 satisfy $\gamma_{0_i} = \gamma_{0_j} = 1$ whenever $(G_0)_{ij} = 1$ for $1 \leq i, j \leq p$.

Condition 1 essentially assumes that the connected variables through the

underlying true DAG are active. Under this condition, compared with modeling the variable and DAG separately, i.e. $b = 0$, incorporating network information into variable selection through the MRF prior with $b > 0$ will increase the posterior probability assigned to $(\gamma_0, \mathcal{D}_0)$ as illustrated in the following theorem. Proof for Theorem 5 will again be provided in the Supplementary material.

Theorem 5. *Let $\pi_1(\gamma_0, \mathcal{D}_0 | Y, X)$ be the posterior probability evaluated at $(\gamma_0, \mathcal{D}_0)$ under $b > 0$ and $\pi_2(\gamma_0, \mathcal{D}_0 | Y, X)$ be the posterior probability evaluated at $(\gamma_0, \mathcal{D}_0)$ under $b = 0$. The following holds:*

$$\pi_1(\gamma_0, \mathcal{D}_0 | Y, X) > \pi_2(\gamma_0, \mathcal{D}_0 | Y, X).$$

Theorem 5 implies that, under Condition 1, our method achieves joint strong selection consistency without the condition on b stated in Assumption 3, which means the hyperparameter b in the MRF prior does not need to go to zero.

5. Numerical Studies

5.1 Posterior inference

For given positive real values a_0 and $b_0 > 0$, let $IG(a_0, b_0)$ be the inverse-gamma distribution with the shape parameter a_0 and scale parameter b_0 .

5.1 Posterior inference

Then, similar to (3.7), the joint posterior distribution of γ and \mathcal{D} based on (3.1)–(3.6) and $\sigma^2 \sim IG(a_0, b_0)$ is

$$\begin{aligned} & \pi(\gamma, \mathcal{D} \mid Y, X) \\ \propto & \pi(\gamma \mid \mathcal{D}) \pi(\mathcal{D}) \frac{z_{\mathcal{D}}(U + X^T X, n + \alpha(\mathcal{D}))}{z_{\mathcal{D}}(U, \alpha(\mathcal{D}))} \\ & \times \det(I_{|\gamma|} + \tau^2 X_{\gamma}^T X_{\gamma})^{-\frac{1}{2}} \left\{ b_0 + \frac{1}{2} Y^T (I_n + \tau^2 X_{\gamma} X_{\gamma}^T) Y \right\}^{-\frac{n+2a_0}{2}}. \end{aligned}$$

We suggest using a Metropolis-Hastings within Gibbs sampling for posterior inference:

1. Set the initial values $\gamma^{(1)}$ and $\mathcal{D}^{(1)}$.
2. For each $s = 2, \dots, S$,
 - (a) sample $\gamma^{new} \sim q_{\gamma}(\cdot \mid \gamma^{(s-1)})$;
 - (b) set $\gamma^{(s)} = \gamma^{new}$ with the probability

$$p_{acc, \gamma} = \min \left\{ 1, \frac{\pi(\gamma^{new} \mid \mathcal{D}^{(s-1)}, Y, X) q_{\gamma}(\gamma^{(s-1)} \mid \gamma^{new})}{\pi(\gamma^{(s-1)} \mid \mathcal{D}^{(s-1)}, Y, X) q_{\gamma}(\gamma^{new} \mid \gamma^{(s-1)})} \right\},$$

otherwise set $\gamma^{(s)} = \gamma^{(s-1)}$;

- (c) sample $\mathcal{D}^{new} \sim q_{\mathcal{D}}(\cdot \mid \mathcal{D}^{(s-1)})$;
- (d) set $\mathcal{D}^{(s)} = \mathcal{D}^{new}$ with the probability

$$p_{acc, \mathcal{D}} = \min \left\{ 1, \frac{\pi(\mathcal{D}^{new} \mid \gamma^{(s)}, Y, X) q_{\mathcal{D}}(\mathcal{D}^{(s-1)} \mid \mathcal{D}^{new})}{\pi(\mathcal{D}^{(s-1)} \mid \gamma^{(s)}, Y, X) q_{\mathcal{D}}(\mathcal{D}^{new} \mid \mathcal{D}^{(s-1)})} \right\},$$

otherwise set $\mathcal{D}^{(s)} = \mathcal{D}^{(s-1)}$.

