

Learning Gaussian DAGs from Network Data

Hangjian Li

LIHANGJIAN123@UCLA.EDU

Oscar Hernan Madrid Padilla

OSCAR.MADRID@STAT.UCLA.EDU

Qing Zhou

ZHOU@STAT.UCLA.EDU

Department of Statistics and Data Science

University of California, Los Angeles

Los Angeles, CA 90095, USA

Editor: Joris M. Mooij

Abstract

Structural learning of directed acyclic graphs (DAGs) or Bayesian networks has been studied extensively under the assumption that the data are independent. We propose a new Gaussian DAG model for dependent data which assumes the observations are correlated according to an undirected network. Under this model, we develop a method to estimate the DAG structure given a topological ordering of the nodes. The proposed method jointly estimates the Bayesian network and the correlations among observations by optimizing a scoring function based on penalized likelihood. We show that under some mild conditions, the proposed method produces consistent estimators after one iteration. Extensive numerical experiments also demonstrate that, by jointly estimating the DAG structure and the sample correlation, our method achieves much higher accuracy in structure learning. When the node ordering is unknown, through experiments on synthetic and real data, we show that our algorithm can be used to estimate the correlations between samples, with which we can de-correlate the dependent data to significantly improve the performance of classical DAG learning methods.

Keywords: Bayesian networks, causal discovery, matrix normal distribution, network data

1. Introduction

Bayesian networks (BNs) with structure given by a directed acyclic graph (DAG) are a popular class of graphical models in statistical learning and causal inference. Extensive research has been done to develop new methods and theories to estimate DAG structures and its parameters from data. In this study, we focus on the Gaussian DAG model defined as follows. Let $\mathcal{G} = (V, E)$ be a DAG that represents the structure of a BN for p random variables X_1, \dots, X_p . The vertex set $V = \{1, \dots, p\}$ represents the set of random variables and the edge set $E = \{(j, i) \in V \times V : j \rightarrow i\}$ represents the directed edges in \mathcal{G} . Let $\Pi_i = \{j \in V : (j, i) \in E\}$ denote the parent set of the vertex i . A data matrix $X \in \mathbb{R}^{n \times p}$ is generated by the following Gaussian linear structural equations induced by \mathcal{G} :

$$X_j = \sum_{k \in \Pi_j} \beta_{kj} X_k + \varepsilon_j, \quad \varepsilon_j = (\varepsilon_{1j}, \dots, \varepsilon_{nj}) \sim \mathcal{N}_n(0, \omega_j^2 I_n), \quad (1)$$

for $j = 1, \dots, p$, where X_j is the j th column in X , ω_j^2 the error variance, and $B = (\beta_{kj})_{p \times p}$ is the weighted adjacency matrix (WAM) of \mathcal{G} such that $\beta_{kj} \neq 0$ if and only if $(k, j) \in E$, and $\beta_{jj} = 0$. The errors $\{\varepsilon_j\}$ are independent and ε_j is independent of X_k for $k \in \Pi_j$. The goal is to estimate the structure of \mathcal{G} from X , which is equivalent to estimating the support of B .

A key assumption under (1) is that the rows of X are jointly independent since the covariance matrix of each ε_j is diagonal. Under such i.i.d. assumption, many structure learning algorithms for DAGs have been developed, which can be largely categorized into three groups: score-based, constraint-based, and hybrid of the two. Score-based methods search for the optimal DAG by maximizing a scoring function such as minimum description length (Roos, 2017), BIC (E. Schwarz, 1978), and Bayesian scores (Heckerman et al., 1995; Cooper and Herskovits, 1992) with various search strategies, such as order-based search (Scanagatta et al., 2016; Schmidt et al., 2007; Ye et al., 2021), greedy search (Ramsey et al., 2017; Chickering, 2003), and coordinate descent (Fu and Zhou, 2013; Aragam and Zhou, 2015; Gu et al., 2019). Constraint-based methods, such as the PC algorithm in Spirtes et al. (2000), perform conditional independence tests among variables to construct a skeleton and then proceed to orient some of the edges. There are also hybrid methods such as in Tsamardinos et al. (2006) and Gasse et al. (2014) that combine the above two approaches.

In real applications, however, it is common for observations to be dependent as in network data, which violate the i.i.d. assumption for the aforementioned methods. For example, when modeling the characteristics of an individual in a social network, the observed characteristics from different individuals can be dependent because they belong to the same social group such as friends, family, and colleagues who often share similar features. Another example appears when modeling a gene regulatory network from individuals that are potentially genetically linked. When estimating brain functional networks, we often have a matrix of fMRI measurements for each individual, $X \in \mathbb{R}^{T \times \nu}$, across T time points, and ν brain regions of interest. The existence of correlations across both time points and brain regions often renders the estimates of standard graphical modeling methods inaccurate (Kundu and Risk, 2020). Motivated by these applications, we are interested in developing a Gaussian DAG model that can take into account the dependence between observations. Based on this model, we will develop a learning algorithm that can simultaneously infer the DAG structure and the sample dependencies. Moreover, since many real-world networks are sparse, we also want our method to be able to learn a sparse DAG and scale to a large number of vertices. A sparsity constraint on the estimated DAG can also effectively prevent overfitting and greatly improve the computational efficiency. Lastly, we would like to have theoretical guarantees on the consistency and finite-sample accuracy of the estimators. With these requirements in mind, we seek to

1. Develop a novel Bayesian network model for network data;
2. Develop a method that can jointly estimate a sparse DAG and the sample dependencies under the model;
3. Establish finite-sample error bound and consistency of our estimators;
4. Achieve good empirical performance on both synthetic and real data sets.

When the data are no longer i.i.d., we need a more general model than (1) to capture between-individual dependence, so what we propose is a *union model* where, in addition to using directed edges to describe in-sample relationships, we will use undirected edges to describe cross-sample relationships. Because X is defined by the Gaussian noise vectors ε_j according to the structural equations in (1), dependence among the rows of X may be introduced by modeling the covariance structure among the variables $\varepsilon_{1j}, \dots, \varepsilon_{nj}$ in ε_j . Based on this observation, we will use an undirected graph H to define the sparsity pattern in the precision matrix of ε_j . When H is an empty graph, the variables in ε_j are independent as in the classical Gaussian DAG model. However, when H is not empty, X follows a more complex matrix normal distribution, and the variance is defined by the product of two covariance matrices, one for the DAG \mathcal{G} and the other for the undirected graph H . Now, one can view our proposed model as a more general model that consists of both the directed DAG \mathcal{G} in (1) and the undirected graph H , where cross-sample dependence is characterized by H and in-sample features follow a distribution specified in (1). Estimating the structure of the DAG \mathcal{G} as well as other model parameters under the sparsity regularization in both graphs is a challenging task. We will start by assuming that a topological ordering π of \mathcal{G} is given so that the search space for DAGs can be largely reduced. However, due to the presence of the second graph for network data, the usual likelihood-based objective function used in traditional score-based methods is nonconvex. Constraint-based methods do not naturally extend to network data, either due to the dependence between individuals in X , which complicates conditional independence tests. In order to find a suitable objective function and develop an optimization algorithm, we exploit the biconvex nature of a regularized likelihood score function and develop an effective blockwise coordinate descent algorithm with a nice convergence property. If the topological ordering of the DAG is unknown, it is impossible to identify a unique DAG from the data due to the Markov equivalence of DAGs (Chickering, 2003). Moreover, due to the lack of independence, it is very difficult to estimate the equivalence class defined by \mathcal{G} . In this case, we take advantage of an invariance property of the matrix normal distribution. Under some sparsity constraint on \mathcal{G} , we show that even with a random ordering, we can still get a good estimate of the covariance of ε_j , which can be used to de-correlate X so that existing DAG learning algorithms can be applied to estimate an equivalence class of \mathcal{G} .

The remainder of the paper is structured as follows. In Section 2 we introduce a novel Gaussian DAG model for network data and discuss its connections with some existing models. We propose a structural learning algorithm for the model and go through its details in Section 3. Section 4 is devoted to our main theoretical results, as well as their implications under various high-dimensional asymptotic settings. Section 5 reports numerical results of our method with detailed comparisons with some competing methods on simulated data. Section 6 presents an application of our method on a real single-cell RNA sequencing data set. All proofs are deferred to the Appendix.

Notations. For the convenience of the reader, we now summarize some notations to be used throughout the paper. We write \mathcal{G}^* and H^* for the true DAG and the true undirected graph, respectively. Let $\Omega := \text{diag}(\omega_j^2)$ be a $p \times p$ diagonal matrix of error variances, B^* denote the true WAM of \mathcal{G}^* , and $s := \sup_j \|\beta_j^*\|_0$ denote the maximum number of parents of any node in \mathcal{G}^* . Furthermore, X_j denotes the j th column of X for $j = 1, \dots, p$, and x_i denotes the i th row of X for $i = 1, \dots, n$. Given two sequences f_n and g_n , we

write $f_n \lesssim g_n$ if $f_n = O(g_n)$, and $f_n \asymp g_n$ if $f_n \lesssim g_n$ and $g_n \lesssim f_n$. Denote by $[p]$ the index set $\{1, \dots, p\}$. For $x \in \mathbb{R}^n$, we denote by $\|x\|_q$ its ℓ_q norm for $q \in [0, \infty]$. For $A \in \mathbb{R}^{n \times m}$, $\|A\|_2 = \sup_v \{\|Av\|_2 : \|v\|_2 \leq 1, v \in \mathbb{R}^m\}$ is the operator norm of A , $\|A\|_f$ is the Frobenius norm of A , $\|A\|_\infty = \max_{i,j} |a_{ij}|$ is the element-wise maximum norm of A , and $\|A\|_\infty = \max_{i \in [n]} \sum_{j=1}^m |a_{ij}|$ is the maximum row-wise ℓ_1 norm of A . Denote by $\sigma_{\min}(A)$ and $\sigma_{\max}(A)$, respectively, the smallest and largest singular values of a matrix A . Let $|S|$ be the size of a set S .

2. A Novel DAG Model for Dependent Data

We model sample dependency through an undirected graph H on n vertices, with each vertex representing an observation $x_i, i \in [n]$, and the edges representing the conditional dependence relations among them. More explicitly, let $A(H)$ be the edge set of H so that

$$(i, j) \notin A(H) \Rightarrow x_i \perp\!\!\!\perp x_j | x_{[n] \setminus \{i, j\}}, \quad \forall i \neq j.$$

Suppose we observe not only the dependent samples $\{x_i\}_{i=1}^n$ but also the graph (network) H . We generalize the structural equation model (SEM) in (1) to

$$X_j = \sum_{k \in \Pi_j} \beta_{kj} X_k + \varepsilon_j, \quad \varepsilon_j = (\varepsilon_{1j}, \dots, \varepsilon_{nj}) \sim \mathcal{N}_n(0, \omega_j^2 \Sigma), \quad (2)$$

where $\Sigma \in \mathbb{R}^{n \times n}$ is positive definite. The support of the precision matrix $\Theta = (\Sigma)^{-1}$ is restricted by $\text{supp}(\Theta) \subseteq A(H)$, where $\text{supp}(\Theta) = \{(i, j) \mid \Theta_{ij} \neq 0\}$ and for notional brevity we put $(i, i) \in A(H)$ for all $i \in [n]$. Note that when $\Sigma = I_n$, the SEM (2) reduces to (1). Hence, the classical Gaussian DAG model in (1) is a special case of our proposed model (2). Under the more general model (2), we face a more challenging structural learning problem: Given dependent data X generated from a DAG \mathcal{G} and the undirected graph H that encodes the sample dependencies, we want to estimate the DAG coefficients B , the noise variance $\Omega = \text{diag}(\omega_j^2)$, and the precision matrix Θ of the samples. Before introducing our method, let us first look at some useful properties of model (2).

2.1 Related Graphical Models

Let us demonstrate the distinction between (1) and (2) using a data matrix $X = (x_{ij})_{n \times p}$ with $n = 5$ units and $p = 4$ variables. Under SEM (1), we model $x_1 = (x_{11}, x_{12}, x_{13}, x_{14}), \dots, x_5 = (x_{51}, x_{52}, x_{53}, x_{54})$ using the same DAG, as shown in Figure 1a, and assume that they are independent. In contrast, the model proposed in (2) allows units to be dependent by relaxing the independence assumption among $x_{1j}, x_{2j}, x_{3j}, x_{4j}, x_{5j}$ for $j \in \{1, 2, 3, 4\}$. The dependence between units is induced by the dependence between background variables $\varepsilon_{1j}, \dots, \varepsilon_{5j}$ and is modeled by an undirected graph as shown in Figure 1b. In general, the variables x_{i1}, \dots, x_{ip} in each unit satisfy the same conditional independence constraints defined by a DAG, while the background variables $\varepsilon_{1j}, \dots, \varepsilon_{nj}$ across the n units are dependent. When estimating the DAG structure with such data, the correlations among individuals will reduce the effective sample size. Therefore, we need to take into account the distribution of the correlated ε_i . Note that our model excludes edges between two variables of any two different units. That is, we do not consider an edge of any type between x_{ij} and $x_{k\ell}$ for $i \neq k$ and $j \neq \ell$.

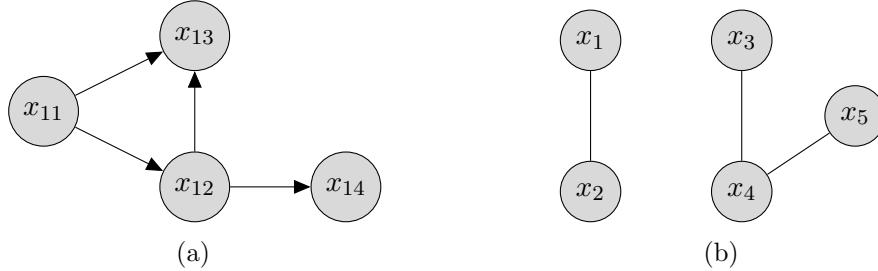


Figure 1: Graphical representation of the model (2) as a product of a DAG over features and an undirected graph over units (individuals). (a) The same DAG structure is shared by the features of every individual. (b) An undirected network of five individuals to represent cross-unit dependence.

The model can also be viewed as a special case of a chain graph (Lauritzen, 1996; Lauritzen and Richardson, 2002), where the same variable across all units, x_{1j}, \dots, x_{nj} , form a chain component. However, the key structure in our model is that our graph over all variables and all units, $\{x_{ij}, i \in [n], j \in [p]\}$, is a Kronecker product of a DAG over p variables and an undirected graph over n units (more details in Section 2.2). This special structure is utilized in both our estimation algorithm introduced in Section 3 and theoretical analysis in Section 4. If we regard the problem as a general structure learning of a chain graph over all $n \times p$ nodes, we would have only a sample of size one, from which one may not be able to construct consistent estimators in general. Another related model that has been studied for interference analysis is the *segregated graph model* proposed by Shpitser (2015), which represents the equivalence classes of latent variable chain graphs. Bhattacharya et al. (2020) considered a closely related problem of estimating causal effects under *partial interference* (Hudgens and Halloran, 2008). They assume the unit-level causal DAG is known and proposed a method to estimate causal effects under unit dependence. In comparison, we assume that the DAG structure is unknown (see Section 3.2), but instead the network structure among units is a subgraph of a given undirected graph. In fact, we only need to know a block structure among the units, and then the network structure within each block can be learned from the data. Due to these different assumptions, the structure learning algorithms in this work are very different from the one proposed by Bhattacharya et al. (2020). In general, although there is some similarity between the model in (2) and the models proposed in the interference causal inference literature, the motivation for our work is quite different. We consider the structural equation model (1) at the single-unit level as a map from extraneous variables $\varepsilon_1, \dots, \varepsilon_p$ to the observed variables X_1, \dots, X_p . Then we introduce association (symmetric and undirected) among the extraneous variables based on a network (undirected graph).

2.2 Matrix Normal Distribution

Our model (2) defines a matrix normal distribution for X . To see this, note that $\varepsilon = (\varepsilon_{ij})_{n \times p}$ in (2) follows a matrix normal distribution:

$$\varepsilon \sim \mathcal{N}_{n,p}(0, \Sigma, \Omega) \Leftrightarrow \text{vec}(\varepsilon) \sim \mathcal{N}_{np}(0, \Omega \otimes \Sigma),$$

where $\text{vec}(\cdot)$ is the vectorization operator and \otimes is the Kronecker product. Then, the random matrix X satisfies

$$X \sim \mathcal{N}_{n,p}(0, \Sigma, \Psi), \quad (3)$$

where $\Psi = (I - B)^{-\top} \Omega (I - B)^{-1}$. We fix $\omega_1 = 1$ so that the parameters (Ω, Σ) are identifiable under model (2). From the properties of a matrix normal distribution, we can easily prove the following lemma which will come in handy when estimating the row covariance matrix Σ from different orderings of nodes. Given a permutation π of the set $[p]$, define P_π as a permutation matrix such that $hP_\pi = (h_{\pi^{-1}(1)}, \dots, h_{\pi^{-1}(p)})$ for any row vector $h = (h_1, \dots, h_p)$.

Lemma 1 *If X follows the model (2), then for any permutation π of $[p]$ we have*

$$XP_\pi \sim \mathcal{N}_{n,p}(0, \Sigma, P_\pi^\top \Psi P_\pi).$$

Although matrix normal distributions have been studied extensively in the past, the structural learning problem we consider here is quite unique. First of all, previous studies on the matrix normal model usually assume that we observe m copies of X and that MLE exists when $m \geq \max\{p/n, n/p\} + 1$ (Dutilleul, 1999). In our case, we only observe one copy of X and thus the MLE does not exist without additional sparsity constraints. Allen and Tibshirani (2010) proposed to use ℓ_1 regularization to estimate the covariance matrices when $m = 1$, but the estimation relies on the assumption that the model is transposable, which means that the two components (Σ, Ψ) of the covariance are symmetric and can be estimated in a symmetric way. However, in the model (2), the two covariance components have different structural constraints and cannot be estimated in the same way. Lastly, practitioners are often interested in estimating large Bayesian networks with hundreds or more nodes under certain sparsity assumptions on the WAM B . For example, for methods that minimize a score function to estimate the covariances, adding a sparsity regularization term on $\Psi = (I - B)^{-\top} \Omega (I - B)^{-1}$ to the score function does not necessarily lead to a sparse estimate of B . In this paper, we propose a new DAG estimation method under the assumption that both the underlying undirected network among individuals and the Bayesian network are sparse. We are not interested in estimating Ψ but a sparse factorization of Ψ represented by the WAM B . This would require imposing sparsity constraints on B itself instead of on Ψ . This is different from the recent work by Tsiligkaridis et al. (2013), Allen and Tibshirani (2010), and Zhou (2014) on the Kronecker graphical lasso.

2.3 Score-equivalence

The likelihood function of the proposed model (2) also satisfies the desired score-equivalence property. To see this, let $\beta_j = (\beta_{1j}, \dots, \beta_{pj})^\top$ be the j th column of the WAM B . Define an $n \times n$ sample covariance matrix of $\varepsilon_1/\omega_1, \dots, \varepsilon_p/\omega_p$ from X as

$$S(\Omega, B) = \frac{1}{p} \sum_{j=1}^p \frac{1}{\omega_j^2} (X_j - X\beta_j)(X_j - X\beta_j)^\top. \quad (4)$$

Then the negative log-likelihood $\ell(B, \Omega, \Theta \mid X)$ from (2) is given by

$$2\ell(B, \Omega, \Theta \mid X) = n \log \det \Omega - p \log \det \Theta + p \text{tr}(\Theta S(\Omega, B)). \quad (5)$$

Due to the dependence among observations, it is unclear whether the well-known score-equivalence property for Gaussian DAGs (Chickering, 2003) still holds for our model. Let $(\hat{B}(\mathcal{G}), \hat{\Omega}(\mathcal{G}), \hat{\Theta}(\mathcal{G}))$ denote the MLE of (B, Ω, Θ) given a DAG \mathcal{G} and the support restriction on Θ . Then, the following theorem confirms the score-equivalence property for our DAG model.

Theorem 2 (*Score equivalence*) *Suppose \mathcal{G}_1 and \mathcal{G}_2 are two Markov equivalent DAGs on the same set of p nodes. If the MLEs $(\hat{B}(\mathcal{G}_m), \hat{\Omega}(\mathcal{G}_m), \hat{\Theta}(\mathcal{G}_m))$, $m = 1, 2$, exist for the matrix $X = (x_{ij})_{n \times p}$, then the two MLEs will give the same log-likelihood value,*

$$\ell(\hat{B}(\mathcal{G}_1), \hat{\Omega}(\mathcal{G}_1), \hat{\Theta}(\mathcal{G}_1) | X) = \ell(\hat{B}(\mathcal{G}_2), \hat{\Omega}(\mathcal{G}_2), \hat{\Theta}(\mathcal{G}_2) | X).$$

This property justifies the evaluation of estimated DAGs using a common model selection criterion such as AIC and BIC. For examples, we show in Section 5 that one can use BIC scores to select the optimal penalty level for our proposed DAG estimation algorithm.

3. Methods

We have discussed the properties of our novel DAG model for dependent data and the unique challenges faced by the structural learning task. In this section, we develop a new method to estimate the parameters in model (2). Our estimator is defined by the minimizer of a score function that derives from a penalized log-likelihood. In order to explain our method, let us start from the penalized negative log-likelihood function:

$$f(B, \Omega, \Theta) := 2\ell(B, \Omega, \Theta | X) + \rho_1(B) + \rho_2(\Theta), \quad B \in \mathcal{D}, \quad (6)$$

where \mathcal{D} is the space of WAMs for DAGs and ρ_1 and ρ_2 are some penalty functions. This loss function is difficult to minimize due to the non-convexity of ℓ and the exponentially large search space of DAGs. One way to reduce the search space is to assume a given topological ordering. Recall that a permutation π of p elements is a bijective map from $[p]$ to $[p]$. A *topological ordering* of a DAG \mathcal{G} with p vertices is a permutation π such that for every edge $(i, j) \in E(\mathcal{G})$, $\pi^{-1}(i) < \pi^{-1}(j)$. Recall that a WAM B is defined as $(\beta_{kj})_{p \times p}$ such that $\beta_{kj} \neq 0$ if and only if $(k, j) \in E(\mathcal{G})$; therefore, given a topological ordering π , we can define a set $\mathcal{D}(\pi)$ of WAMs compatible with π such that all $B \in \mathcal{D}(\pi)$ are strictly upper triangular after permuting its rows and columns according to π . Given a topological ordering π , the loss function (6) becomes

$$\begin{aligned} f(B, \Omega, \Theta) = & -p \log \det \Theta + n \log \det \Omega + \sum_{j=1}^p \frac{1}{\omega_j^2} \|LX_j - LX\beta_j\|_2^2 \\ & + \rho_1(B) + \rho_2(\Theta), \quad B \in \mathcal{D}(\pi), \end{aligned} \quad (7)$$

where L is the Cholesky factor of Θ (i.e. $\Theta = L^\top L$). If $\rho_1(\cdot)$ and $\rho_2(\cdot)$ are convex loss functions and the noise covariance matrix $\Omega = \text{diag}(\omega_j^2)$ is known, (7) will be a bi-convex function in (B, Θ) , which can be minimized using iterative methods such as coordinate descent. Tseng (2001) showed that the coordinate descent algorithm in bi-convex problems converges to a stationary point. Inspired by this observation, we propose the following two-step algorithm:

Step 1: Pre-estimate Ω to get $\widehat{\Omega} = \text{diag}(\widehat{\omega}_j^2)$.

Step 2: Estimate \widehat{B} and $\widehat{\Theta}$ by minimizing a biconvex score function derived from the penalized negative log-likelihood conditioning on $\widehat{\omega}_j$.

Many existing noise estimation methods for high-dimensional linear models can be used to estimate $\widehat{\Omega}$ in Step 1 such as scaled lasso/MCP (Sun and Zhang, 2012), natural lasso (Yu and Bien, 2019), and refitted cross-validation (Fan et al., 2012). We will present the natural estimator of Ω and discuss some other alternatives in Section 4.2. Importantly, the statistical properties of the chosen estimator $\widehat{\Omega}$ in Step 1 will affect the properties of the $\widehat{\Theta}$ and \widehat{B} we get in Step 2, and thus we must choose the estimator carefully. We leave the detailed discussion of the theoretical properties of $\widehat{\Omega}$ and their implications to Section 4. Suppose $\widehat{\Omega}$ is given, we propose the following estimator for Step 2:

$$\begin{aligned} (\widehat{\Theta}, \widehat{B}(\pi)) = \arg \min_{\Theta \succ 0, B \in \mathcal{D}(\pi)} & \left\{ -p \log \det \Theta + \sum_{j=1}^p \frac{1}{\widehat{\omega}_j^2} \|LX_j - LX\beta_j\|_2^2 \right. \\ & \left. + \frac{\lambda_1}{\widehat{\omega}_j^2} \|\beta_j\|_1 + \lambda_2 \|\Theta\|_1 \right\}. \end{aligned} \quad (8)$$

The ℓ_1 regularization on $\beta_j/\widehat{\omega}_j^2$ not only helps promote sparsity in the estimated DAG but also prevents the model from overfitting variables that have small variances. The ℓ_1 regularization on Θ ensures that the estimator is unique and can improve the accuracy of $\widehat{\Theta}$ by controlling the error carried from the previous step. We will discuss how to control the estimation errors in more detail in Section 4.

In Section 3.1, we assume a topological ordering π^* of the true DAG \mathcal{G}^* is known. In this case, we will order the columns of X according to π^* so that for each j , only the first $j-1$ entries in β_j can be nonzero. When minimizing (8), we fix $\beta_{kj} = 0$ for $k \geq j$ and the resulting \widehat{B} is guaranteed to be upper-triangular. If π^* is unknown, we show in Section 3.2 how the score function in (8) is still useful for estimating Θ^* and describe a method of de-correlation so that standard DAG learning methods can be applied on the de-correlated data.

3.1 Block Coordinate Descent

We denote an estimate of the true precision matrix Θ^* at iteration t by $\widehat{\Theta}^{(t)}$. We also write $\widehat{L}^{(t)}$ and L^* for the Cholesky factors of $\widehat{\Theta}^{(t)}$ and Θ^* , respectively. Since (8) is biconvex, it can be solved by iteratively minimizing over Θ and B , i.e., using block coordinate descent. Consider the t th iteration of block coordinate descent. Fixing $\widehat{\Theta}^{(t)}$, the optimization problem in (8) becomes the standard Lasso problem (Tibshirani, 1996) for each j :

$$\widehat{\beta}_j^{(t+1)} = \arg \min_{\beta_j} \frac{1}{2n} \|\widehat{L}^{(t)} X_j - \widehat{L}^{(t)} X \beta_j\|_2^2 + \lambda_n \|\beta_j\|_1, \quad \lambda_n = \lambda_1/(2n), \quad (9)$$

where $\widehat{\Theta}^{(t)} = \widehat{L}^{(t)\top} \widehat{L}^{(t)}$ is the Cholesky decomposition. Since the columns of X are ordered according to π , we can set $\widehat{\beta}_{ij}^{(t+1)} = 0$ for $i = j, j+1, \dots, p$ and reduce the dimension of

feasible β_j in (9) to $j - 1$. In particular, $\hat{\beta}_1^{(t+1)}$ is always a zero vector. Fixing $\hat{B}^{(t+1)}$, solving for $\hat{\Theta}^{(t+1)}$ is equivalent to a graphical Lasso problem with fixed support (Ravikumar et al., 2011)

$$\hat{\Theta}^{(t+1)} = \arg \min_{\Theta \succ 0, \text{ supp}(\Theta) \subseteq A(H^*)} -\log \det \Theta + \text{tr}(\hat{S}^{(t+1)} \Theta) + \lambda_p \|\Theta\|_1, \quad (10)$$

where $\hat{S}^{(t+1)} = S(\hat{\Omega}, \hat{B}^{(t+1)})$ and $\lambda_p = \lambda_2/p$. The details of the method are given in Algorithm 1.

Algorithm 1: Block coordinate descent (BCD) algorithm

Input: $X, \Theta^{(0)}, \hat{\Omega}, \rho, A(H^*), T$
while $\max \left\{ \|\hat{\Theta}^{(t+1)} - \hat{\Theta}^{(t)}\|_f, \|\hat{B}^{(t+1)} - \hat{B}^{(t)}\|_f \right\} > \rho$ and $t < T$ **do**
 for $j = 1, \dots, p$ **do**
 $\hat{\beta}_j^{(t+1)} \leftarrow$ Lasso regression (9)
 end
 $\hat{\Theta}^{(t+1)} \leftarrow$ graphical Lasso with support restriction (10)
 $t \leftarrow t + 1$
end
Output: $\hat{B} \leftarrow \hat{B}^{(t)}, \hat{\Theta} \leftarrow \hat{\Theta}^{(t)}$

As shown in Proposition 3, Algorithm 1 will converge to a stationary point of the objective function (8). The stationary point here is defined as a point where all directional directives are nonnegative (Tseng, 2001).

Proposition 3 *Let $\{(\hat{B}^{(t)}, \hat{\Theta}^{(t)}) : t = 1, 2, \dots\}$ be a sequence generated by Algorithm 1 for any $\lambda_1, \lambda_2 > 0$. Then for almost all $X \in \mathbb{R}^{n \times p}$, every cluster point of $\{(\hat{B}^{(t)}, \hat{\Theta}^{(t)})\}$ is a stationary point of the objective function in (8).*

3.2 Estimating DAGs with Unknown Ordering

Given any permutation π of $[p]$, there exists a DAG \mathcal{G}_π such that (i) π is a topological sort of \mathcal{G}_π and (ii) the joint distribution \mathbb{P} of the p random variables admits a recursive factorization according to \mathcal{G}_π (van de Geer and Bühlmann, 2013; Lauritzen, 1996, Chap. 3). Under the assumption that the true DAG \mathcal{G}^* is sparse, i.e., the number of nonzero entries in β_j^* is at most s for all j , for any random ordering π' we choose, the corresponding DAG $\mathcal{G}_{\pi'}$ is also likely to be sparse so that the number of parents for each node is less than some positive constant s' . Formally, we assume that $s' \ll n/\log p$ uniformly for all permutations π' . Similar assumptions have been made in other works on DAG learning from high-dimensional data, such as Condition 3.4 in van de Geer and Bühlmann (2013) and Theorem 4.1 in Aragam et al. (2019a). Let us randomly pick a permutation π' and apply Algorithm 1 to $X_{\pi'} := X P_{\pi'}$, where $(X P_{\pi'})_{ij} = X_{i\pi'(j)}$. Under the above assumption, the sparsity s' is small compared to the sample size n , and therefore, the estimate $\hat{\beta}'_{ij}$ we get from solving the Lasso problem (9) will be consistent as well (we discuss the error bound on $\hat{\beta}_{ij}$ in detail

in Section 4, Eq.(15)). Moreover, since the covariance Θ^* is invariant to permutations by Lemma 1, the resulting estimate $\widehat{\Theta}$ under the random ordering π' will also be a good estimate of Θ^* (Eq. (18)). Note that the DAG $\mathcal{G}_{\pi'}$ is not necessarily Markov equivalent to the true DAG \mathcal{G}^* . We merely assume that $\mathcal{G}_{\pi'}$ is sparse so that Θ can be accurately estimated even with a random order. With the Cholesky factor \widehat{L} of $\widehat{\Theta}$, we de-correlate the rows of X and treat

$$\widehat{X} = \widehat{L}X, \quad (11)$$

as the new data. Because the row correlations in \widehat{X} vanish, we can apply existing structure learning methods which require independent observations to learn the underlying DAG. We find that this de-correlation step is able to substantially improve the accuracy of structure learning by well-known state-of-the-art methods, such as the greedy equivalence search (GES) (Chickering, 2003) and the PC algorithm (Spirtes et al., 2000). See Section 5 for more details.

4. Main Theoretical Results

In this section, we present our main theoretical results for Algorithm 1 assuming a true ordering is given. We state our main theorem, Theorem 4, in Section 4.1. Section 4.2 is devoted to the error bounds of $\widehat{\Omega}$ using a *natural estimator*, along with some important corollaries. Finally, in Section 4.3, we compare the error rates of our estimators with those in related problems.

Before we start, let us introduce some additional notations used in this section. Let the errors of \widehat{L} and $\widehat{\Theta}$ be defined as $\widehat{\Delta}_{chol} := \widehat{L} - L^*$ and $\widehat{\Delta}_{prec} := \widehat{\Theta} - \Theta^*$. Let $\widehat{\Delta}_j := \widehat{\beta}_j - \beta_j^* \in \mathbb{R}^p$ denote the estimation error of the j th column of B^* . Let

$$\bar{\beta} = \sup_{1 \leq i, j \leq p} |\beta_{ij}^*|, \quad \bar{\omega} = \sup_{1 \leq j \leq p} \omega_j^*, \quad \bar{\psi}^2 = \sup_{1 \leq j \leq p} \Psi_{jj}^*,$$

where $\Psi^* = (I - B^*)^{-\top} \Omega^* (I - B^*)^{-1}$. In the proofs, we also use $X_{i \cdot}$ and $X_{\cdot j}$ to denote the i th row and j th column of X , respectively. Let $\widetilde{X} = L^* X \sim \mathcal{N}(0, \Psi^* \otimes I_n)$ and $\tilde{\varepsilon} = L^* \varepsilon$. Then the rows of \widetilde{X} , i.e. $\tilde{x}_i, i \in [n]$, are i.i.d. from $\mathcal{N}(0, \Psi^*)$.

Recall from Section 3 that $\widehat{\Omega}$ is pre-estimated at Step 1 in our two-step learning procedure. As we will discuss in more detail in Section 4.1, the accuracy of $\widehat{\Theta}$ obtained in Step 2 using Algorithm 1 depends on the accuracy of $\widehat{\Omega}$. Thus, let us define $r(\widehat{\Omega})$ as a measure of the estimation error of $1/\widehat{\omega}_j^2$ as follows:

$$r(\widehat{\Omega}) := \sup_{1 \leq j \leq p} \left| \frac{1}{\widehat{\omega}_j^2} - \frac{1}{\omega_j^{*2}} \right|. \quad (12)$$

Existing methods for estimating the error variance in linear models, such as scaled Lasso (Sun and Zhang, 2012), square-root lasso (Belloni et al., 2011), and natural lasso (Yu and Bien, 2019), often assume independence among samples, which is not necessarily true under our network setting. We will discuss additional assumptions as well as possible choice of consistent estimators of Ω under these assumptions in Section 4.2 under the network setting. But our main result in Section 4.1, in particular the consistency of the estimated DAG parameters \widehat{B} , does not require these additional assumptions (see Remark 10 for more discussion).

4.1 Error Bounds and Consistency of $\widehat{B}^{(1)}$ and $\widehat{\Theta}^{(1)}$

Applying Algorithm 1 with a pre-estimated $\widehat{\Omega}$ as input gives us $\widehat{B}^{(t)}$ and $\widehat{\Theta}^{(t)}$. In this section, we study the error bounds and consistency of $\widehat{B}^{(t)}$ and $\widehat{\Theta}^{(t)}$. Although Algorithm 1 is computationally efficient and has a desirable convergence behavior as described in Proposition 3, there are technical difficulties in establishing the consistency of $\widehat{B}^{(\infty)}$ and $\widehat{\Theta}^{(\infty)}$ after convergence, due to the dependence between $(\widehat{B}^{(t)}, \widehat{\Theta}^{(t)})$ across iterations. However, we show that as long as we have a suitable initial estimate satisfying Assumption 1, Algorithm 1 can produce consistent estimators after one iteration, i.e., $(\widehat{B}^{(1)}, \widehat{\Theta}^{(1)})$ is consistent.

Assumption 1 *There exists a constant $0 < M \leq \sigma_{\min}^2(L^*)$, such that the initial estimate $\widehat{\Theta}^{(0)}$ satisfies*

$$\|\widehat{\Theta}^{(0)} - \Theta^*\|_2 \leq M.$$

Assumption 1 states that the initial estimate $\widehat{\Theta}^{(0)}$ is inside an operator norm ball centered at the true parameter Θ^* with a radius smaller than a constant M . The constant M is less than or equal to the smallest eigenvalue of Θ^* . Since $\sigma_{\min}^2(L^*)$ is a constant, Assumption 1 requires that the initial estimate $\widehat{\Theta}^{(0)}$ is not far from the truth. This is much weaker than, for example, assuming $\widehat{\Theta}^{(0)}$ is a consistent estimator. This condition may be written as $\|\widehat{\Theta}^{(0)} - \Theta^*\|_2 \leq c\lambda_{\min}(\Theta^*)$ for some $c \in (0, 1]$. Intuitively, $\lambda_{\min}(\Theta^*)$ can be understood as the minimum “radius” of Θ^* among all directions. In this sense, this assumption gives a rough scale on how accurate the initial estimate needs to be. This assumption may not be easy to verify in practice, but since it only requires $\widehat{\Theta}^{(0)}$ to be within an ℓ_2 -ball of constant radius around Θ^* , it is not difficult for Assumption 1 to be met if Θ^* is sparse and normalized. Under Assumption 1 we can establish finite-sample error bounds for $(\widehat{B}^{(1)}, \widehat{\Theta}^{(1)})$. Let m denote the maximum degree of the undirected graph H^* and $s = \sup_{1 \leq j \leq p} \|\beta_j^*\|_0$. Define

$$\bar{R}(s, p, n) := \max \left\{ 6\bar{\omega}r(\widehat{\Omega}), \frac{72\bar{\omega}\bar{\psi}s}{b} \sqrt{\frac{\log p \log(\max\{n, p\})^2}{n}} \right\}, \quad (13)$$

where $r(\widehat{\Omega})$ is defined in (12) and $b > 0$ is a constant (see Lemma 7). Following the setup in Ravikumar et al. (2011), the set of non-zero entries in the precision matrix is denoted as $\text{supp}(\Theta^*) := \{(i, j) \mid \Theta_{ij}^* \neq 0\}$. Let us use the shorthand S and S^c to denote the support and its complement in the set $[n] \times [n]$, respectively. Define the following quantities:

$$\begin{aligned} \kappa_{\Sigma^*} &:= \|\Sigma^*\|_{\infty} = \max_{1 \leq i \leq n} \sum_{j=1}^n |\Sigma_{ij}^*|, \\ \kappa_{\Gamma^*} &:= \|(\Gamma_{SS}^*)^{-1}\|_{\infty}, \quad \Gamma_{SS}^* = [\Theta^{*-1} \otimes \Theta^{*-1}]_{SS}. \end{aligned} \quad (14)$$

The quantities κ_{Γ^*} and κ_{Σ^*} defined in (14) measure, respectively, the scale of the entries in Σ^* and the inverse Hessian matrix Γ_{SS}^{*-1} of the graphical Lasso log-likelihood function (10). The L_2 error bound of $\widehat{\Theta}$ (Theorem 4) depends on these constants. In general, they may scale with n and p . To simplify the asymptotic results, we assume the quantities are bounded by a constant as $n, p \rightarrow \infty$ (Corollary 8 and 9). See Ravikumar et al. (2011) for a related discussion.