The inference for the DAG \mathcal{D} , the steps 2-(c) and 2-(d) in the above algorithm, can be parallelized for each column. For more details, we refer to Cao et al. (2019c) and Lee et al. (2018). We used the proposal kernel $q_\gamma(\cdot | \gamma')$ which gives a new set γ^{new} by changing a randomly chosen nonzero component in γ' to 0 with probability 0.5 or by changing a randomly chosen zero component to 1 randomly with probability 0.5. The same proposal kernels were used for each column of \mathcal{D} .

5.2 Simulation Studies

In this section, we demonstrate the performance of the proposed method in various settings. We closely follow but slightly modify the simulation settings in Peterson et al. (2016).

Suppose that we have $X_i = (X_{i1}, \dots, X_{ip})^T \stackrel{i.i.d.}{\sim} N_p(0, \Sigma_0), i = 1, \dots, n$, where $\Sigma_0^{-1} = L_0(D_0)^{-1}L_0^T$, $n = 100$ and $p = 240$. If we consider p as the number of genes, we have 240 genes in this case. Among 240 genes, we assume that there are 40 transcription factors (TFs) and each TF regulates 5 genes. Let TF_j be the index for the j th TF and $(TF_1, TF_2, \dots, TF_{40}) = (6, 12, \dots, 240)$. It corresponds to the DAG \mathcal{D}_0 , the support of L_0 , such that $pa_{TF_j-k}(\mathcal{D}_0) = \{TF_j\}$ for $j = 1, \dots, 40$ and $k = 1, \dots, 5$. Suppose that the TFs independently follow the normal distribution, that is, $X_{TF_j} \stackrel{ind}{\sim}$

5.2 Simulation Studies

$N(0, d_{TF_j})$, where $d_{TF_j} \stackrel{i.i.d.}{\sim} Unif(3, 5)$, for $j = 1, \dots, 40$. We further assume that, given X_{TF_j} , the conditional distribution of the gene X_j that TF_j regulates is $N(X_{TF_j}, d_j)$, where $d_j \stackrel{i.i.d.}{\sim} Unif(3, 5)$ for $j = 1, \dots, 240$. It corresponds to the true modified Cholesky parameter (L_0, D_0) such that $(L_0)_{TF_j, TF_j-k} = 1$ and $D_0 = diag(d_j)$ for $j = 1, \dots, 40$ and $k = 1, \dots, 5$. We simulate the data from

$$Y = X\beta_0 + \epsilon,$$

where $X = (X_1, \dots, X_n)^T$ and $\epsilon \sim N_n(0, \sigma_\epsilon^2 I_n)$ and $\sigma_\epsilon^2 = \|\beta_0\|_2^2/4$. We investigate four settings for the true coefficient vector β_0 as described in Li and Li (2008) and Peterson et al. (2016). In the first setting, it is assumed that $\beta_{0, TF_{1:4}} = (5, -5, 3, -3)^T$, $\beta_{0, TF_j-k} = \beta_{0, TF_j}/\sqrt{10}$ for $j = 1, 2, 3, 4$ and $k = 1, \dots, 5$, and $\beta_{0,j} = 0$ for $j = 25, \dots, 240$. This setting implies that the genes in the same cluster have the same signs for the coefficients. In the second setting, the true coefficient β_0 is the same as the first setting except that the signs are reversed for the two genes that TF_j regulates, i.e., $\beta_{0, TF_j-k} = -\beta_{0, TF_j}/\sqrt{10}$ for $j = 1, 2, 3, 4$ and $k = 1, 2$. This setting implies that the genes in the same cluster might have different signs for the coefficients. The third and fourth settings the same as the first and second settings expect considering 10 instead of $\sqrt{10}$. Thus, they consider smaller signals. We call this simulation setting Scenario 1.