Theorem 4 Consider a sample matrix X from model (2). Let $\widehat{\Theta}^{(1)}, \widehat{B}^{(1)}$ be the estimates after one iteration of the Algorithm 1, given the initial estimator $\widehat{\Theta}^{(0)}$ satisfying Assumption 1. Suppose $b > 0$ as defined in Lemma 7. Pick the regularization parameters in (9) and (10) such that

$$\begin{aligned}\lambda_n &\geq 12\bar{\psi}\bar{\omega} \left(\sqrt{\frac{2\log p}{n}} + \sqrt{\frac{2\log 2 + 4\log p}{n}} \right), \\ \lambda_p &\geq 40\sqrt{2} \sqrt{\frac{\tau \log n - \log 4}{p}} + \bar{R}(s, p, n),\end{aligned}$$

where $\tau > 2$ and $r(\widehat{\Omega})$ is defined in (12). Let $\bar{\kappa} = \sigma_{\min}(\Psi^*)$. Then for some positive constant c_1 , we have

$$\sup_j \|\widehat{\beta}_j^{(1)} - \beta_j^*\|_2 \leq \frac{\sqrt{s}}{c_1 \bar{\kappa}} \lambda_n, \quad (15)$$

with probability at least $(1 - 2/p)^2 - \{1/(\exp\{n/32\} - 1) + 1/n^{\tau-2} + 5n^2/\max\{n, p\}^4\}$. If in addition n, p satisfy

$$3200 \log(4n^\tau) \max\{160, 24mC\} \leq p, \quad (16)$$

and the following holds,

$$\max \left\{ r(\widehat{\Omega}), \frac{4s\bar{\omega}^3}{b} \sqrt{\frac{\log p \log^2 \max\{n, p\}}{n}} \right\} \leq 1/(24C), \quad (17)$$

where $C = \max\{\kappa_{\Sigma^*} \kappa_{\Gamma^*}, \kappa_{\Sigma^*}^3 \kappa_{\Gamma^*}^2\}$, we also have

$$\|\widehat{\Theta}^{(1)} - \Theta^*\|_2 \leq 4\kappa_{\Gamma^*} m \lambda_p, \quad (18)$$

with the same probability.

We leave the detailed proof for Theorem 4 to the Appendix. Note that $\widehat{\Omega}$ is assumed to be pre-specified in this result. In addition, assume $\bar{\kappa}, \bar{\psi}, \bar{\omega}$ stay bounded as well. Then, under the conditions in Theorem 4, we have for fixed positive constants c_2, c_3, c_4 that

$$\begin{aligned}\sup_j \|\widehat{\beta}_j^{(1)} - \beta_j^*\|_2^2 &\leq c_2 s \frac{\log p}{n}, \\ \|\widehat{\Theta}^{(1)} - \Theta^*\|_2 &\leq c_3 m \left(\sqrt{\frac{\tau \log n}{p}} + \max \left\{ r(\widehat{\Omega}), c_4 s \sqrt{\frac{\log p \log^2 \max\{n, p\}}{n}} \right\} \right),\end{aligned} \quad (19)$$

with high probability.

Remark 5 The error bounds in Theorem 4 hold for any finite samples and for any values of the sparsity measures s (maximum number of parents in the DAG) and m (maximum degree in the undirected graph). If the DAG is dense so that $s = p$, then the corresponding error bounds are given by the right sides of (19) with $s = p$. For example, $\sup_j \|\widehat{\beta}_j^{(1)} - \beta_j^*\| \leq (c_2 p \log p/n)^{1/2}$. This error bound would be small only if $p \ll n$. This is expected when the DAG is not sparse. Similarly, we can obtain the error bounds when m is comparable to n (i.e. the undirected graph is dense).

4.2 Error bound on $\widehat{\Omega}$ under Block-diagonal Θ^*

The error bound of $\widehat{\Theta}$ in (19) depends on $r(\widehat{\Omega})$, the convergence rate of the error variance estimator $\widehat{\Omega}$. In order to achieve the desired rate for the consistent estimator of Θ , we need to make additional sparsity assumptions on Θ . Thus, we restrict our attention mainly to the sparse undirected graphs H consisting of N connected components, which implies that the row precision matrix Θ is block-diagonal:

$$\Theta = \begin{pmatrix} \Theta_1 & & & \\ & \Theta_2 & & \\ & & \ddots & \\ & & & \Theta_N \end{pmatrix}. \quad (20)$$

The support of Θ inside each diagonal block Θ_i could be dense. This type of network is often seen in applications where individuals in the network form clusters: nodes in the same cluster are densely connected, and those from different clusters tend to be more independent from each other. The underlying network H will be sparse if the individuals are from a large number of small clusters. In other words, the sparsity of H depends mainly on the number of diagonal blocks in Θ . More general network structures are also considered in the numerical experiments in Section 5. If we assume that the network of the samples is *block diagonal* and the samples form many small clusters, we would be able to collect independent samples from different clusters. This intuition suggests that existing methods are readily applicable in our setting to obtain consistent estimates of Ω^* , as long as there are enough independent clusters in the undirected network H^* .

Given the block-diagonal structure of Θ in (20), there are a few ways to estimate Ω . We use the *natural estimator* introduced by Yu and Bien (2019). We estimate $\widehat{\omega}_j^2$ using independent samples in X according to the block structure of Θ^* . Let $\Upsilon \subseteq [n]$ be a row index set and X^Υ denote the submatrix formed by selecting rows from a matrix $A_{n \times m}$ with row index $i \in \Upsilon$. We draw one sample from each block and form a smaller $N \times p$ design matrix X^Υ . It is not difficult to see that $X_j^\Upsilon = X^\Upsilon \beta_j^* + \varepsilon_j^\Upsilon$. Next define the natural estimator of ω_j^{*2} for $j \in [p]$ as in Yu and Bien (2019):

$$\widehat{\omega}_j^2 = \min_{\beta_j} \left\{ \frac{1}{N} \|X_j^\Upsilon - X^\Upsilon \beta_j\|_2^2 + 2\lambda_N \|\beta_j\|_1 \right\}, \quad (21)$$

where $\lambda_N > 0$ is a tuning parameter. Alternative methods, such as scaled lasso (Sun and Zhang, 2012) and Stein's estimator (Bayati et al., 2013), can also be used to estimate ω_j^2 .

Formally, suppose that Θ^* has a block-diagonal structure defined in (20). Let N be the number of blocks. For simplicity, we assume that Θ^* consists of N blocks as in (20) hereafter. If we use the *natural estimator* from Yu and Bien (2019) introduced above to get $\widehat{\omega}_j$, then we have the following error bound:

Lemma 6 *Let X be generated from (2) and assume Θ^* is block-diagonal with N blocks. Recall that $s = \sup_j \|\beta_j^*\|_0$. Let $\widehat{\Omega}$ be the natural estimator defined in (21) with*

$$\lambda_N = 12\bar{\psi}\bar{\omega} \left(\sqrt{\frac{2 \log p}{N}} + \sqrt{\frac{2 \log 2 + 4 \log p}{N}} \right),$$

then with probability at least $1 + 1/p^2 - 3/p$ ($p \geq 2$),

$$\sup_{1 \leq j \leq p} \left| \hat{\omega}_j^2 - \omega_j^{*2} \right| \leq \lambda_N s \bar{\beta} + 5 \bar{\omega} \sqrt{\frac{\log 2 + \log p}{N}}.$$

Lemma 6 shows that the maximum error of $\hat{\omega}_j$ is upper bounded by $C \cdot s \sqrt{\log p/N}$ where $C > 0$ is constant and therefore is consistent as long as $s \sqrt{\log p/N} \rightarrow 0$. When p is fixed and N increases, there are more diagonal blocks in Θ^* , indicating more independent samples, and thus $\hat{\omega}_j^2$ will be more accurate. When $N = n$, Θ^* becomes a diagonal matrix and the rows of X become i.i.d. We can easily show that the following lemma holds:

Lemma 7 *Suppose $\inf_{1 \leq j \leq p} w_j^* \asymp 1$ and $\frac{1}{2} \inf_{1 \leq j \leq p} \omega_j^{*2} - \sup_{1 \leq j \leq p} \left| \hat{\omega}_j^2 - \omega_j^{*2} \right| > b > 0$. Then*

$$r(\hat{\Omega}) \leq \frac{1}{b^4} \sup_{1 \leq j \leq p} \left| \hat{\omega}_j^2 - \omega_j^{*2} \right|.$$

When $s(\log p/N)^{1/2}$ is small enough, Lemma 6 implies that b can be taken as a positive constant with high probability.

If $\hat{\Omega}$ satisfies the convergence rate specified in Lemmas 6 and 7, i.e., $r(\hat{\Omega}) \lesssim s \sqrt{\log p/N}$, then the sample constraints in (16) are satisfied as long as

$$m \log n \lesssim p, \quad s^2 \log p \lesssim N, \quad \text{and } s^2 \log^3 \max\{n, p\} \lesssim n. \quad (22)$$

Here, p, s, m all may approach infinity as $n \rightarrow \infty$. As a result, we have the following two asymptotic results based on Theorem 4. The first considers the scaling $p \gg n$ under which DAG estimation is high-dimensional. The second considers the case $n \gg p$ so that the estimation of Θ^* is a high-dimensional problem.

Corollary 8 *Suppose that the sample size and the number of blocks satisfy*

$$p \gg N \log^2 p \gtrsim n \gtrsim N \gg \log p \rightarrow \infty.$$

Assume $\bar{\beta}, \bar{\omega}, \bar{\psi}, \kappa_{\Gamma^}, \kappa_{\Sigma^*} < \infty$ as $n, p \rightarrow \infty$ and $r(\hat{\Omega}) \lesssim s \sqrt{\log p/N} \lesssim 1$. Then under the same assumptions as Theorem 4, we have*

$$\begin{aligned} \sup_j \|\hat{\beta}_j^{(1)} - \beta_j^*\|_2^2 &= O_p \left(s \frac{\log p}{n} \right), \\ \|\hat{\Theta}^{(1)} - \Theta^*\|_2 &= O_p \left(ms \sqrt{\frac{\log^3 p}{n}} \right). \end{aligned}$$

Corollary 9 *Suppose that the sample size and the block numbers satisfy*

$$n \gg s^2 p \log p \log n \gtrsim N \gtrsim s^2 p \rightarrow \infty.$$

Assume $\bar{\beta}, \bar{\omega}, \bar{\psi}, \kappa_{\Gamma^}, \kappa_{\Sigma^*} < \infty$ as $n, p \rightarrow \infty$ and $r(\hat{\Omega}) \lesssim s \sqrt{\log p/N} \lesssim 1$. Then under the same assumptions as Theorem 4, we have*

$$\begin{aligned} \sup_j \|\hat{\beta}_j^{(1)} - \beta_j^*\|_2^2 &= O_p \left(s \frac{\log p}{n} \right), \\ \|\hat{\Theta}^{(1)} - \Theta^*\|_2 &= O_p \left(m \sqrt{\frac{\log n}{p}} \right). \end{aligned}$$

Remark 10 Although we derived the consistency of $\widehat{\Theta}^{(1)}$ and $\widehat{B}^{(1)}$ in the above two corollaries in the setting where Θ^* is block-diagonal, these consistency properties still hold even when Θ^* is not block-diagonal. The main purpose of the block-diagonal setting is to provide an example in which we can conveniently control the error of $\widehat{\Omega}$. But in practice Θ^* does not have to be block-diagonal. In particular, we did not assume any block-diagonal structure of Θ^* for the error bounds in Theorem 4. It can be seen that the error bound of $\widehat{B}^{(1)}$ does not depend on the error of $\widehat{\Omega}$ at all. Hence, the accuracy of $\widehat{\Omega}$ has no impact on the accuracy of $\widehat{B}^{(1)}$. The error bound of $\widehat{\Theta}^{(1)}$ in (19) is determined by the trade-off among three terms, one of which is the error rate $r(\widehat{\Omega})$ as in (12). This is also supported by our numerical results. In Section 5, we demonstrate, with both simulated and real networks where Θ^* is not block-diagonal, that our proposed BCD method can still accurately estimate Θ^* and B^* whenever a relatively accurate $\widehat{\Omega}$ is provided.

4.3 Comparison to Other Results

If the data matrix X consists of i.i.d. samples generated from Gaussian linear SEM (1), assuming the topological sort of the vertices is known, the DAG estimation problem in (8) is reduced to solving the standard Lasso regression in (9) with $\widehat{L}^{(t)} = I_n$, and thus independent of the initial $\widehat{\Theta}^{(0)}$ estimator. Under the *restricted eigenvalue condition* and letting $\lambda_n \asymp \sqrt{\log p/n}$, it is known the Lasso estimator has the following optimal rate for ℓ_2 error (van de Geer and Bühlmann, 2009):

$$\sup_j \|\widehat{\beta}_j - \beta_j^*\|_2^2 = O_p \left(s \frac{\log p}{n} \right).$$

When the data are dependent, Theorem 4 shows that the estimator from Algorithm 1 can achieve the same optimal rate if we make the extra assumptions above. In particular, what we need is a reasonably good initial $\widehat{\Theta}^{(0)}$ estimate such that $\|\widehat{\Theta}^{(0)} - \Theta^*\|_2 \leq M$ for a small positive constant M .

On the other hand, if the underlying DAG is an empty graph and $\Omega^* = I_p$, the problem of estimating Θ^* can be solved using graphical Lasso in (10) because the data (columns in X) are i.i.d. The sample variance \widehat{S} would also be an unbiased estimator of Σ^* . In this case, Ravikumar et al. (2011) showed that

$$\|\widehat{\Theta} - \Theta^*\|_2 = O_p \left(m \sqrt{\frac{\log n}{p}} \right).$$

This result does not require knowing $\text{supp}(\Theta^*)$ but assumes a *mutual incoherence* condition on the Hessian of the log likelihood function. In our case, $\widehat{S}^{(1)}$ is biased due to the accumulated errors from the previous Lasso estimate, as well as $\widehat{\Omega}$. As a result, there is an extra bias term $\bar{R}(s, n, p)$ in $\|\widehat{S}^{(1)} - \Sigma^*\|_\infty$ (see Lemma 22 in Appendix A.3) compared to the i.i.d. setting:

$$\|\widehat{S}^{(1)} - \Sigma^*\|_\infty = \bar{\delta}_f(p, n^\tau) + \bar{R}(s, n, p),$$

where $\bar{\delta}_f(p, n^\tau) \asymp \sqrt{\log n/p}$ is the classical graphical Lasso error rate, and

$$\bar{R}(n, p, s) \asymp \max \left\{ r(\widehat{\Omega}), s \sqrt{\log p \log \max\{n, p\}/n} \right\}$$

depends on the estimation errors of $\widehat{B}^{(1)}$ and $\widehat{\Omega}$. When $n \gg p$ and $r(\widehat{\Omega})$ is dominated by $\sqrt{\log(n)/p}$, we get the same rate for the ℓ_2 consistency of $\widehat{\Theta}$ (Corollary 9) under a slightly more strict constraint on sample size (22). If $n \ll p$, then the ℓ_2 error rate is determined by $\max\{r(\widehat{\Omega}), s(\log^3 p/n)^{1/2}\}$. Suppose Θ^* is block-diagonal. If the number of blocks N is much smaller than n , then the ℓ_2 rate will be dominated by $r(\widehat{\Omega}) \asymp s\sqrt{\log p/N}$, which could be slower than the optimal graphical Lasso rate. But that is expected due to the error introduced in the DAG estimates $\widehat{B}^{(1)}$ and $\widehat{\Omega}$.

5. Numerical Experiments

Under the assumption that observations generated from a DAG model are dependent, we will evaluate the performance of the block coordinate descent (BCD) algorithm, i.e., Algorithm 1, in recovering the DAG compared to traditional methods that treat data as independent. We expect that the BCD method would give more accurate structural estimation than the baselines by taking the dependence information into account. When a topological ordering of the true DAG is known, we can identify a DAG from data using BCD. When ordering is unknown, the BCD algorithm may still give an accurate estimate of row correlations that are invariant to node-wise permutations according to Lemma 1. The estimated row correlation matrix can then be used to de-correlate the data so that traditional DAG learning algorithms would be applicable. We will demonstrate this idea of de-correlation with numerical results as well.

5.1 Simulated Networks

We first perform experiments on simulated networks for both ordered and unordered cases. To apply the BCD algorithm, we need to set values for λ_1 and λ_2 in (8). Since the support of Θ^* is restricted to the edge set of H^* , we simply fixed λ_2 to a small value ($\lambda_2 = 0.01$) in all experiments. For each data set, we computed a solution path from the largest $\lambda_{1\max}$, for which we get an empty DAG, to $\lambda_{1\min} = \lambda_{1\max}/100$. The optimal λ_1 was then chosen by minimizing the BIC score over the DAGs on the solution path.

We generate random DAGs with p nodes and fix the total number of edges s_0 in each DAG to $2p$. The entries in the weighted adjacency matrix B^* of each DAG were drawn uniformly from $[-1, -0.1] \cup [0.1, 1]$, and ω_j^* 's were sampled uniformly from $[0.1, 2]$. In our simulations of Θ^* , we first considered networks with a clustering structure, i.e., Θ^* was block-diagonal as in (20). We fixed the size of the clusters to 20 or 30, and within each cluster, the individuals were correlated according to the following four covariance structures.

- Toeplitz: $\Sigma_{ij}^* = 0.3^{|i-j|/5}$.
- Equal correlation: $\Sigma_{ij}^* = 0.7$ if $i \neq j$, and $\Sigma_{ii}^* = 1$.
- Star-shaped: $\Theta_{1j}^* = \Theta_{i1}^* = a$, $i, j \geq 2$, $a \in (0, 1)$, and $\Theta_{ii}^* = 1$.
- Autoregressive (AR): $\Theta_{ij}^* = 0.7^{|i-j|}$ if $|i - j| \leq \lceil b/4 \rceil$; $\Theta_{ij}^* = 0$ otherwise, where b is the cluster size.

Toeplitz covariance structure implies that the observations are correlated as in a Markov chain. The equal correlation structure represents the cases when all observations are fully

connected in a cluster. Star-shaped and AR structures capture intermediate dependence levels. In addition to these block-diagonal covariances, we also considered a more general covariance structure defined through *stochastic block models* (SBM), in which H^* consists of several clusters and nodes within a cluster have a higher probability of being connected than those from different clusters. More explicitly, we generated Θ^* as follows:

1. Let $\mathcal{B}_1, \dots, \mathcal{B}_L$ be L clusters with varying sizes that form a partition of $\{1, \dots, n\}$, where the number of clusters L ranges from 5 to 10 in our experiments. Define a probability matrix $P \in \mathbb{R}^{n \times n}$ where $P_{ij} = 0.5$ if $i, j \in \mathcal{B}_l, l \in \{1, \dots, L\}$; otherwise, $P_{ij} = 0.1$.
2. Construct the adjacency matrix A of H^* :

$$A_{ij} \sim \text{Bern}(P_{ij}).$$

3. Sample $\Theta'_{ij} \sim \text{Unif}[-5, 5]$ if $A_{ij} = 1$. Otherwise, $\Theta'_{ij} = 0$. To ensure a positive-definite Θ^* , we then perform the following transformations to get Θ^* :

$$\begin{aligned} \tilde{\Theta} &= (\Theta' + \Theta'^\top)/2, \\ \Theta^* &= \tilde{\Theta} - \left(\sigma_{\min}(\tilde{\Theta}) - 0.01 \right) \cdot I_n. \end{aligned} \tag{23}$$

In the stochastic block model, two nodes from different clusters in H^* are connected with probability 0.1, so Θ^* is not block-diagonal in general. As explained in Section 4.1, our proposed BCD algorithm does not require Θ^* to be block-diagonal in practice to produce accurate estimates of B^* and Θ^* . Our numerical experiments will support this point and demonstrate the robustness of the BCD method.

We compared the BCD algorithm with its competitors in both high-dimensional ($p > n$) and low-dimensional ($p < n$) settings with respect to DAG learning. For each (n, p) and each type of covariances, we simulated 10 random DAGs and then generated one data set following equation (2) for each DAG. Thus, we had 10 results for each of the $2 \times 5 = 10$ simulation settings. In the end, we averaged the results over the 10 simulations under each setting for comparison.

5.1.1 LEARNING WITH GIVEN ORDERING

Assuming that the nodes in a DAG are sorted according to a given topological order, we compared our BCD algorithm against a baseline setting that fixes $\Theta^* = I_n$. In other words, the baseline algorithm ignores the dependencies between observations when estimating the DAG with BCD. The block sizes in Θ^* were set to 20 in all cases except SBM whose block sizes ranged from 5 to 25. Among other estimates, both algorithms return an estimated weighted adjacency matrix \hat{B} for the optimal λ_1 selected by BIC. For the BCD algorithm, we use \hat{B} and $\hat{\Theta}$ for $\hat{B}^{(\infty)}$ and $\hat{\Theta}^{(\infty)}$ after convergence (see Algorithm 1). Note that since $\hat{\Theta}^{(0)}$ is initialized to I_n by default in the BCD algorithm, the estimated \hat{B} from the baseline algorithm is the same as the estimate $\hat{B}^{(1)}$ from BCD after one iteration.

We also included three other DAG learning algorithms in our comparison. The first is the Kronecker graphical Lasso (KGLasso) algorithm (Allen and Tibshirani, 2010; Tsiligkaridis

et al., 2013) mentioned in Section 2, which estimates both $\widehat{\Psi}$ and $\widehat{\Theta}$ via graphical Lasso in alternating fashion. When estimating Θ^* , KGLasso also makes use of its block-diagonal structure. After KGLasso converges, we perform the Cholesky factorization on $\widehat{\Psi} = (I - \widehat{B})^{-\top} \widehat{\Omega} (I - \widehat{B})^{-1}$ according to the given order to obtain \widehat{B} and $\widehat{\Omega}$. A distinction between BCD and KGLasso is that KGLasso imposes a sparsity regularization on Ψ instead of B , so the comparison between these two will highlight the importance of imposing sparsity directly on the Cholesky factor. The other two algorithms are PC (Spirtes et al., 2000) and Greedy Equivalence Search (Chickering et al., 2004), both of which were applied with the node ordering given as prior knowledge. Note that GES in this case is similar to a greedy hill-climbing algorithm, since restricting to a given ordering, an equivalence class reduces to a single DAG.

Given the estimate \widehat{B} from a method, we hard-thresholded the entries in \widehat{B} at a threshold value $\bar{\tau}$ to obtain an estimated DAG. To compare the methods, we chose $\bar{\tau}$ such that they predicted roughly the same number of edges (E). Then we calculated the number of true positives (TP), false positives (FP) and false negatives (FN, missing edges), and two overall accuracy metrics: Jaccard index ($TP / (TP + s_0)$) and structural Hamming distances ($SHD = FP + FN$). Note that, there were no reserved edges (i.e., estimated edges whose orientation is incorrect) because the ordering of the nodes was given. Detailed comparisons are summarized in Table 1 and Table 2. Our DAG model (2) defines a matrix normal distribution in eq (3), for which the row-wise inverse covariance matrix is $\Theta^* = (\Sigma^*)^{-1}$. The matrix normal matches exactly the model assumption for KGLasso. So it is expected to be accurate in estimating Θ^* and Ψ^* . However, KGLasso showed lower accuracy (JI and SHD) in estimating the DAG (B^*), because it does not impose sparsity on the DAG coefficients B . In general, the BCD algorithm outperformed the competitors by having more true positives and fewer false positives in all cases. Because the KGLasso method does not impose sparsity directly on the DAG structure, it suffered from having too many false negatives after thresholding when $p > n$. PC and GES had better performance than KGLasso in terms of FDR, SHD and Jaccard index. When $p < n$, the correlations between observations had a more significant impact on the estimation accuracy for DAGs. As a result, by taking the correlations into account, BCD performed better than the baseline, PC, and GES. In particular, BCD substantially reduced the number of missing edges (FNs) and FDR, compared to baseline. Both BCD and KGLasso yielded accurate estimates of $\widehat{\Theta}$ when $n < p$ except for the SBM. When $n > p$, as the sample size p for estimating $\Theta^* \in \mathbb{R}^{n \times n}$ decreased relative to dimension n , $\widehat{\Theta}$ became less accurate. The difference in the accuracy of $\widehat{\Theta} = \widehat{\Theta}^{(\infty)}$ and $\widehat{\Theta}^{(1)}$ was not significant.

Figure 2 shows the ROC curves of all the methods over a sequence of $\bar{\tau}$ under the 10 settings. The $\bar{\tau}$ sequence contains 30 equally spaced values in $[0, 0.5]$. The BCD algorithm uniformly outperformed the others in terms of the area under the curve (AUC) with substantial margins when $n < p$. When $n > p$, the BCD still performed better than the other four most of the time, but its lead over KGLasso was not as significant in some cases. This was largely due to insufficient regularization on $\widehat{\Theta}$. Fixing $\lambda_2 = 0.01$ in this case implies $\lambda_p = 0.01/p = 0.0001$ in the graphical Lasso step (10) of BCD, resulting in severe overestimates of the magnitude of the entries in Θ^* . After we increased λ_p to 0.1 which is still quite small, the BCD indeed outperformed the other methods by much larger margins. KGLasso also performed much better when $n > p$ as shown in Table 2 and Figure 2. This

Θ -Network (n, p, s_0)	Method	E	FN	TP	FDR	JI	SHD	err($\hat{\Theta}$) (err($\hat{\Theta}^{(1)}$))
(150, 300, 600)	BCD	686.2 (98.9)	214.0 (19.3)	386.0 (19.3)	0.355 (0.065)	0.443 (0.021)	514.2 (65.8)	0.00034 (0.00032)
	Baseline	642.4 (79.0)	240.3 (20.5)	359.7 (20.5)	0.383 (0.052)	0.410 (0.01)	523.0 (40.5)	—
	KGLasso	756.2 (146.7)	504.9 (6.3)	95.1 (6.3)	0.822 (0.035)	0.080 (0.006)	1166.0 (135.8)	0.00019
	PC	333.1 (6.3)	426.0 (3.6)	174.0 (3.6)	0.476 (0.012)	0.229 (0.005)	585.1 (7.7)	—
	GES	734.8 (80.9)	232.4 (18.8)	367.6 (18.8)	0.472 (0.028)	0.383 (0.007)	599.6 (45.7)	—
(200, 400, 800)	BCD	535.0 (6.2)	306.4 (4.7)	493.6 (4.7)	0.077 (0.004)	0.586 (0.005)	347.8 (4.3)	0.00143 (0.00833)
	Baseline	549.5 (9.0)	425.2 (4.3)	374.8 (4.3)	0.317 (0.008)	0.384 (0.004)	599.9 (6.2)	—
	KGLasso	550.6 (11.2)	634.3 (3.3)	165.7 (3.3)	0.698 (0.004)	0.139 (0.002)	1019.2 (7.8)	0.01617
	PC	536.2 (6.7)	554.4 (3.1)	245.6 (3.1)	0.542 (0.006)	0.225 (0.003)	845.0 (6.6)	—
	GES	534.0 (5.5)	496.8 (3.1)	303.2 (3.1)	0.432 (0.005)	0.294 (0.003)	727.6 (4.9)	—
(200, 400, 800)	BCD	543.1 (3.7)	301.1 (2.4)	498.9 (2.4)	0.081 (0.003)	0.591 (0.003)	345.3 (2.7)	0.00006 (0.00051)
	Baseline	515.7 (6.7)	338.3 (2.1)	461.7 (2.1)	0.103 (0.01)	0.541 (0.003)	392.3 (5.2)	—
	KGLasso	495.8 (13.9)	504.3 (7.5)	295.7 (7.5)	0.403 (0.006)	0.295 (0.006)	704.4 (5.7)	0.00233
	PC	544.1 (4.0)	375.7 (2.3)	424.3 (2.3)	0.220 (0.005)	0.461 (0.003)	495.5 (3.9)	—
	GES	543.6 (3.5)	329.7 (2.8)	470.3 (2.8)	0.135 (0.005)	0.539 (0.004)	403.0 (4.5)	—
(100, 200, 400)	BCD	253.1 (5.7)	193.8 (2.9)	206.2 (2.9)	0.184 (0.008)	0.461 (0.004)	240.7 (1.5)	0.00247 (0.00219)
	Baseline	245.8 (8.5)	208.9 (2.6)	191.1 (2.6)	0.217 (0.018)	0.420 (0.004)	263.6 (4.6)	—
	KGLasso	270.4 (9.3)	271.1 (3.1)	128.9 (3.1)	0.521 (0.007)	0.237 (0.004)	412.6 (4.1)	0.02587
	PC	253.5 (6.1)	237.9 (1.8)	162.1 (1.8)	0.358 (0.015)	0.330 (0.005)	329.3 (6.1)	—
	GES	314.6 (2.8)	218.2 (2.2)	181.8 (2.2)	0.422 (0.007)	0.341 (0.005)	351.0 (4.2)	—
(100, 300, 600)	BCD	343.3 (16.9)	313.6 (8.5)	286.4 (8.5)	0.159 (0.016)	0.435 (0.008)	370.5 (10.2)	0.51266 (0.52297)
	Baseline	344.6 (16.4)	338.4 (6.5)	261.6 (6.5)	0.233 (0.02)	0.383 (0.008)	421.4 (10.2)	—
	KGLasso	301.3 (18.9)	510.7 (3.0)	89.3 (3.0)	0.696 (0.016)	0.110 (0.003)	722.7 (16.5)	0.46201
	PC	302.7 (4.1)	438.3 (1.9)	161.7 (1.9)	0.465 (0.008)	0.218 (0.003)	579.3 (4.9)	—
	GES	411.3 (5.8)	335.2 (3.0)	264.8 (3.0)	0.356 (0.006)	0.355 (0.004)	481.7 (4.5)	—

Table 1: Results for ordered DAGs on simulated data when $n < p$. E is the number of estimated edges in the DAG. FN is the number of false negatives in the estimated DAG. TP is the number of true positives in the estimated DAG. $FDR = FP/(FP + TP)$ is the false discovery rate. $JI = TP/(FP + s_0)$ is the Jaccard Index. $SHD = FP + FN$ is the structural hamming distance. We highlight in boldface the best performance in terms of FDR, JI, and SHD. The last column shows the ℓ_2 -estimation errors of $\hat{\Theta}$ and $\hat{\Theta}^{(1)}$ normalized by the true support size. The numbers in the brackets are errors after one iteration of BCD. Each number corresponds to the average (standard error) over 10 simulations.

Θ -Network (n, p, s_0)	Method	E	FN	TP	FDR	JI	SHD	err($\hat{\Theta}$) (err($\hat{\Theta}^{(1)}$))
equi-cor (200, 100, 200)	BCD	143.0 (9.4)	75.1 (6.4)	124.9 (6.4)	0.119 (0.014)	0.570 (0.022)	93.2 (3.7)	0.00211 (0.00199)
	Baseline	143.6 (9.5)	102.8 (4.2)	97.2 (4.2)	0.311 (0.024)	0.393 (0.010)	149.2 (3.0)	—
	KGLasso	142.7 (9.3)	129.5 (1.8)	70.5 (1.8)	0.490 (0.030)	0.260 (0.007)	201.7 (7.6)	0.00207
	PC	141.8 (8.7)	135.9 (2.5)	64.1 (2.5)	0.539 (0.019)	0.231 (0.007)	213.6 (6.1)	—
	GES	191.2 (3.4)	103.3 (1.3)	96.7 (1.3)	0.493 (0.008)	0.329 (0.005)	197.8 (3.4)	—
toeplitz (200, 100, 200)	BCD	167.1 (7.3)	56.7 (4.0)	143.3 (4.0)	0.137 (0.015)	0.639 (0.010)	80.5 (1.4)	0.51735 (0.68703)
	Baseline	167.3 (7.2)	104.8 (2.1)	95.2 (2.1)	0.426 (0.015)	0.350 (0.005)	176.9 (4.0)	—
	KGLasso	167.1 (7.2)	68.8 (2.2)	131.2 (2.2)	0.206 (0.023)	0.557 (0.007)	104.7 (3.7)	0.014
	PC	162.6 (5.8)	130.0 (1.4)	70.0 (1.4)	0.566 (0.013)	0.240 (0.005)	222.6 (4.9)	—
	GES	235.2 (4.8)	109.7 (1.3)	90.3 (1.3)	0.615 (0.008)	0.262 (0.005)	254.6 (4.8)	—
star (200, 100, 200)	BCD	186.9 (7.3)	50.6 (4.0)	149.4 (4.0)	0.199 (0.015)	0.629 (0.010)	88.1 (1.4)	0.54769 (0.39316)
	Baseline	184.3 (7.2)	64.4 (2.1)	135.6 (2.1)	0.263 (0.015)	0.546 (0.005)	113.1 (4.0)	—
	KGLasso	186.1 (7.2)	66.1 (2.2)	133.9 (2.2)	0.279 (0.023)	0.531 (0.007)	118.3 (3.7)	0.18472
	PC	170.0 (2.5)	91.5 (1.3)	108.5 (1.3)	0.361 (0.009)	0.415 (0.006)	153.0 (2.5)	—
	GES	187.4 (3.9)	68.9 (2.2)	131.1 (2.2)	0.300 (0.005)	0.511 (0.006)	125.2 (1.8)	—
AR(5) (200, 100, 200)	BCD	185.9 (6.2)	52.2 (1.5)	147.8 (1.5)	0.200 (0.018)	0.622 (0.009)	90.3 (4.0)	0.01644 (0.01184)
	Baseline	187.7 (6.9)	63.8 (1.2)	136.2 (1.2)	0.268 (0.021)	0.544 (0.011)	115.3 (5.8)	—
	KGLasso	186.0 (6.7)	59.8 (1.5)	140.2 (1.5)	0.240 (0.02)	0.572 (0.01)	105.6 (4.7)	0.01038
	PC	175.0 (1.8)	90.7 (1.7)	109.3 (1.7)	0.375 (0.009)	0.412 (0.008)	156.4 (3.1)	—
	GES	185.9 (6.6)	78.9 (2.3)	121.1 (2.3)	0.345 (0.011)	0.457 (0.005)	143.7 (3.0)	—
SBM (300, 100, 200)	BCD	140.1 (3.5)	72.2 (2.4)	127.8 (2.4)	0.086 (0.008)	0.602 (0.009)	84.5 (1.7)	0.31189 (0.31448)
	Baseline	139.9 (4.0)	85.7 (2.4)	114.3 (2.4)	0.181 (0.01)	0.507 (0.009)	111.3 (2.1)	—
	KGLasso	128.8 (11.0)	161.1 (3.5)	38.9 (3.5)	0.689 (0.023)	0.134 (0.011)	251.0 (8.5)	0.32263
	PC	139.5 (3.3)	113.8 (1.6)	86.2 (1.6)	0.379 (0.017)	0.341 (0.009)	167.1 (4.4)	—
	GES	146.4 (2.5)	90.6 (1.95)	109.4 (1.95)	0.252 (0.011)	0.462 (0.009)	127.6 (2.9)	—

Table 2: Results for ordered DAGs on simulated data when $n > p$.

is expected because when n is large compared to p , the dependence between individuals will have a greater impact on the accuracy of the estimation of DAGs. Since KGLasso is designed to iteratively estimate Θ^* and Ψ^* , the more accurate estimates of Θ^* as reported in Table 2 compensated for the relatively inaccurate \hat{B} .

We also compared test data log-likelihood among the five methods. Specifically, in each setting, we generated a test sample matrix X_{test} from the true distribution for each of the 10 repeated simulations and computed $-L(\hat{B}, \hat{\Theta}, \hat{\Omega} \mid X_{\text{test}})$ using the estimates from a method following equation (5). Figure 3 shows the boxplots of the test data log-likelihood, normalized by \sqrt{np} after subtracting the median of the baseline method: $\ell_{\text{plot}} = (\ell_0 - \text{median}(\ell_0^{\text{baseline}})) / \sqrt{np}$, where ℓ_0 is the original test data log-likelihood. The top row shows the test log-likelihood when $n < p$, where we did not include the data for KGLasso in four cases because its test data log-likelihood values were too small to fit in the same plot. The bottom row shows the results for $n > p$. For both cases, we see that the test data log-likelihood of the BCD method (in green) is consistently higher than that of the other methods.