5.2 Simulation Studies

We also investigate a different simulation scenario, say Scenario 2, where the signals in β_0 are small. In this case, there are $p = 150$ genes, 30 TFs and 4 regularized genes for each TF. The precision matrix $\Sigma_0^{-1} = L_0(D_0)^{-1}L_0^T$ is generated by $d_j \stackrel{i.i.d.}{\sim} Unif(2, 5)$ and $(L_0)_{TF_j, TF_j-k} \stackrel{i.i.d.}{\sim} Unif(0.3, 0.7)$. The variance of ϵ is chosen as $\sigma_\epsilon^2 = \|\beta_0\|_2^2$. We consider four settings for the true coefficient vector β_0 . In the first and third settings, β_0 is generated by $\beta_{0,j} \stackrel{i.i.d.}{\sim} Unif(0.5, 1)$ and $\beta_{0,j} \stackrel{i.i.d.}{\sim} Unif(0.2, 1)$ for $j = 1, \dots, 20$, respectively, and $\beta_{0,j} = 0$ for $j = 21, \dots, 150$. In the second and fourth settings, we only change the signs of nonzero entries of β_0 randomly. We call this simulation setting Scenario 2.

Lastly, we consider a setting where the network structure of the covariate X is a undirected graph. We generate the covariates $\tilde{X}_i \stackrel{i.i.d.}{\sim} N_p(0, \Sigma_0)$, $i = 1, \dots, n$, where $n = 100$, $p = 150$, $\Sigma_0 = \tilde{\Sigma}_0 + \{0.01 - eig_1(\tilde{\Sigma}_0)\}I_p$ and

$$(\tilde{\Sigma}_0)_{ij} = \begin{cases} 2 \max\left(1 - \frac{|i-j|}{10}, 0\right), & \text{if } |i-j| \leq 5 \\ 0, & \text{otherwise.} \end{cases}$$

Note that Σ_0 is positive definite. Furthermore, to consider the mis-specified ordering case, we randomly shuffle columns of $\tilde{X} = (\tilde{X}_1, \dots, \tilde{X}_n)^T$ to construct X . We simulate the data from $Y = X\beta_0 + \epsilon$, where $\epsilon \sim N_n(0, \sigma_\epsilon^2 I_n)$ and $\sigma_\epsilon^2 = \|\beta_0\|_2^2/4$. Two settings for the true coefficient vector β_0 are considered. In the first setting, β_0 is generated by $\beta_{0,j} \stackrel{i.i.d.}{\sim} Unif(0.5, 1)$ for

5.2 Simulation Studies

$j = 1, \dots, 10$ and $\beta_{0,j} = 0$ for $j = 11, \dots, 150$. In the second setting, we only change the signs of nonzero entries of β_0 randomly. We call this simulation setting Scenario 3, and the simulation results for this setting are reported at Table 3.

We compare the performance of our joint selection method with other existing variable selection methods: Lasso (Tibshirani, 1996), elastic net (Zou and Hastie, 2005) and the Bayesian joint selection method proposed by Peterson et al. (2016). The tuning parameters in Lasso and elastic net were chosen by 10-fold cross-validation. For Bayesian methods, as discussed by Peterson et al. (2016), we suggest using the hyperparameters $a = 2.75$ and $b = 0.5$ for the MRF prior as default. Furthermore, to show the benefits of joint modeling, we also tried the setting with $b = 0$ which corresponds to the Bayesian method modeling the variable and DAG separately. The other hyperparameters were set at $a_0 = 0.1, b_0 = 0.01, \tau^2 = 1, q = 0.005, U = I_p$ and $\alpha_i(\mathcal{D}) = \nu_i(\mathcal{D}) + 10$ for all $i = 1, \dots, p$. The initial state for γ was set at p -dimensional zero vector, i.e., the empty model, while the initial state for \mathcal{D} was chosen by the CSCS method (Khare et al., 2017). For posterior inference, 5,000 posterior samples were drawn after a burn-in period of 5,000. The indices having posterior inclusion probability larger than 0.5 were included in the final model. The resulting model is called the median

5.2 Simulation Studies

probability model, and when there is a model with posterior probability larger than $1/2$, it coincides with the posterior mode Barbieri and Berger (2004). Since we have proved the joint strong selection consistency (Theorem 4), the two models are asymptotically equivalent in our setting. Thus, although other approaches (for example, see Scott and Carvalho (2008)) can be adapted to give a reasonable estimate of the posterior mode, we use the median probability model as a convenient but asymptotically equivalent alternative.