5.1.2 LEARNING WITH DE-CORRELATION

When true ordering is unknown, we focus on estimating the row-wise covariance Σ^* . Given $\hat{\Sigma}$ we can de-correlate the data by Equation (11) and apply existing structural learning methods. In this study, we compared the performances of three structure learning methods before and after de-correlation: GES (Chickering, 2003) and sparsebn (Aragam et al., 2019b)

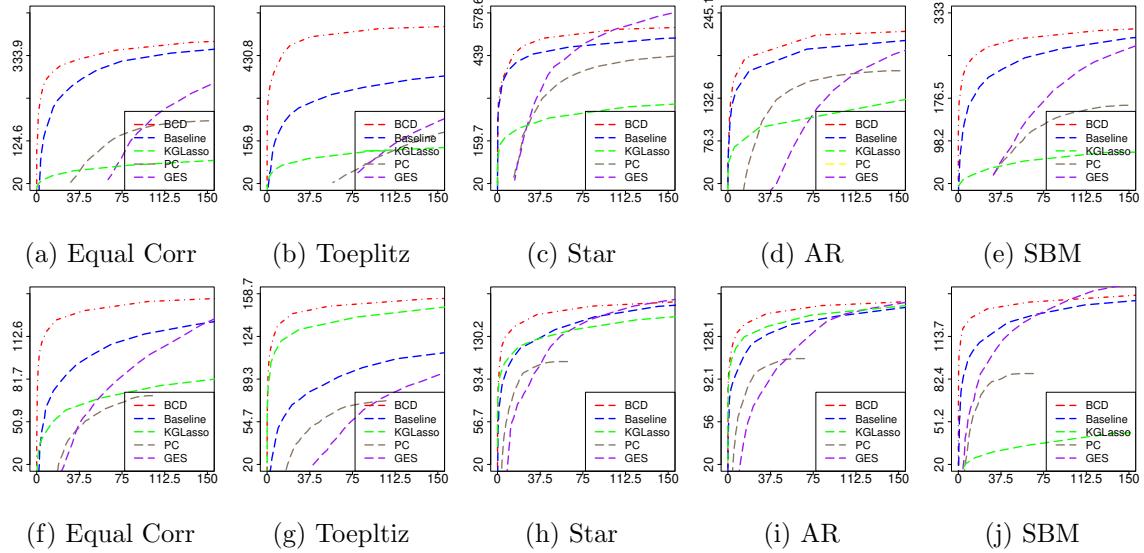


Figure 2: ROC curves of BCD, baseline, and KGLasso on simulated and sorted DAGs: x-axis reports the number of false positive edges and y-axis true positive edges. Top row: $n < p$. Bottom row: $n > p$. Each data point in the ROC curves corresponds to the average over 10 simulations.

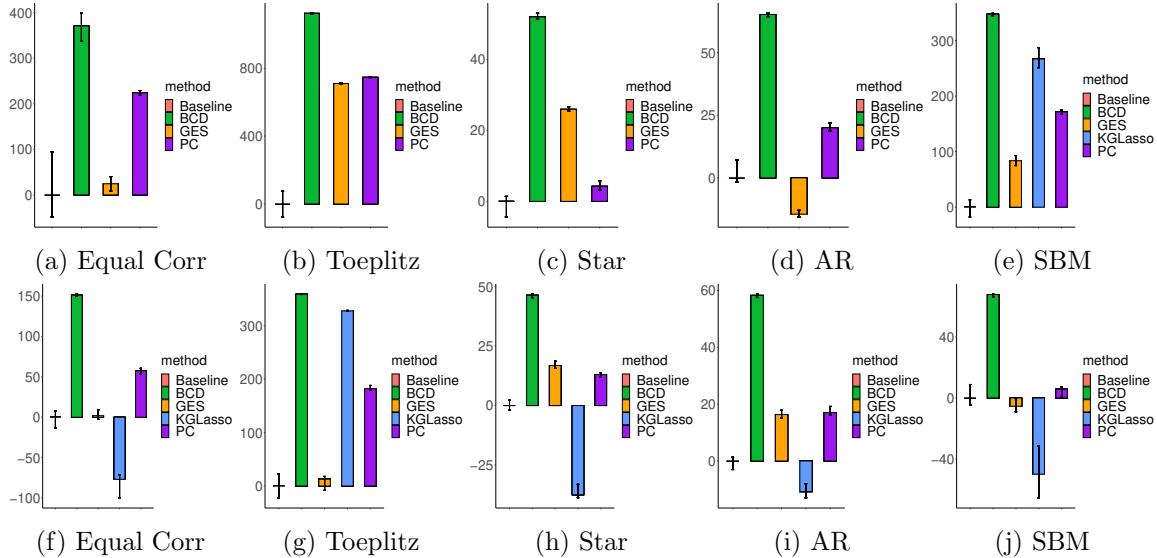


Figure 3: Normalized test data log-likelihood of BCD and alternative methods on simulated sorted DAGs. Top row: $n < p$. Bottom row: $n > p$. Bar height represents the median log-likelihood from the 10 repeated experiments with error bars for the 25% and 75% quantiles.

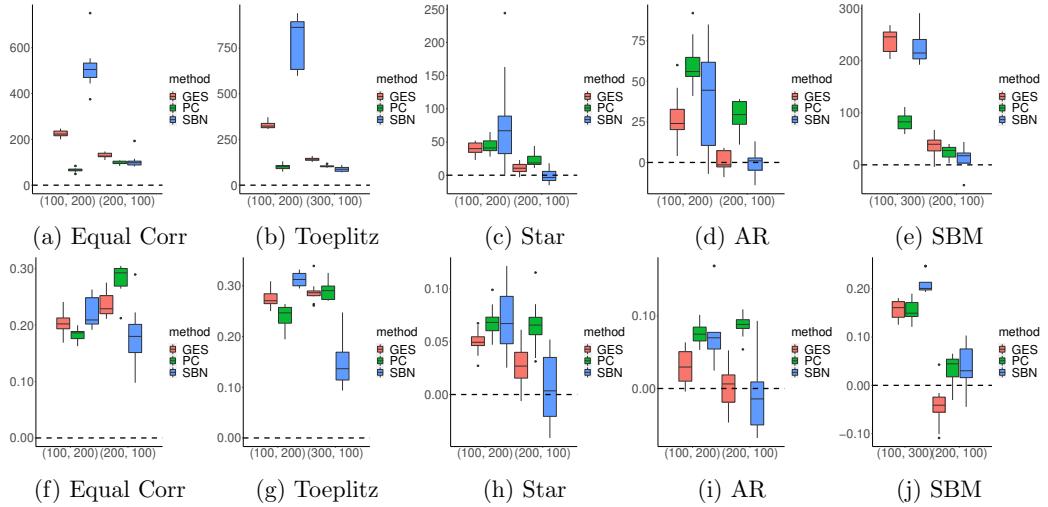


Figure 4: Decrease in SHD (top row) and increase in Jaccard index (bottom row) via de-correlation on simulated unsorted DAGs, with x-axis reporting the value of (n, p) . In each panel, the three boxplots on the left and the three on the right correspond to the cases of $n < p$ and $n > p$, respectively. Each boxplot contains 10 data points from 10 simulations.

which are score-based methods implemented respectively in the R packages `rcausal` (Ramsey et al., 2017) and `sparsebn`, and PC (Spirtes et al., 2000) which is a constraint-based method implemented in `pcaig` (Kalisch et al., 2012). All three methods rely on the independent data assumption, so we expect the de-correlation step to improve their performances significantly. Unlike the previous comparison, the ordering of the nodes is unknown, so GES and PC return an estimated CPDAG (completed acyclic partially directed graph) instead of a DAG. Thus, in the following comparisons, we converted both the estimated DAG from `sparsebn` and the true DAG to CPDAGs, so that all the reported metrics are computed with respect to CPDAGs.

As before, we divide the cases into $n < p$ and $n > p$. The block size for the four block-diagonal Θ was fixed to 30. The estimated Cholesky factor \widehat{L} of $\widehat{\Theta}$ used for de-correlating X in (11) was calculated by our BCD algorithm with tuning parameter λ_1 selected by BIC. Figure 4 shows the decrease in SHD and increase in the Jaccard index via the de-correlation of GES, PC and `sparsebn` on 10 random DAGs, generated under each row-covariance structure and each sample size. For almost all types of covariances and (n, p) settings we considered, there is a significant improvement of all three methods in estimating the CPDAG structures after de-correlation. Additional tables with detailed results can be found in the Supplementary Material. Before de-correlation, GES and `sparsebn`, both score-based methods, tend to significantly overestimate the number of edges, resulting in high false positives, so does PC in some of the cases. After de-correlation, both GES and `sparsebn` had significant improvements and outperformed PC, as long as $\widehat{\Theta}$ was accurately estimated. The test data log-likelihood (normalized by \sqrt{np}) of all three algorithms also increased significantly after de-correlation as shown in Figure 5.

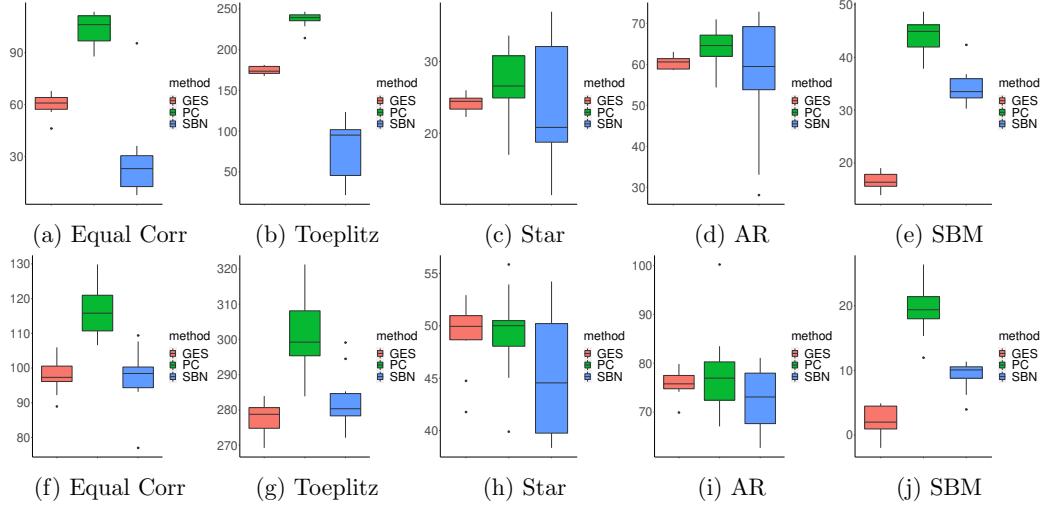


Figure 5: Increase in the normalized test data log-likelihood after decorrelation on simulated unsorted DAGs. Top row: $n < p$. Bottom row: $n > p$.

5.2 Experiments on Real Network Structures

In this section, we examine the performance of the BCD algorithm on real network structures. We took four real DAGs from the bnlearn repository (Scutari, 2010): `Andes`, `Hailfinder`, `Barley`, `Hepar2`, and two real undirected networks from `tnet` (Opsahl, 2009): `facebook` (Opsahl and Panzarasa, 2009) and `celegans_n306` (Watts and Strogatz, 1998). Only the structures (supports) of these real networks were used, and the parameters of the edges were simulated as follows. Given a DAG structure, we sampled the coefficients β_j^* uniformly from $[-1, -0.1] \cup [0.1, 1]$. Given the support of Θ^* , we generate Θ'_{ij} uniformly from $[-5, 5]$. Then, we apply the transformations in (23) to get Θ^* . In order to increase the size of the underlying DAG and show the scalability of the algorithm, we duplicated the DAGs above to form larger networks. In Sections 5.2.1 and 5.2.2, we again consider undirected networks consisting of several disconnected subgraphs, corresponding to a block-diagonal structure in Θ^* . Each of the subgraphs was sub-sampled from the original real network. In Section 5.2.3 we present experiments on more general Θ^* without a block-diagonal structure. The ω_j^* were uniformly sampled from $[0.1, 2]$ as before. With these parameters, we generated observational samples X following the structural equation (2).

5.2.1 LEARNING WITH GIVEN ORDERING

Similar to the previous section, we first look at the results on ordered DAGs. We considered four different combinations of network structures, as shown in Table 3, with both $n > p$ and $n < p$. Because KGLasso does not make use of the given ordering in its estimation of the column-wise covariance matrix, it became very slow when p is large. Therefore, we did not include it in our comparisons. GES performed significantly worse than the others in some cases, so we removed it from those plots. Note that when n is much smaller than p , as in the case of (`Andes`, `facebook`) in Table 3, the BIC score that GES maximizes over may not be a good scoring function, resulting in much worse estimation. BCD continued to

outperform the baseline method by modeling the sample correlation. This improvement was more prominent when $n > p$, where BCD significantly reduced the number of false positive edges, achieving higher JI and lower SHD compared to the baseline and to itself in the $n < p$ case. Figure 6 compares the test data log-likelihood across 10 simulations, and BCD scored significantly higher test data log-likelihood in all the cases. The ROC curves are provided in a figure in the Supplementary Material. Both figures indicate that the BCD method indeed gives better DAG estimates than the baseline and the other standard methods.

DAG	Θ -Network (n, p, s_0)	Method	E	FN	TP	FDR	JI	SHD
Andes (2)	facebook (100, 446, 676)	BCD	440.9 (15.9)	314.7 (7.8)	361.3 (7.8)	0.176 (0.014)	0.478 (0.006)	394.3 (3.1)
		Baseline	436.6 (15.2)	334.2 (4.8)	341.8 (4.8)	0.211 (0.022)	0.444 (0.006)	429.0 (9.8)
		PC	441.8 (16.2)	385.2 (6.5)	290.8 (6.5)	0.338 (0.012)	0.351 (0.005)	536.2 (6.6)
		GES	895.7 (10.4)	504.4 (2.5)	171.6 (2.5)	0.808 (0.002)	0.123 (0.001)	1228.5 (7.2)
	facebook (500, 446, 676)	BCD	500.0 (5.9)	197.3 (5.5)	478.7 (5.5)	0.043 (0.002)	0.686 (0.008)	218.6 (5.3)
		Baseline	499.1 (5.7)	206.4 (1.9)	469.6 (1.9)	0.058 (0.008)	0.666 (0.003)	235.9 (3.1)
		PC	499.2 (6.6)	252.2 (3.3)	423.8 (3.3)	0.150 (0.006)	0.564 (0.003)	327.6 (3.1)
		GES	507.6 (2.9)	295.8 (1.4)	380.2 (1.4)	0.251 (0.003)	0.473 (0.002)	423.2 (2.5)
	celegan_n306 (100, 224, 264)	BCD	178.5 (4.5)	113.9 (2.5)	150.1 (2.5)	0.157 (0.009)	0.513 (0.006)	142.3 (1.7)
		Baseline	179.1 (4.8)	117.8 (1.8)	146.2 (1.8)	0.180 (0.016)	0.493 (0.006)	150.7 (3.1)
		PC	178.5 (5.0)	154.6 (2.4)	109.4 (2.4)	0.386 (0.005)	0.328 (0.005)	223.7 (1.2)
		GES	435.4 (8.6)	128.1 (1.7)	135.9 (1.7)	0.687 (0.007)	0.242 (0.005)	427.6 (8.9)
	Hailfinder (4)	BCD	198.6 (1.3)	72.7 (1.0)	191.3 (1.0)	0.037 (0.002)	0.705 (0.003)	80.0 (0.98)
		Baseline	199.0 (1.5)	74.7 (1.2)	189.3 (1.2)	0.049 (0.005)	0.692 (0.005)	84.4 (1.6)
		PC	198.6 (1.3)	95.6 (1.2)	168.4 (1.2)	0.152 (0.006)	0.573 (0.006)	125.8 (2.3)
		GES	198.2 (1.5)	93.8 (1.4)	170.2 (1.4)	0.141 (0.006)	0.583 (0.006)	121.8 (2.3)
	facebook (100, 192, 336)	BCD	261.7 (6.7)	133.8 (2.7)	202.2 (2.7)	0.225 (0.012)	0.511 (0.005)	193.3 (3.4)
		Baseline	261.8 (8.3)	140.3 (2.6)	195.7 (2.6)	0.247 (0.019)	0.488 (0.009)	206.4 (6.9)
		PC	219.3 (2.4)	180.0 (1.7)	156.0 (1.7)	0.288 (0.006)	0.391 (0.004)	243.3 (2.5)
		GES	314.5 (4.9)	159.3 (1.5)	176.7 (1.5)	0.437 (0.008)	0.373 (0.004)	297.1 (4.7)
	Barley (4)	BCD	260.5 (4.1)	91.3 (3.9)	244.7 (3.9)	0.061 (0.003)	0.696 (0.011)	107.1 (3.9)
		Baseline	260.2 (4.7)	97.6 (2.5)	238.4 (3.9)	0.083 (0.008)	0.666 (0.004)	119.4 (1.6)
		PC	259.9 (4.2)	135.6 (1.0)	200.4 (1.0)	0.227 (0.013)	0.507 (0.007)	195.1 (4.9)
		GES	260.5 (4.1)	104.2 (3.2)	231.8 (3.2)	0.110 (0.005)	0.636 (0.008)	132.9 (2.9)
	celegan_n306 (100, 280, 492)	BCD	366.7 (13.7)	234.5 (6.2)	257.5 (6.2)	0.295 (0.010)	0.428 (0.006)	343.7 (3.2)
		Baseline	373.9 (14.1)	238.0 (2.7)	254.0 (2.7)	0.314 (0.022)	0.416 (0.007)	357.9 (10.9)
		PC	308.4 (3.4)	327.6 (2.4)	164.4 (2.4)	0.466 (0.009)	0.259 (0.005)	471.6 (5.4)
		GES	555.8 (7.3)	236.6 (1.4)	255.4 (1.4)	0.540 (0.004)	0.322 (0.002)	537.0 (5.3)
	Hepar2 (4)	BCD	417.5 (3.4)	122.2 (3.1)	369.8 (3.1)	0.114 (0.003)	0.685 (0.006)	169.9 (3.3)
		Baseline	417.1 (4.0)	126.0 (2.1)	366.0 (2.1)	0.122 (0.006)	0.674 (0.005)	177.1 (3.3)
		PC	411.4 (2.9)	225.9 (1.3)	266.1 (1.3)	0.353 (0.005)	0.418 (0.003)	371.2 (3.3)
		GES	418.6 (3.4)	114.3 (2.6)	377.7 (2.6)	0.098 (0.003)	0.709 (0.004)	155.2 (2.3)

Table 3: Results for ordered DAGs on real network data. Block size is 20 for $n < p$ and 50 for $n > p$. The number under each DAG reports the number of times it is duplicated to form a large DAG. All numbers represent the average (standard error) over 10 simulations.

5.2.2 LEARNING WITH DE-CORRELATION

When the ordering of the DAG nodes is not given, we compared the effect of de-correlation as in Section 5.1.2. All network parameters were generated in the same way as before but we randomly shuffled the columns of X . The decrease in the structural Hamming distance

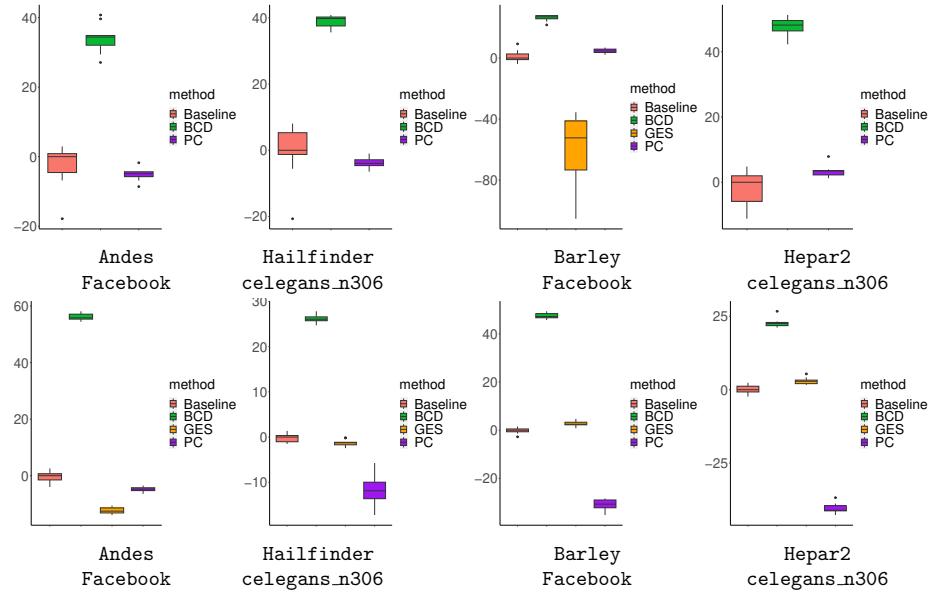


Figure 6: Test data log-likelihood normalized by \sqrt{np} on real sorted DAGs. Top row: $n < p$. Bottom row: $n > p$.

and increase in Jaccard index from de-correlation over 10 simulations are summarized as boxplots in Figure 7. PC performed uniformly better after de-correlation compared to before. GES and sparsebn also improved after de-correlation in most cases. The changes in the test data log-likelihood are shown in Figure 8, which are positive for almost all data sets, except two outliers (removed from plots) of sparsebn in the second and fourth panels in the top row.

5.2.3 LEARNING UNDER GENERAL COVARIANCE STRUCTURE

In the following experiments, we generate the support of Θ^* without the block-diagonal constraint by directly sampling the real undirected networks `facebook` and `celegans_n306`. In other words, the underlying undirected network may have only one connected subgraph where all individuals are dependent. This setup poses a major challenge particularly for the estimation of Θ because its support becomes much larger, and, as a result, we will need to impose stronger regularization in the graphical Lasso step when $n > p$. Moreover, we use Lasso regression loss (Hastie et al., 2015) to estimate Ω :

$$\hat{\omega}_j^2 = \min_{\beta} \frac{1}{2n} \|X_j - X_{\Pi_j} \beta\|_2^2 + \lambda \|\beta\|_1$$

as if the data were i.i.d., where $\lambda = \sqrt{\log(p)/n}$ and $\Pi_j = \{1, \dots, j-1\}$ is the index set of the potential parent nodes of j . This approximation performed well in our experiments. For simplicity, we focus on the setting $p > n$ so that we can still fix $\lambda_2 = 0.01$. Proceeding as before, we generate B^* from real DAGs with duplications: `Andes`, `Hailfinder`, `Barley`, and `Hepar2`.

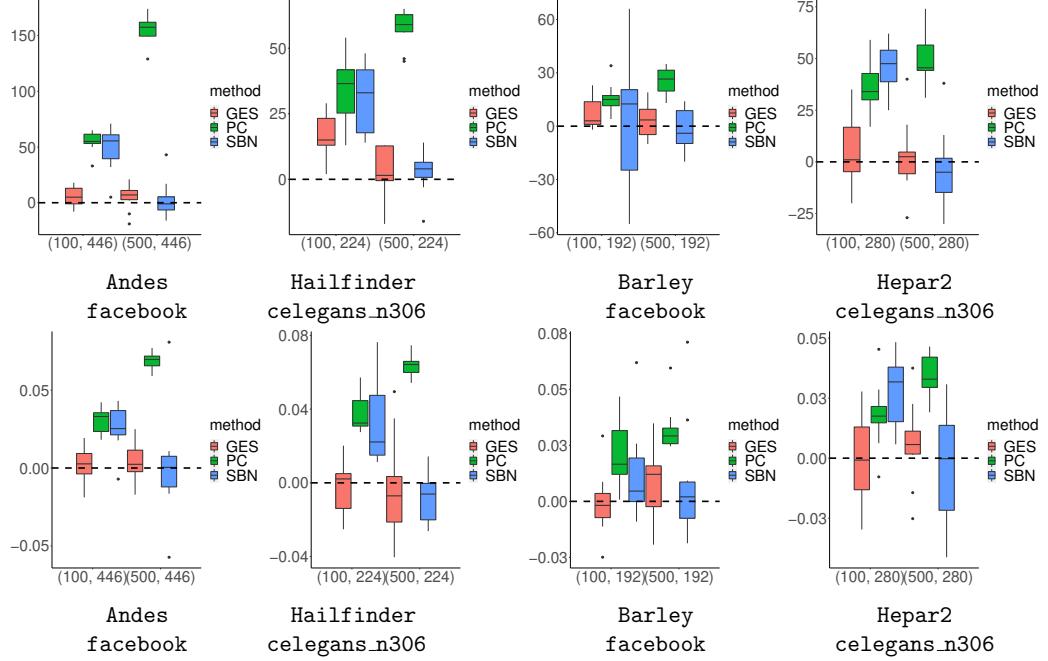


Figure 7: Experiments on real unsorted DAGs. Decrease in SHD (top row) and increase in JI (bottom row) via de-correlation for real networks, where the x-axis reports the value of (n, p) . In each panel, the three boxplots on the left and the three on the right correspond to cases of $n < p$ and $n > p$, respectively.

DAG	Θ -Network	Method	(n, p, s_0)	E	FN	TP	FDR	JI	SHD
Andes (2)	facebook	BCD	(100, 446, 676)	424.1 (16.4)	319.9 (6.9)	356.1 (6.9)	0.155 (0.016)	0.478 (0.004)	387.9 (3.6)
		Baseline	(100, 446, 676)	422.2 (16.2)	326.2 (5.6)	349.8 (5.6)	0.165 (0.019)	0.467 (0.003)	398.6 (6.2)
		PC	(100, 446, 676)	424.1 (16.4)	378.6 (4.9)	297.4 (4.9)	0.293 (0.016)	0.371 (0.003)	505.3 (7.9)
		GES	(100, 446, 676)	1782.0 (11.9)	314.5 (2.2)	361.5 (2.2)	0.797 (0.002)	0.173 (0.002)	1735.0 (13.4)
Hailfinder (4)	celegans_n306	BCD	(100, 224, 264)	140.8 (6.5)	153.3 (2.8)	110.7 (2.8)	0.206 (0.019)	0.376 (0.007)	183.4 (2.5)
		Baseline	(100, 224, 264)	139.0 (5.8)	154.8 (2.2)	109.2 (2.2)	0.205 (0.028)	0.372 (0.009)	184.6 (4.9)
		PC	(100, 224, 264)	141.5 (6.3)	160.0 (3.4)	104.0 (3.4)	0.261 (0.013)	0.344 (0.009)	197.5 (2.4)
		GES	(100, 224, 264)	312.3 (3.2)	136.6 (1.3)	127.4 (1.3)	0.592 (0.004)	0.284 (0.003)	321.5 (2.9)
Barley (4)	facebook	BCD	(100, 192, 336)	252.6 (10.1)	147.3 (4.6)	188.7 (4.6)	0.248 (0.014)	0.471 (0.007)	211.2 (3.0)
		Baseline	(100, 192, 336)	249.4 (10.0)	155.1 (3.0)	180.9 (3.0)	0.268 (0.020)	0.448 (0.005)	223.6 (5.1)
		PC	(100, 192, 336)	206.7 (2.7)	187.2 (1.1)	148.8 (1.1)	0.279 (0.009)	0.378 (0.004)	245.1 (2.6)
		GES	(100, 192, 336)	288.4 (4.2)	161.2 (1.3)	174.8 (1.3)	0.393 (0.006)	0.389 (0.002)	274.8 (2.5)
Hepar2 (6)	celegans_n306	BCD	(100, 420, 738)	492.3 (12.7)	386.7 (6.2)	351.3 (6.2)	0.285 (0.008)	0.399 (0.005)	527.7 (4.5)
		Baseline	(100, 420, 738)	490.8 (12.9)	397.8 (4.1)	340.2 (4.1)	0.303 (0.018)	0.384 (0.007)	548.4 (12.9)
		PC	(100, 420, 738)	454.1 (4.0)	467.8 (1.7)	270.2 (1.7)	0.405 (0.007)	0.293 (0.003)	651.7 (5.5)
		GES	(100, 420, 738)	970.7 (7.3)	358.3 (2.6)	379.7 (2.6)	0.609 (0.004)	0.286 (0.003)	949.3 (8.5)

Table 4: Results for ordered DAGs on real network data without block structure.

First, we assume that the DAG ordering is known and compare the BCD method against the baseline, PC, and GES methods. In the four cases that we considered, BCD gave better estimates of the DAG structure in terms of the Jaccard index and structural Hamming distance, as shown in Table 4. Next, without assuming a known DAG ordering, we compare the performance of GES, PC, and sparsebn before and after de-correlation. The top row in

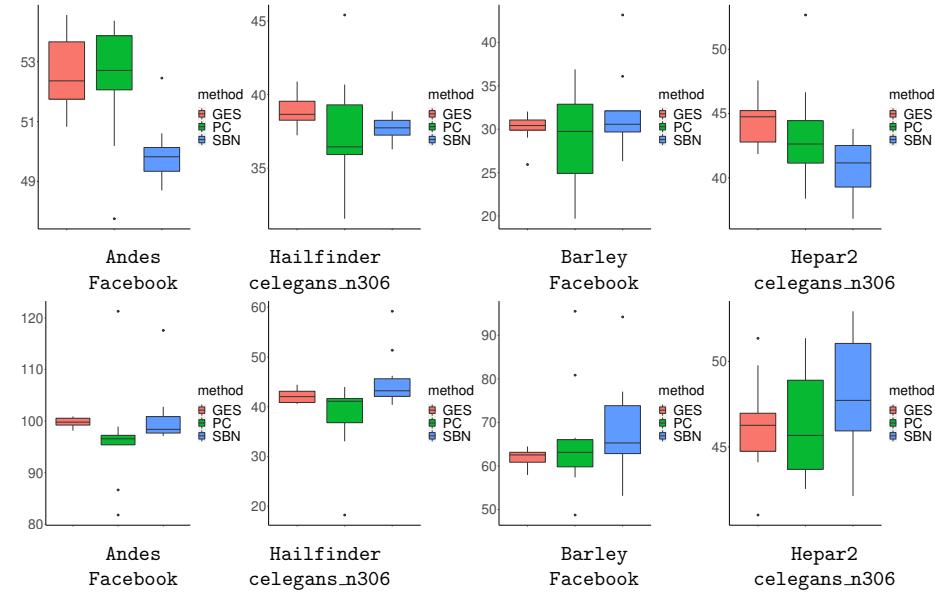


Figure 8: Increases in test data log-likelihood on real unsorted DAGs. Top row: $n < p$. Bottom row: $n > p$. Some outliers in the top panels are not shown for a better view of the boxplots.

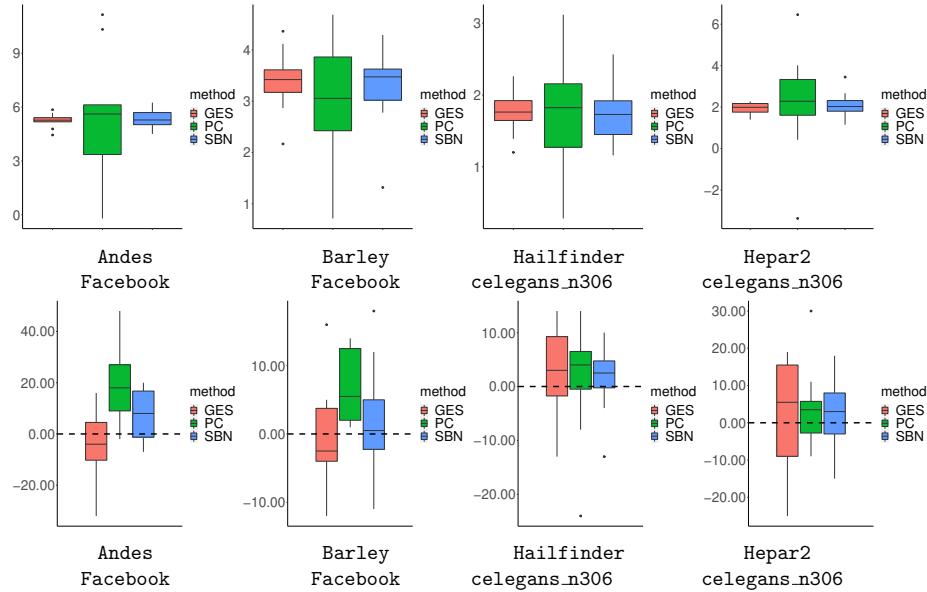


Figure 9: Results on real unsorted DAGs with general Θ . Top row: increase in normalized test data log-likelihood after de-correlation. Bottom row: Decrease in SHD after de-correlation.

Figure 9 shows the increase in the normalized test data log-likelihood after de-correlation, and the increases are positive across the 10 simulations for each of the four scenarios. The bottom row shows the distribution of the decrease in SHD across 10 simulations after

de-correlation. In most cases, all three methods gave much more accurate estimates after de-correlation. We defer the additional tables and figures containing more detailed results to the Supplementary Material. The above results confirm that our methods can indeed improve the accuracy in DAG estimation even though Θ^* is not block-diagonal, as suggested by the theoretical results in Section 4.

6. Application on RNA-seq Data

Gene regulatory networks (GRNs) enable biologists to examine the causal relations in gene expression during different biological processes, and are usually estimated from gene expression data. Recent advances in single-cell RNA sequencing technology have made it possible to trace cellular lineages during differentiation and to identify new cell types by measuring the gene expression of thousands of individual cells. A key question now is whether we can discover the GRN that controls cellular differentiation and drives transitions from one cell type to another using this type of data. Such GRNs can be interpreted as causal networks among genes, where nodes correspond to different genes, and a directed edge encodes a direct causal effect of one gene on another.

The RNA-seq data set used in this section were generated by Chu et al. (2016), and is accessible through the Gene Expression Omnibus (GEO) series accession number GSE75748. The data set contains measurements of gene expression of around 20,000 genes from $n = 1018$ cells. The cells are progenitor cells differentiated from human embryonic stem (ES) cells. The cell types correspond to different lineage-specific progenitors, undifferentiated ES cells, and foreskin fibroblasts (as a control). As lineage-specific progenitor cells were differentiated from the same population of ES cells, it is reasonable to assume dependence among the cells. This was also confirmed with the principal component analysis by Chu et al. (2016) (their Figure 1b). Before conducting the experiment, we processed the data according to Li and Li (2018) by imputing missing values and applying log transformation. In this study, we focus on estimating a GRN among $p = 51$ target genes selected by Chu et al. (2016), while the rest of the genes, which we call background genes, are used to estimate an undirected network for all 1018 cells.

6.1 Pre-estimate the Undirected Network

An essential input to Algorithm 1 is H^* , the undirected network of observations (cells in this case). In the dataset, progenitor cells differentiated from human ES cells included neuronal progenitor cells (NPCs, $n = 173$), endoderm derivatives cells (DE, $n = 138$), endothelial cells (ECs, $n = 105$), and trophoblast-like cells (TBs, $n = 69$). Undifferentiated H1 ($n = 212$) and H9 ($n = 162$) human ES cells and human foreskin fibroblasts (HFFs, $n = 159$) as control were also included. In total, the 1018 single cells in the dataset consist of the above seven distinct cell types. It is reasonable to assume that the similarity between cells of different types is minimal and we posit that the network of the 1018 cell consists of at least 7 connected components, i.e. $N \geq 7$, where N denotes the number of diagonal blocks in Θ^* as in Section 4. Although cells from different blocks may not be completely independent, their dependence should be much weaker than those in the same clusters. We used the block structure as an approximation, which greatly reduced the model and computational complexities. Since it is unlikely that all cells of the same type are strongly associated

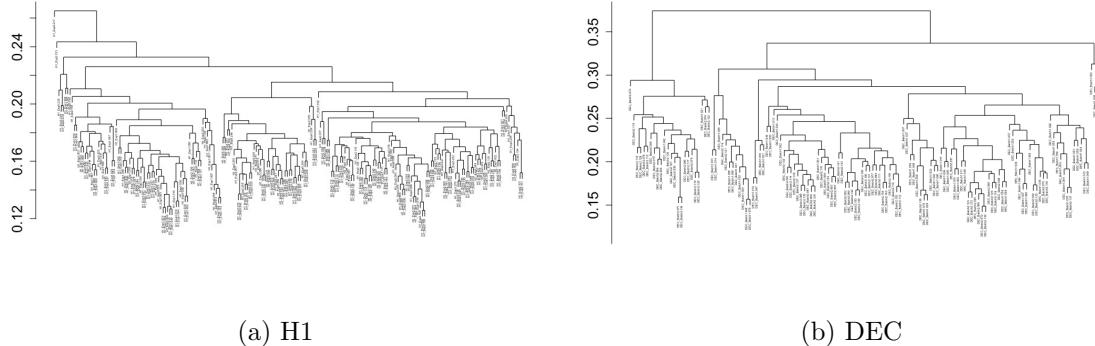


Figure 10: Cluster dendograms of H1 and DEC cells from hierarchical clustering. The y-axis represents $1 - \rho$ and the leaf nodes are individual cells.

with one another, we further divided each type of cells into smaller clusters by applying the classical clustering algorithm on the background genes. More specifically, we randomly selected 8000 genes from the background genes and applied hierarchical clustering on each type of cells. In this experiment, we used hierarchical clustering with complete-linkage and a distance metric between two cells defined as $1 - \rho$, where ρ is the correlation between their observed gene expression levels. We verified our choice of clustering algorithm by applying it on the entire data set, and it clearly grouped the cells into 7 groups, coinciding with the 7 cell types. At the end of the hierarchical clustering step, we needed to pick cut-off levels in order to finish clustering. Because the levels of dependence among cells are quite different across cell types as shown in Figure 10, we pick the cutoff points separately for each cell type. Generally, we chose the cutoff thresholds such that the largest cluster is smaller than $p = 51$, so λ_2 can be set to 0.01. By shifting the cutoff levels, we also obtained different number N of blocks in Θ^* . In the end, this clustering process returned an adjacency matrix A of the estimated network defined by the N clusters. In our experiments, we compared results from three choices of $N \in \{383, 519, 698\}$, much larger than the number of cell types. These refined small clusters will capture more subtle dependence relations among cells. The cluster size varied from 1 (singleton clusters) to 43 across the three values of N .

6.2 Model Evaluation

In this experiment, the input to the BCD algorithm is a data matrix $X_{1018 \times 51}$ and the network A estimated by the above hierarchical clustering. The matrix of error variances $\widehat{\Omega}$ was obtained following the method described in Section 4.2. The output is a solution path of $(\widehat{B}, \widehat{\Theta})$ for a range of λ_1 's. We computed the corresponding MLEs $(\widehat{B}^{MLE}, \widehat{\Omega}^{MLE})$ given the support of each \widehat{B} on the solution path. We picked the $(\widehat{B}^{MLE}, \widehat{\Omega}^{MLE})$ with the smallest BIC from the solution path as in Section 5 and used the corresponding $\widehat{\Theta}$ to de-correlate X . Table 5 shows the results of GES, PC, and sparsebn before and after de-correlation. In each case, we computed the BIC of the estimated GRN and used the Likelihood Ratio (LR) test to determine whether the increase in log-likelihood from de-correlation is significant or not.