To evaluate the performance of variable selection, the sensitivity, specificity, area under the curve (AUC), Matthews correlation coefficient (MCC), the number of errors (#Error) and mean-squared prediction error (MSPE) are reported at Tables 1, 2 and 3. The criteria are defined as

$$\begin{aligned}\text{Sensitivity} &= \frac{TP}{TP + FN}, \\ \text{Specificity} &= \frac{TN}{TN + FP}, \\ \text{MCC} &= \frac{TP \times TN - FP \times FN}{\sqrt{(TP + FP)(TP + FN)(TN + FP)(TN + FN)}}, \\ \text{\#Error} &= FP + FN, \\ \text{MSPE} &= \frac{1}{n_{\text{test}}} \sum_{i=1}^{n_{\text{test}}} (\hat{Y}_i - Y_{\text{test},i})^2,\end{aligned}$$

where TP, TN, FP and FN are true positive, true negative, false positive and false negative, respectively. The AUC is calculated based on the true

5.2 Simulation Studies

positive rate (Sensitivity) and the false positive rate (1–Specificity) for Bayesian methods with varying thresholds. To draw the AUC, for each threshold, the indices having posterior inclusion probability larger than a given threshold were included in the final model. The AUCs for the regularization methods are omitted. We denote $\hat{Y}_i = X_i^T \hat{\beta}$, where $\hat{\beta}$ is the estimated coefficient based on each method. For Bayesian methods, the usual least square estimates based on the selected support were used as $\hat{\beta}$. We generated test samples $Y_{\text{test},1}, \dots, Y_{\text{test},n_{\text{test}}}$ with $n_{\text{test}} = 100$ to calculate the MSPE.

Based on Tables 1 and 2, we notice that Bayesian joint selection methods tend to have better specificity and MCC, while the regularization methods (Lasso and elastic net) have better sensitivity. As discussed by Peterson et al. (2016), this seems natural because the regularization methods based on cross-validation tend to include many redundant variables. It leads to relatively larger number of errors for the regularization methods compared with those for the Bayesian joint selection methods. We also found that the joint Bayesian selection method proposed in this paper (Joint.CL ($b = 1/2$)) works better than that proposed by Peterson et al. (2016) (Joint.P) in terms of performance measures in Tables 1 and 2 except the AUC. In fact, the two Bayesian joint selection methods are quite similar to each other except

5.2 Simulation Studies

Table 1: The summary statistics for Scenario 1 are represented for each setting. Different setting means different choice of the true coefficient β_0 . Sens and Spec are sensitivity and specificity, respectively. Joint.CL: the Bayesian joint selection method proposed in this paper. Joint.P: the Bayesian joint selection method suggested by Peterson et al. (2016). Elastic: elastic net.

	Setting 1						Setting 2					
	Sens	Spec	AUC	MCC	#Error	MSPE	Sens	Spec	AUC	MCC	#Error	MSPE
Joint.CL ($b = \frac{1}{2}$)	0.8750	0.9861	0.9937	0.8611	6	69.1445	0.8750	0.9954	0.9894	0.9049	4	56.4885
Joint.CL ($b = 0$)	0.7500	0.9815	0.9601	0.7605	10	96.9889	0.3333	1.0000	0.9058	0.5571	16	142.2708
Joint.P	0.8750	0.9861	0.9838	0.8611	6	71.0443	0.7500	0.9954	0.9958	0.8282	7	73.7870
Lasso	1.0000	0.8056	.	0.5412	42	45.5522	0.7083	0.8519	.	0.4170	39	106.0526
Elastic	1.0000	0.9352	.	0.7685	14	41.8631	0.8750	0.8426	.	0.5122	37	92.6665
	Setting 3						Setting 4					
	Sens	Spec	AUC	MCC	#Error	MSPE	Sens	Spec	AUC	MCC	#Error	MSPE
Joint.CL ($b = \frac{1}{2}$)	0.2083	0.9907	0.8493	0.3549	21	42.5213	0.3750	1.0000	0.7373	0.5922	15	30.3394
Joint.CL ($b = 0$)	0.1667	0.9907	0.7117	0.3025	22	42.7116	0.3333	0.9954	0.7619	0.5191	17	35.0479
Joint.P	0.2500	0.9907	0.8559	0.4023	20	40.3569	0.2917	0.9954	0.8954	0.4797	18	35.7181
Lasso	1.0000	0.8241	.	0.5648	38	32.1919	0.6667	0.8102	.	0.3362	49	40.7437
Elastic	1.0000	0.9444	.	0.7935	12	29.3908	0.6250	0.8935	.	0.4261	32	34.9673

the graph structure they consider. In these simulation scenarios, the DAG structure seems more appropriate because clearly there are parents (TFs genes) and children (regularized genes for each TF). Thus, our method would be preferable in this case. Furthermore, based on asymptotic results, one can expect that our method will give accurate inference results as we have more observations, while asymptotic properties of the Bayesian method proposed by Peterson et al. (2016) are still in question. Lastly,

5.2 Simulation Studies

Table 2: The summary statistics for Scenario 2 are represented for each setting. Different setting means different choice of the true coefficient β_0 .