Method	N	$\Delta(BIC)$	$LR \chi^2$	df	p-value	Z-scores
GES	383	6762.79	10148.79	3386	$< 10^{-5}$	82.2
PC		18171.20	21596.19	3425	$< 10^{-5}$	219.5
sparsebn		9341.59	12722.59	3381	$< 10^{-5}$	113.6
GES	519	-1723.65	8.34	1732	> 0.99	-29.2
PC		4635.81	6388.81	1753	$< 10^{-5}$	78.3
sparsebn		244.61	1976.61	1732	3.3×10^{-5}	4.2
GES	698	-114.97	573.03	688	> 0.99	-3.1
PC		2856.79	3552.79	696	$< 10^{-5}$	76.6
sparsebn		1198.33	1875.33	677	$< 10^{-5}$	32.5

Table 5: Compare goodness-of-fit of standard DAG learning methods with and without de-correlation of the data. N denotes the number of clusters among cells. $\Delta(BIC) = \text{BIC}(\text{baseline}) - \text{BIC}(\text{decor})$ is the change in BIC scores before and after de-correlation (positive value means improved model fitting). $LR \chi^2$ is the likelihood ratio statistic (24) followed by its degree of freedom (df). p values are computed based on χ^2 -test. We also include the Z-score = $(LR - df)/\sqrt{2df}$.

The LR test statistic is defined as follows:

$$LR = 2 \left[\log p(X | \hat{\Theta}, \hat{B}_{decor}^{MLE}, \hat{\Omega}_{decor}^{MLE}) - \log p(X | I_n, \hat{B}_{baseline}^{MLE}, \hat{\Omega}_{baseline}^{MLE}) \right], \quad (24)$$

where $(\hat{B}_{baseline}^{MLE}, \hat{\Omega}_{baseline}^{MLE})$ and $(\hat{B}_{decor}^{MLE}, \hat{\Omega}_{decor}^{MLE})$ denote the MLEs given the estimated graph structures from GES, PC, and sparsebn before and after de-correlation, respectively. If the baseline model (before de-correlation) is true, then the LR statistic follows approximately a χ^2 distribution with degrees of freedom

$$df = \frac{|\text{supp}(\hat{\Theta})| - n}{2} + |\text{supp}(\hat{B}_{decor}^{MLE})| - |\text{supp}(\hat{B}_{baseline}^{MLE})|.$$

As reported in Table 5, for most cases, we saw significant improvements, in terms of both the BIC and the χ^2 statistic, in all three DAG estimation methods by de-correlating X using the estimated $\hat{\Theta}$ from the BCD algorithm. To better quantify the significance, we also include a standardized Z-score of the LR statistic, $Z = (LR - df)/\sqrt{2df}$, as df and $2df$ are the mean and the variance of the LR statistic under the null. Since most standard DAG structure learning methods ignore data dependence, this application shows that considering such dependence would improve goodness of fit of DAG models. This confirms the dependence among individual cells and implies that our proposed network model fits this real-world data better.

Figure 11 shows the estimated CPDAGs after de-correlation for the case $N = 383$, which corresponds to the minimum BIC for all three methods in our experiments. It is interesting to note that a directed edge NANOG \rightarrow POU5F1, between the two master regulators in embryonic stem cells, appears in all three estimated CPDAGs, consistent with previously reported gene regulatory networks (Chen et al., 2008; Zhou et al., 2007).

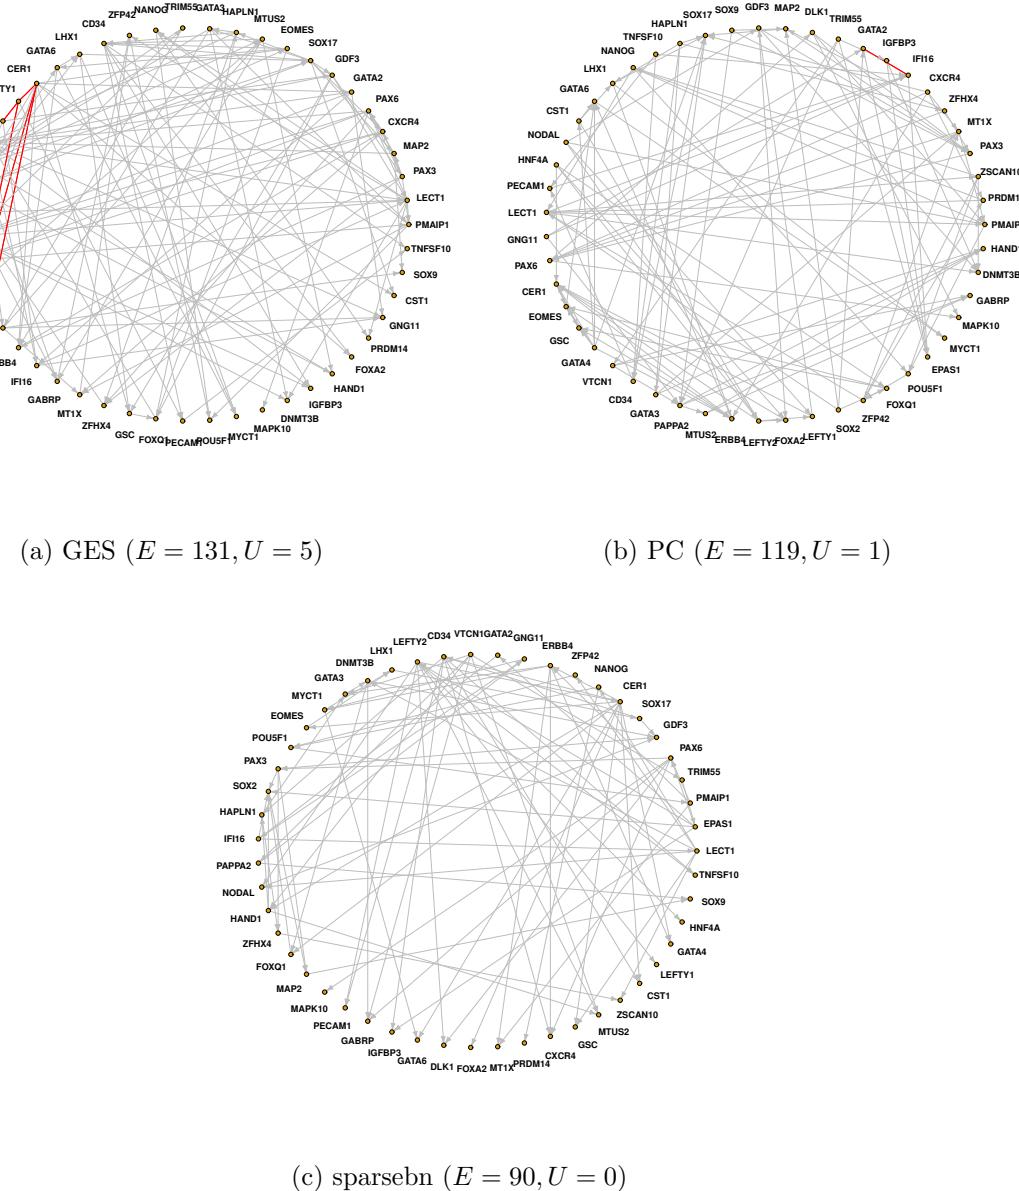


Figure 11: Estimated gene regulatory networks (CPDAGs) after de-correlation, with E edges and U undirected edges colored in red.

7. Discussion

In this paper our main goal is to generalize the existing Gaussian DAG model to dependent data. We proposed to model the covariance between observations by assuming a non-diagonal covariance structure of the noise vectors. Our main contributions include the development of a consistent structural learning method for the DAG and the sample network under sparsity assumptions and finite-sample guarantees for the estimators.

Our proposed BCD algorithm is built upon existing Lasso regression and graphical Lasso covariance estimation methods. When a topological ordering of the true DAG is known, it estimates the covariance between the observations Σ and the WAM of the DAG B in an iterative way. The method is fast and often converges in a few iterations. Our theoretical analysis shows that the estimates after one iteration are ℓ_2 -consistent under various asymptotic frameworks including both $n \ll p$ and $n \gg p$, assuming a proper initialization of the precision matrix $\widehat{\Theta}^{(0)}$. The estimate of the DAG WAM $\widehat{B}^{(1)}$ achieves the optimal rate as Lasso estimators. The estimate of the precision matrix $\widehat{\Theta}^{(1)}$ achieves the same optimal rate as the graphical Lasso method when $n \gg p$ and there are sufficiently many independent subgroups within the data. Otherwise, it has a slightly worse rate due to the bias of the sample covariance matrix. We have shown that covariance Σ is invariant under permutations of the DAG nodes. Therefore, when the ordering is unknown but the true DAG is sparse, our BCD algorithm can still give a good estimate of Θ which can be used to de-correlate the data. In addition to the theoretical analysis, we conducted extensive experiments on both synthetic and real data to compare our method with existing methods. When a true ordering of the DAG was given, the BCD algorithm significantly improved the structural estimation accuracy compared to the baseline method which ignored the sample dependency. When the ordering was unknown, classical DAG learning methods, such as GES, PC, and sparsebn, can all be greatly improved with respect to structural learning of CPDAGs by using our proposed de-correlation method based on the BCD algorithm. In all cases, our estimation methods under the proposed network Gaussian DAG model yielded significantly higher test data log-likelihood compared to other competing methods, indicating better predictive modeling performance.

There are several unexplored directions from our research. First, the current error bounds and consistency results are based on a known topological ordering of the true DAG. In practice, however, it can be hard to obtain the order in advance. It would be interesting to see if we can combine the method of estimating partial orders such as Niinimäki et al. (2016) with our method and extend the theoretical results. Second, part of the reason the current model relies on a known DAG ordering is the lack of experimental data. From purely observational data, it is impossible to orient some of the edges and find a topological ordering of the true DAG. In the next step, we would like to extend our method to handle both observational and experimental data sets. Finally, there are recent methods that use continuous optimization for DAG learning without imposing the acyclicity constraint, such as NOTEARS (Zheng et al., 2018). It is a promising future direction to incorporate such ideas into DAG learning on network data.

Acknowledgments

This work was supported by NSF grants DMS-1952929 and DMS-2305631 (to Q.Z.).

Appendix A. Technical Details of Section 4

A.1 Some Auxiliary Results

Here we introduce four lemmas that we use to establish the error bounds of $\widehat{\Theta}^{(1)}$ and $\widehat{B}^{(1)}$. Let us start by deriving an upper bound on the ℓ_2 deviation of $\widehat{L}^{(0)}$ from L^* under Assumption 1.

Lemma 11 *Suppose Assumption 1 holds and let L^* and $\widehat{L}^{(t)}$ be the Cholesky factors of Θ^* and $\widehat{\Theta}^{(t)}$, respectively. Let $\widehat{\Delta}_{chol}^{(t)} = \widehat{L}^{(t)} - L^*$. Then,*

$$\|\widehat{\Delta}_{chol}^{(0)}\|_2 \leq \frac{M}{2\sigma_{\min}(L^*)},$$

where $\sigma_{\min}(L^*)$ is the smallest singular value of L^* , and M is from Assumption 1.

To generalize the basic bound on $\|\widehat{\beta}^{lasso} - \beta_j^*\|_2$ from Buhlmann and van de Geer (2011) to dependent data, we need to control the ℓ_∞ -norm of an empirical process component $2X^\top \widehat{\Theta}^{(0)} \varepsilon_j/n$. Let us start with the case when the data are independent. Define the following events, where $\tilde{X} = L^*X$ represents the independent data as explained in Section 4:

$$\mathcal{E} := \bigcap_{k=1}^p \left\{ \|\tilde{X}_k\|_2 \leq 6\bar{\psi}\sqrt{n} \right\}. \quad (25)$$

Then the following lemma follows from \tilde{X} being sub-Gaussian.

Lemma 12 *Let $\alpha \geq 2$ be an integer. If $n > 2\sqrt{2\alpha} \log p$, then the event \mathcal{E} defined in (25) holds with probability at least $1 - 1/p^{\alpha-1}$.*

Next, let $\tilde{X}_{[j-1]}$ denote the first $j-1$ columns in \tilde{X} and define the following events that depend on λ_n

$$\begin{aligned} \mathcal{I}_j &:= \left\{ 2\|\tilde{X}_{[j-1]}^\top \tilde{\varepsilon}_j\|_\infty/n \leq \lambda_n \right\}, \quad j = 1, \dots, p \\ \mathcal{T} &:= \bigcap_{j=1}^p \mathcal{I}_j, \end{aligned} \quad (26)$$

where $\tilde{\varepsilon}_j = L^* \varepsilon_j$ for $j \in [p]$.

Lemma 13 *Let \tilde{X}_k consist of n i.i.d sub-Gaussian random variables with parameter $\bar{\psi}^2$ for $k = 1, \dots, p$. If $n > 4 \log p$ and*

$$\lambda_n = 12\bar{\psi}\bar{\omega} \left(\sqrt{\frac{2 \log p}{n}} + \sqrt{\frac{2 \log 2 + 4 \log p}{n}} \right),$$

then the probability of \mathcal{T} satisfies

$$\mathbb{P}(\mathcal{T}) \geq \left(1 - \frac{1}{p} \right)^2.$$

Lemma 13 implies that if $\lambda_n \asymp \sqrt{\frac{\log p}{n}}$, then the error terms will be uniformly under control with high probability, especially when both n and p are large.

Lemma 14 (Maximal inequality) *Let $x_i = (x_{i1}, \dots, x_{ip})$ be a random vector where each element x_{ij} is sub-Gaussian with parameter $\bar{\psi}^2$, then*

$$\mathbb{P} \left(\|x_i\|_\infty \geq 2\bar{\psi}\sqrt{\log p} \right) \leq 2/p.$$

From model (3), it is clear that each row in \tilde{X} is sub-Gaussian with parameter $\bar{\psi}^2$. By Lemma 14, we have $\|\tilde{x}_i\|_\infty \lesssim \sqrt{\log p}$ w.h.p.

A.2 Error Bound of $\hat{B}^{(1)}$

The estimation error bound for the classical Lasso problem where samples are i.i.d. was established by choosing a penalty parameter that dominates the measurement error term. Specifically, as shown in Lemma 13, this can be achieved with high probability by setting $\lambda_n \asymp \sqrt{\frac{\log p}{n}}$. In order to prove the consistency of $\hat{B}^{(1)}$ from Algorithm 1, we need to control a similar error term which depends on $\hat{\Theta}^{(0)}$. Notably, such error can be controlled under the same rate as λ_n :

Theorem 15 (Control the empirical process) *Let λ_n be the same as in Lemma 13. Suppose the initial estimator $\hat{\Theta}^{(0)}$ satisfies Assumption 1. Then for $n > 4\log p$,*

$$\mathbb{P} \left(\sup_{j \in [p]} 2\|X^\top \hat{\Theta}^{(0)} \varepsilon_j\|_\infty / n \leq \lambda_n \right) \geq \left(1 - \frac{1}{p}\right) \left(1 - \frac{2}{p}\right). \quad (27)$$

Next, we show that the random matrix $\hat{L}^{(0)}X$ satisfies the Restricted Eigenvalue (RE) condition (Wainwright, 2019) w.h.p. Towards that end, we define the event \mathcal{K} as in Theorem 7.16 from Wainwright (2019) given as

$$\mathcal{K} := \left\{ \|\tilde{X}\beta\|_2^2 / n \geq \tilde{c}_1 \|\sqrt{\Psi^*} \beta\|_2^2 - \tilde{c}_2 \rho^2(\Psi^*) \frac{\log p}{n} \|\beta\|_1^2 \right\}, \quad (28)$$

where $\rho^2(\Psi^*)$ is the maximum diagonal entry of Ψ^* , $\tilde{X} = L^*X$ is the de-correlated data, and $\tilde{c}_1 < 1 < \tilde{c}_2$ are positive constants.

Lemma 16 (Restricted eigenvalue condition) *Consider a random matrix $X \in \mathbb{R}^{n \times p}$ drawn from $\mathcal{N}_{n \times p}(0, \Sigma^*, \Psi^*)$. Let $\hat{\Theta}^{(0)}$ be the initial estimate of $\Theta^* = \Sigma^{*-1}$ satisfying Assumption 1, $\hat{L}^{(0)}$ be the Cholesky factor of $\hat{\Theta}^{(0)}$, and $\rho^2(\Psi^*)$ be the maximum diagonal entry of Ψ^* . Then under event \mathcal{K} defined in (28), there are universal positive constants $c_1 < 1 < c_2$ such that*

$$\frac{\|\hat{L}^{(0)}X\beta\|_2^2}{n} \geq c_1 \|\sqrt{\Psi^*} \beta\|_2^2 - c_2 \rho^2(\Psi^*) \frac{\log p}{n} \|\beta\|_1^2, \quad (29)$$

for all $\beta \in \mathbb{R}^p$.

Theorem 7.16 from Wainwright (2019) shows that the event \mathcal{K} happens with high probability. This event is a restriction on the design matrix \tilde{X} and it holds with high probability for a variety of matrix ensembles. With Theorem 15 and Lemma 16, it is possible to prove an *oracle inequality* for the dependent Lasso problem, which yields a family of upper bounds on the estimation error.

Theorem 17 (Lasso oracle inequality) *Consider the Lasso problem in (9) for $t = 0$. Suppose the inequality (29) and the event in (27) hold. Let $\bar{\kappa} = \sigma_{\min}(\Psi^*)$. For $j \in [p]$ and any $\beta_j^* \in \mathbb{R}^p$, if*

$$\lambda_n \geq 12\bar{\psi}\bar{\omega} \left(\sqrt{\frac{2\log p}{n}} + \sqrt{\frac{2\log 2 + 4\log p}{n}} \right),$$

then any optimal solution $\hat{\beta}_j^{(1)}$ satisfies:

$$\|\hat{\beta}_j^{(1)} - \beta_j^*\|_2^2 \leq \frac{144\lambda_n^2}{c_1^2\bar{\kappa}^2}|S| + \frac{32\lambda_n}{c_1\bar{\kappa}}\|\beta_{j,S^c}^*\|_1 + \frac{128c_2}{c_1}\frac{\rho^2(\Psi^*)}{\bar{\kappa}}\frac{\log p}{n}\|\beta_{j,S^c}^*\|_1^2, \quad (30)$$

for any subset S with cardinality $|S| \leq \frac{c_1}{64c_2}\frac{\bar{\kappa}}{\rho^2(\Psi^*)}\frac{n}{\log p}$. Let $\hat{L}^{(0)}$ be the Cholesky factor of $\hat{\Theta}^{(0)}$. Then,

$$\frac{1}{n}\left\|\hat{L}^{(0)}X(\hat{\beta}_j^{(1)} - \beta_j^*)\right\|_2^2 \leq 6\lambda_n\|\beta_j^*\|_1. \quad (31)$$

Theorem 17 implies, with high probability, that $\sup_{j \in [p]} \|\hat{\beta}_j^{(1)} - \beta_j^*\|_2^2 \leq \frac{768\lambda_n^2}{c_1^2\bar{\kappa}^2}s \asymp s\frac{\log p}{n}$, where s is the maximum in-degree of the true DAG.