	Setting 1						Setting 2					
	Sens	Spec	AUC	MCC	#Error	MSPE	Sens	Spec	AUC	MCC	#Error	MSPE
Joint.CL ($b = \frac{1}{2}$)	0.7500	0.9923	0.9362	0.8174	6	20.9925	0.6000	1.0000	0.8933	0.7518	8	15.8789
Joint.CL ($b = 0$)	0.6000	1.0000	0.9200	0.7518	8	29.6691	0.6500	0.9923	0.8790	0.7506	8	23.3007
Joint.P	0.6500	1.0000	0.9842	0.7854	7	15.4705	0.5000	1.0000	0.9081	0.6814	10	19.2450
Lasso	1.0000	0.8308	.	0.6290	22	14.8092	0.9000	0.7692	.	0.4877	32	13.4260
Elastic	0.9500	0.9077	.	0.7201	13	18.9942	0.8000	0.8615	.	0.5371	22	14.5779

	Setting 3						Setting 4					
	Sens	Spec	AUC	MCC	#Error	MSPE	Sens	Spec	AUC	MCC	#Error	MSPE
Joint.CL ($b = \frac{1}{2}$)	0.7500	1.0000	0.9537	0.8498	5	6.6246	0.6500	1.0000	0.8398	0.7854	7	7.4111
Joint.CL ($b = 0$)	0.4000	1.0000	0.9631	0.6051	12	20.3681	0.3000	1.0000	0.7962	0.5204	14	12.7521
Joint.P	0.6500	1.0000	0.9811	0.7854	7	11.6528	0.4500	1.0000	0.9057	0.6441	11	9.2049
Lasso	0.9500	0.8154	.	0.5754	25	8.2451	0.8500	0.7462	.	0.4299	36	7.7223
Elastic	0.9500	0.8923	.	0.6912	15	10.7742	0.7000	0.8846	.	0.5032	21	7.8241

the results show that our joint modeling (Joint.CL ($b = 1/2$)) significantly improves the performance of variable selection compared with modeling the variable and DAG separately (Joint.CL ($b = 0$)). These suggest that joint modeling approach actually improves the performance of variable selection by borrowing information from the DAG structure.

Table 3 shows the results for Scenario 3, where the true network structure for X is a undirected graph and the ordering is mis-specified. Even in this case, our joint modeling method provides comparable performance to that of Peterson et al. (2016), which is designed for undirected graphs. Similar to Scenarios 1 and 2, regularization methods do not work well compared

Table 3: The summary statistics for Scenario 3 are represented for each setting. Different setting means different choice of the true coefficient β_0 .

	Setting 1						Setting 2					
	Sens	Spec	AUC	MCC	#Error	MSPE	Sens	Spec	AUC	MCC	#Error	MSPE
Joint.CL ($b = \frac{1}{2}$)	1.0000	0.9357	0.9964	0.7018	9	1.6875	1.0000	0.9500	0.9821	0.7475	7	1.6469
Joint.CL ($b = 0$)	1.0000	0.9429	0.9786	0.7237	8	1.7331	1.0000	0.9429	0.9786	0.7237	8	1.7331
Joint.P	1.0000	0.9500	0.9857	0.7475	7	1.6838	1.0000	0.9500	0.9821	0.7475	7	1.6838
Lasso	1.0000	0.3571	.	0.1890	90	1.8121	1.0000	0.3571	.	0.1890	90	1.8121
Elastic	1.0000	0.6714	.	0.3463	46	1.6770	1.0000	0.6286	.	0.3184	52	1.6969

with Bayesian methods in our settings.

6. Discussion

In this paper, we work in a regression setting, where the predictors are both relevant to a response variable of interest and functionally related to one another via a Gaussian DAG model. In particular, we consider a hierarchical multivariate regression model with DAG-Wishart priors on the covariance matrix for the predictors, spike and slab priors on regression coefficients, independent Bernoulli priors for each edge in the DAG, and a MRF prior linking the variable indicators to the graph structure. Under high-dimensional settings and standard regularity assumptions, when the underlying variance σ^2 is available, we establish both posterior ratio consistency and strong selection consistency for estimating the variable and the graph for the covariates jointly. When the underlying response variance is

unknown and an appropriate inverse gamma prior is placed on σ^2 , we also establish the joint selection consistency under the same regularity conditions. Finally, through simulation studies, we demonstrate that the model studied in this paper can outperform existing state-of-the-art methods in selecting network-structured predictors including both penalized likelihood and Bayesian approaches in several settings. For future studies, we intend to explore other types of priors over the graph space and on the regression coefficients to see if the consistency and better simulation performance can both be achieved under weakened assumptions.

Supplementary Materials

Supplementary material includes the proofs for main results and other auxiliary results.

Acknowledgements

We thank Dr. Christine Peterson for sharing the code to implement joint Bayesian variable and graph selection method in Peterson et al. (2016). We would like to thank two referees for their valuable comments which have led to improvements of an earlier version of the paper. This research was supported by the Simons Foundation's collaboration grant (No.635213) and the National Research Foundation of Korea (NRF) grant funded by the

Korea government (MSIT) (No.2019R1F1A1059483).

References

- Banerjee, S. and S. Ghosal (2014). Posterior convergence rates for estimating large precision matrices using graphical models. *Electron. J. Stat.* 8, 2111–2137.
- Banerjee, S. and S. Ghosal (2015). Bayesian structure learning in graphical models. *J. Multivariate Anal.* 136, 147–162.
- Barbieri, M. M. and J. O. Berger (2004). Optimal predictive model selection. *Ann. Statist.* 32(3), 870–897.
- Ben-David, E., T. Li, H. Massam, and B. Rajaratnam (2016). High dimensional Bayesian inference for gaussian directed acyclic graph models. *Technical Report* <http://arxiv.org/abs/1109.4371>.
- Bickel, P. J. and E. Levina (2008). Regularized estimation of large covariance matrices. *Ann. Statist.* 36, 199–227.
- Cao, X., K. Khare, and M. Ghosh (2019a). Consistent Bayesian sparsity selection for high-dimensional Gaussian DAG models with multiplicative and beta-mixture priors. <https://arxiv.org/abs/1903.03531>.
- Cao, X., K. Khare, and M. Ghosh (2019b). High-dimensional posterior consistency for hierarchical non-local priors in regression. *Bayesian Anal.*, to appear.
- Cao, X., K. Khare, and M. Ghosh (2019c). Posterior graph selection and estimation consistency

REFERENCES

- for high-dimensional Bayesian DAG models. *Ann. Statist.* 47(1), 319–348.
- Castillo, I., J. Schmidt-Hieber, and A. van der Vaart (2015). Bayesian linear regression with sparse priors. *Ann. Statist.* 43, 1986–2018.
- Chekouo, T., F. C. Stingo, J. D. Doecke, and K.-A. Do (2015). miRNA–target gene regulatory networks: A Bayesian integrative approach to biomarker selection with application to kidney cancer. *Biometrics* 71(2), 428–438.
- Chekouo, T., F. C. Stingo, M. Guindani, and K.-A. Do (2016). A Bayesian predictive model for imaging genetics with application to schizophrenia. *Ann. Appl. Stat.* 10(3), 1547–1571.
- Dobra, A. (2009). Variable selection and dependency networks for genomewide data. *Biostatistics* 10(4), 621–639.
- El Karoui, N. (2008). Spectrum estimation for large dimensional covariance matrices using random matrix theory. *Ann. Statist.* 36, 2757–2790.
- George, E. I. and R. E. McCulloch (1993). Variable selection via Gibbs sampling. *J. Amer. Statist. Assoc.* 88(423), 881–889.
- Huang, J., N. Liu, M. Pourahmadi, and L. Liu (2006). Covariance selection and estimation via penalised normal likelihood. *Biometrika* 93, 85–98.
- Ishwaran, H., U. B. Kogalur, and J. S. Rao (2005). Spike and slab variable selection: Frequentist and Bayesian strategies. *Ann. Statist.* 33(2), 730–773.
- Johnson, V. E. and D. Rossell (2012). Bayesian model selection in high-dimensional settings.

REFERENCES

- Journal of the American Statistical Association* 107(498), 649–660.
- Khare, K., S. Oh, S. Rahman, and B. Rajaratnam (2017). A convex framework for high-dimensional sparse Cholesky based covariance estimation in Gaussian DAG models. *Preprint, Department of Statistics, University of Florida.*
- Lee, K., J. Lee, and L. Lin (2018). Minimax posterior convergence rates and model selection consistency in high-dimensional DAG models based on sparse Cholesky factors. *Ann. Statist., to appear.*
- Li, C. and H. Li (2008). Network-constrained regularization and variable selection for analysis of genomic data. *Bioinformatics* 24(9), 1175–1182.
- Li, C. and H. Li (2010). Variable selection and regression analysis for graph-structured covariates with an application to genomics. *Ann. Appl. Stat.* 4(3), 1498–1516.
- Li, F. and R. Zhang (2010). Bayesian variable selection in structured high-dimensional covariate spaces with applications in genomics. *J. Amer. Statist. Assoc* 105(491), 1202–1214.
- Liu, F., S. Chakraborty, F. Li, Y. Liu, and A. C. Lozano (2014). Bayesian regularization via graph Laplacian. *Bayesian Anal.* 9(2), 449–474.
- Narisetty, N. and X. He (2014). Bayesian variable selection with shrinking and diffusing priors. *Ann. Statist.* 42, 789–817.
- Pan, W., B. Xie, and X. Shen (2010). Incorporating predictor network in penalized regression with application to microarray data. *Biometrics* 66(2), 474–484.

REFERENCES

- Peng, B., D. Zhu, B. P. Ander, X. Zhang, F. Xue, F. R. Sharp, and X. Yang (2013). An integrative framework for Bayesian variable selection with informative priors for identifying genes and pathways. *PLOS ONE* 8(7), 1–16.
- Peterson, C. B., F. C. Stingo, and M. Vannucci (2016). Joint Bayesian variable and graph selection for regression models with network-structured predictors. *Stat. Med.* 35(7), 1017–1031.
- Pourahmadi, M. (2007). Cholesky decompositions and estimation of a covariance matrix: Orthogonality of variance–correlation parameters. *Biometrika* 94, 1006–1013.
- Scott, J. G. and C. M. Carvalho (2008). Feature-inclusion stochastic search for Gaussian graphical models. *J. Comput. Graph. Statist.* 17(4), 790–808.
- Shin, M., A. Bhattacharya, and V. Johnson (2018). Scalable Bayesian variable selection using nonlocal prior densities in ultrahigh-dimensional settings. *Statist. Sinica* 28, 1053–1078.
- Shojaie, A. and G. Michailidis (2010). Penalized likelihood methods for estimation of sparse high-dimensional directed acyclic graphs. *Biometrika* 97, 519–538.
- Stingo, F. C., Y. A. Chen, M. G. Tadesse, and M. Vannucci (2011). Incorporating biological information into linear models: A Bayesian approach to the selection of pathways and genes. *Ann. Appl. Stat.* 5(3), 1978–2002.
- Stingo, F. C. and M. Vannucci (2010). Variable selection for discriminant analysis with Markov random field priors for the analysis of microarray data. *Bioinformatics* 27(4), 495–501.

REFERENCES

- Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. *J. R. Stat. Soc. Ser. B Stat. Methodol.* 58(1), 267–288.
- Xiang, R., K. Khare, and M. Ghosh (2015). High dimensional posterior convergence rates for decomposable graphical models. *Electron. J. Stat.* 9, 2828–2854.
- Yang, Y., M. J. Wainwright, and M. I. Jordan (2016). On the computational complexity of high-dimensional Bayesian variable selection. *Ann. Statist.* 44(6), 2497–2532.
- Yu, G. and J. Bien (2016). Learning local dependence in ordered data. *arXiv:1604.07451*.
- Zou, H. and T. Hastie (2005). Regularization and variable selection via the elastic net. *J. R. Stat. Soc. Ser. B Stat. Methodol.* 67(2), 301–320.

University of Cincinnati

E-mail: (caox4@ucmail.uc.edu)

Inha University

E-mail: (leekjstat@gmail.com)