A.3 Error Bounds of $\hat{\Theta}^{(1)}$

Recall that s denotes the maximum number of nonzero entries in β_j^* for $j \in [p]$. In order to control $\|\hat{\Theta}^{(1)} - \Theta^*\|_2$, we need to rely on certain type of error bound on $\hat{S}^{(1)} - \Sigma^*$, where $\hat{S}^{(1)}$ is the sample covariance defined in (10) for $t = 0$. Therefore, we adopt the definition of tail condition on the sample covariance from Ravikumar et al. (2011).

Definition 18 (Tail conditions) *We say the $n \times p$ random matrix X from model (2) satisfies tail condition $\mathcal{T}(f, v_*)$ if there exists a constant $v_* \in (0, \infty]$ and a function $f : \mathbb{N} \times (0, \infty) \rightarrow (0, \infty)$ such that for any $(i, j) \in [n] \times [n]$,*

$$\mathbb{P}\left[|\hat{S}_{ij}^{(1)} - \Sigma_{ij}^*| \geq \delta\right] \leq 1/f(p, \delta) \quad \forall \delta \in (0, 1/v_*].$$

We require $f(p, \delta)$ to be monotonically increasing in p , so for a fixed $\delta > 0$, define the inverse function

$$\bar{p}_f(\delta; r) := \arg \max \{p \mid f(p, \delta) \leq r\}.$$

Similarly, f should be increasing in δ for each fixed p , so we define an inverse function in the second argument:

$$\bar{\delta}_f(p; r) := \arg \max \{\delta \mid f(p, \delta) \leq r\}. \quad (32)$$

Under the setting of a Gaussian DAG model, we can derive a sub-Gaussian tail bound.

Lemma 19 *Let X be a sample from our Gaussian DAG model (2). The sample covariance matrix*

$$\widehat{\Sigma} = \frac{1}{p} \sum_{j=1}^p \frac{1}{\omega_j^{2*}} (X_j - X\beta_j^*) (X_j - X\beta_j^*)^\top \quad (33)$$

satisfies the tail bound

$$\mathbb{P} \left(\sup_{i,j} |\widehat{\Sigma}_{ij} - \Sigma_{ij}^*| > \delta \right) \leq 4 \exp \left\{ - \frac{p\delta^2}{3200} \right\},$$

for all $\delta \in (0, 40)$.

Corollary 20 *If $f(p, \delta) = 4 \exp \left\{ \frac{p\delta^2}{3200} \right\}$, then the inverse function $\bar{\delta}_f(p; n^\tau)$ takes the following form,*

$$\bar{\delta}_f(p; n^\tau) = 40\sqrt{2} \sqrt{\frac{\tau \log n - \log 4}{p}}.$$

Based on the tail bound in Corollary 20, we can control the sampling noise $\widehat{\Sigma} - \Sigma^*$ as in Lemma 21.

Lemma 21 (Lemma 8 in Ravikumar et al. 2011) *Define event*

$$\mathcal{A} = \left\{ \|\widehat{\Sigma} - \Sigma^*\|_\infty \leq \bar{\delta}_f(p; n^\tau) \right\}, \quad (34)$$

where $\bar{\delta}_f(p; n^\tau) = 40\sqrt{2} \sqrt{\frac{\tau \log n - \log 4}{p}}$. For any $\tau > 2$ and (n, p) such that $\bar{\delta}_f(p; n^\tau) \leq 1/40$, we have

$$\mathbb{P}[\mathcal{A}^c] \leq \frac{1}{n^{\tau-2}} \rightarrow 0, \quad \text{as } n \rightarrow \infty.$$

Recall $r(\widehat{\Omega})$ defined in (12) and the constant b defined in Lemma 7.

Lemma 22 *Suppose $b > 0$ and*

$$\sup_j \|\hat{\beta}_j^{(1)} - \beta_j^*\|_2^2 \leq c \cdot s \frac{\log p}{n} \leq c \cdot s^2 \frac{\log p}{n} < 1, \quad r(\widehat{\Omega}) \lesssim 1, \quad (35)$$

for a fixed positive constant c . Let $\widehat{W}^{(1)} := \widehat{S}^{(1)} - \Sigma^$. Then, we have*

$$\|\widehat{W}^{(1)}\|_\infty \leq 40\sqrt{2} \sqrt{\frac{\tau \log n}{p}} + \max \left\{ 6\bar{\omega}r(\widehat{\Omega}), \frac{144\bar{\omega}\bar{\psi}s}{b} \sqrt{\frac{c \log p \log^2 \max\{n, p\}}{n}} \right\}$$

with probability at least $1 - 1/n^{\tau-2} - 5n^2/\max\{n, p\}^4$.

Theorem 23 Assume $\widehat{B}^{(1)}$ satisfies (35) and $b > 0$. Let

$$\begin{aligned}\bar{R}(s, p, n) &= \max \left\{ 6\bar{\omega}r(\widehat{\Omega}), \frac{144\bar{\omega}\bar{\psi}s}{b} \sqrt{\frac{\log p \log^2 \max\{n, p\}}{n}} \right\}, \\ \bar{\delta}_f(p; n^\tau) &= 40\sqrt{2} \sqrt{\frac{\tau \log n - \log 4}{p}}.\end{aligned}$$

Consider the graphical Lasso estimate $\widehat{\Theta}^{(1)}$ from Algorithm 1 with $\lambda_p = \bar{\delta}_f(p; n^\tau) + \bar{R}$ for $\tau > 2$. Assume

$$\bar{p}_f(1/\max\{160, 24mC\}; n^\tau) \leq p \quad \text{and} \quad \bar{R} \leq \frac{1}{24mC}, \quad (36)$$

where $C = \max\{\kappa_{\Sigma^*}\kappa_{\Gamma^*}, \kappa_{\Sigma^*}^3\kappa_{\Gamma^*}^2\}$, and m is the maximum degree of the undirected network H^* . Then, with probability at least $1 - 1/n^{\tau-2} - 5n^2/\max\{n, p\}^4$, we have

$$\begin{aligned}\|\widehat{\Theta}^{(1)} - \Theta^*\|_\infty &\leq 4\kappa_{\Gamma^*}(\bar{\delta}_f(p; n^\tau) + \bar{R}), \\ \|\widehat{\Theta}^{(1)} - \Theta^*\|_2 &\leq 4\kappa_{\Gamma^*}(m+1)(\bar{\delta}_f(p; n^\tau) + \bar{R}).\end{aligned}$$

Appendix B. Proofs of Main Results

We collect the proofs of our main theoretical results here, including Theorem 2 in Section 2, Proposition 3 in Section 3, and Theorem 4, Corollary 8 and Corollary 9 in Section 4.

B.1 Proof of Proposition 3

Proof First, notice that, with probability one, X has full column rank. Also, because the Cholesky factor $\widehat{L}^{(t)}$ is always positive definite for each iteration, $\widehat{L}^{(t)}X$ is in *general position* a.s. Note that the two terms of (8) are differentiable (regular) and the whole function is continuous. Furthermore, solving (8) with respect to each variable gives a unique coordinate-wise minimum. Therefore, by Theorem 4.1 (c) in Tseng (2001), the block coordinate descent converges to a stationary point. \blacksquare

B.2 Proof of Theorem 2

Proof According to the matrix normal presentation (3), we write $L(B, \Omega, \Theta) = L(\Psi(B, \Omega), \Theta)$ as the likelihood, where $\Psi(B, \Omega) = (I - B)^{-\top} \Omega (I - B)^{-1}$. Since \mathcal{G}_1 and \mathcal{G}_2 are Markov equivalent, they represent the same set of Gaussian distributions, each parameterized by a covariance matrix Ψ . Let $(\widehat{B}_m, \widehat{\Omega}_m, \widehat{\Theta}_m)$ be any MLE given \mathcal{G}_m for $m = 1, 2$. Then, there exists $(\widetilde{B}_2, \widetilde{\Omega}_2)$, where \widetilde{B}_2 is a WAM for \mathcal{G}_2 , such that $\Psi(\widehat{B}_1, \widehat{\Omega}_1) = \Psi(\widetilde{B}_2, \widetilde{\Omega}_2)$. Therefore,

$$L(\widehat{B}_1, \widehat{\Omega}_1, \widehat{\Theta}_1) = L(\widetilde{B}_2, \widetilde{\Omega}_2, \widehat{\Theta}_1) \leq L(\widetilde{B}_2, \widetilde{\Omega}_2, \widehat{\Theta}_2).$$

By symmetry, $L(\widehat{B}_2, \widehat{\Omega}_2, \widehat{\Theta}_2) \leq L(\widehat{B}_1, \widehat{\Omega}_1, \widehat{\Theta}_1)$, and thus, $L(\widehat{B}_1, \widehat{\Omega}_1, \widehat{\Theta}_1) = L(\widehat{B}_2, \widehat{\Omega}_2, \widehat{\Theta}_2)$. \blacksquare

B.3 Proof of Theorem 4

Proof We first prove the consistency in $\widehat{B}^{(1)}$. Under Assumption 1, Theorem 15 shows that for the given λ_n , the empirical process term of the noises can be uniformly bounded with high probability. Therefore, in order to obtain the conclusion in Theorem 17, we only need the inequality (29) in Lemma 16 to hold. Since the event \mathcal{K} in (28) holds with high probability by Theorem 7.16 in Wainwright (2019), (29) holds by Lemma 16. Next, we show $\widehat{\Theta}^{(1)}$ is consistent by invoking Theorem 23. For the chosen λ_p and under the constraint on (n, p) specified in (16) and (17), the sample size requirement in (36) is satisfied. Therefore, the results follow from Theorem 23. Combining Theorems 17 and 23, we get the desired results. ■

B.4 Proof of Corollary 8

Proof The rate of $\sup_j \|\widehat{\beta}_j^{(1)} - \beta_j^*\|_2^2$ follows directly from the choice of $\lambda_n \asymp \sqrt{\frac{\log p}{n}}$. Since $r(\widehat{\Omega}) = O_p\left(s\sqrt{\frac{\log p}{N}}\right)$ and $p \gg n$,

$$\begin{aligned} \|\widehat{\Theta}^{(1)} - \Theta^*\|_2 &= O_p\left(m\left(\sqrt{\frac{\log n}{p}} + s \max\left\{\sqrt{\frac{\log p}{N}}, \sqrt{\frac{\log^3 p}{n}}\right\}\right)\right) \\ &= O_p\left(ms \max\left\{\sqrt{\frac{\log p}{N}}, \sqrt{\frac{\log^3 p}{n}}\right\}\right) \quad (n \gtrsim N) \\ &= O_p\left(ms\sqrt{\frac{\log^3 p}{n}}\right) \quad (N \log^2 p \gtrsim n). \end{aligned}$$

■

B.5 Proof of Corollary 9

Proof The rate of $\sup_j \|\widehat{\beta}_j^{(1)} - \beta_j^*\|_2^2$ can be derived in the same way as in the proof of Corollary 8. Since $r(\widehat{\Omega}) = O_p\left(s\sqrt{\frac{\log p}{N}}\right)$ and $n \gg p$,

$$\begin{aligned} \|\widehat{\Theta}^{(1)} - \Theta^*\|_2 &= O_p\left(m\left(\sqrt{\frac{\log n}{p}} + s \max\left\{\sqrt{\frac{\log p}{N}}, \sqrt{\frac{\log p \log^2 n}{n}}\right\}\right)\right) \\ &= O_p\left(m\left(\sqrt{\frac{\log n}{p}} + s \max\left\{\sqrt{\frac{\log p}{N}}, \sqrt{\frac{\log p \log^2 n}{n}}\right\}\right)\right). \end{aligned}$$

Since $N \gtrsim s^2 p \implies \sqrt{\frac{\log n}{p}} \gtrsim \sqrt{\frac{s^2 \log p}{N}}$ and $n \gg s^2 p \log p \log n \implies \sqrt{\frac{\log n}{p}} \gtrsim \sqrt{\frac{s^2 \log p \log^2 n}{n}}$, we arrive at

$$\|\widehat{\Theta}^{(1)} - \Theta^*\|_2 = O_p \left(m \sqrt{\frac{\log n}{p}} \right).$$

■

Appendix C. Proofs of Intermediate Results for Theorem 4

We include the proofs for all the intermediates results that lead to Theorem 4 in this section.

C.1 Proof of Theorem 15

Proof For any $j = 1, \dots, p$,

$$\begin{aligned} \frac{2}{n} \|X^\top \widehat{\Theta}^{(0)} \varepsilon_j\|_\infty &= \frac{2}{n} \|X^\top (\Theta^* + \widehat{\Delta}_{prec}^{(0)}) \varepsilon_j\|_\infty \leq \frac{2}{n} \|\tilde{X}^\top \tilde{\varepsilon}_j\|_\infty + \frac{2}{n} \|X^\top \widehat{\Delta}_{prec}^{(0)} \varepsilon_j\|_\infty \\ &\leq \frac{2}{n} \|\tilde{X}^\top \tilde{\varepsilon}_j\|_\infty + \frac{2}{n} \|\tilde{X}^\top L^{*-T} \widehat{\Delta}_{prec}^{(0)} L^{*-1} \tilde{\varepsilon}_j\|_\infty. \end{aligned}$$

Let $\widehat{K}^{(0)} = L^{*-T} \widehat{\Delta}_{prec}^{(0)} L^{*-1}$. Then following Assumption 1,

$$\|\widehat{K}^{(0)}\|_2 \leq \|\widehat{\Delta}_{prec}^{(0)}\|_2 / \sigma_{\min}^2(L^*) \leq M / \sigma_{\min}^2(L^*).$$

Under event \mathcal{E} defined in (25), for $j \in [p]$,

$$\|\widehat{K}^{(0)} \tilde{X}_j\|_2 \leq 6\bar{\psi}M\sqrt{n} / \sigma_{\min}^2(L^*) \leq 6\bar{\psi}\sqrt{n}.$$

For $j \in [p]$, define the event

$$\overline{\mathcal{T}}_j := \left\{ 2\|\tilde{X}^\top \widehat{K}^{(0)} \tilde{\varepsilon}_j\|_\infty / n < \lambda_n / 2 \right\}.$$

Similar to the proof of Lemma 13, we can show

$$\mathbb{P} \left(\bigcup_{j=1}^p \overline{\mathcal{T}}_j^c \mid \mathcal{E} \right) \leq \frac{1}{p},$$

and together with Lemma 12 we have

$$\mathbb{P} \left(\overline{\mathcal{T}} \bigcap \mathcal{T} \right) \geq \mathbb{P} \left(\overline{\mathcal{T}} \bigcap \mathcal{T} \mid \mathcal{E} \right) \mathbb{P}(\mathcal{E}) \geq \left(1 - \frac{2}{p} \right) \left(1 - \frac{1}{p} \right) \geq \left(1 - \frac{2}{p} \right)^2,$$

where \mathcal{T} is defined in (26). ■

C.2 Proof of Lemma 16

Proof We observe that

$$\begin{aligned}
\|\widehat{L}^{(0)}X\theta\|_2/\sqrt{n} &\geq \|L^*X\theta\|_2\sqrt{n} - \|\widehat{\Delta}_{chol}^{(0)}X\theta\|_2/\sqrt{n} \\
&= \|L^*X\theta\|_2/\sqrt{n} - \|\widehat{\Delta}_{chol}^{(0)}L^{*-1}L^*X\theta\|_2/\sqrt{n} \\
&\geq \left(1 - \|\widehat{\Delta}_{chol}^{(0)}\|_2/\sigma_{\min}(L^*)\right) \|L^*X\theta\|_2/\sqrt{n} \\
&\geq \left(1 - \frac{M}{2\sigma_{\min}^2(L^*)}\right) \|\widetilde{X}\theta\|_2/\sqrt{n} \quad (\text{By Assumption 1 and Lemma 11}) \\
&\geq \frac{1}{2\sqrt{n}} \|\widetilde{X}\theta\|_2,
\end{aligned}$$

when n is sufficiently large. Since event \mathcal{K} defined in (28) holds, by Theorem 7.16 in Wainwright (2019), we have

$$\begin{aligned}
\frac{\|\widehat{L}^{(0)}X\theta\|_2^2}{n} &\geq \frac{1}{4} \left(\tilde{c}_1 \|\sqrt{\Psi^*}\theta\|_2^2 - \tilde{c}_2 \rho^2(\Psi^*) \frac{\log p}{n} \|\theta\|_1^2 \right) \\
&= c_1 \|\sqrt{\Psi^*}\theta\|_2^2 - c_2 \rho^2(\Psi^*) \frac{\log p}{n} \|\theta\|_1^2,
\end{aligned}$$

where $c_j = \tilde{c}_j/4, j = 1, 2$. ■

C.3 Proof of Theorem 17

Proof Consider the penalized negative likelihood function from (9):

$$\ell(\beta_j, \lambda_n) = \frac{1}{2n} \|\widehat{L}^{(0)}X_j - \widehat{L}^{(0)}X\beta_j\|_2^2 + \lambda_n \|\beta_j\|_1.$$

For simplicity, we drop the superscript (t) in $\hat{\beta}^{(1)}$ and $\widehat{L}^{(0)}$. Let ρ stand for $\rho(\Psi^*)$, $\beta_j^* \in \mathbb{R}^p$, and $\widehat{\Delta}_j = \hat{\beta}_j - \beta_j^*$. We start from the *basic inequality* (Wainwright, 2019):

$$\ell(\hat{\beta}_j, \lambda_n) \leq \ell(\beta_j^*, \lambda_n) = \frac{1}{2n} \|\widehat{L}\varepsilon_j\|_2^2 + \lambda_n \|\beta_j^*\|_1.$$

After rearranging some terms,

$$0 \leq \frac{1}{2n} \|\widehat{L}X\widehat{\Delta}_j\|_2^2 \leq \frac{\varepsilon_j^\top \widehat{\Theta}X\widehat{\Delta}_j}{n} + \lambda_n \left(\|\beta_j^*\|_1 - \|\hat{\beta}_j\|_1 \right). \quad (37)$$

Next, for any subset $S \subseteq [p]$, we have

$$\|\beta_j^*\|_1 - \|\hat{\beta}_j\|_1 = \|\beta_{j,S}^*\|_1 + \|\beta_{j,S^c}^*\|_1 - \|\beta_{j,S}^* + \widehat{\Delta}_{j,S}\|_1 - \|\widehat{\Delta}_{j,S^c} + \beta_{j,S^c}^*\|_1. \quad (38)$$

Combining (37) with (38), and applying triangle and Hölder's inequalities,

$$\begin{aligned}
 0 \leq \frac{1}{2n} \|\widehat{L}X\widehat{\Delta}_j\|_2^2 &\leq \frac{1}{n} \varepsilon_j^\top \widehat{\Theta}X\widehat{\Delta}_j + \lambda_n \left(\|\widehat{\Delta}_{j,S}\|_1 - \|\widehat{\Delta}_{j,S^c}\|_1 + 2\|\beta_{j,S^c}^*\|_1 \right) \\
 &\leq \|X^\top \widehat{\Theta}\varepsilon_j\|_\infty / n \|\widehat{\Delta}_j\|_1 + \lambda_n \left(\|\widehat{\Delta}_{j,S}\|_1 - \|\widehat{\Delta}_{j,S^c}\|_1 + 2\|\beta_{j,S^c}^*\|_1 \right) \\
 &\leq \frac{\lambda_n}{2} \left(\|\widehat{\Delta}_j\|_1 + 2\|\widehat{\Delta}_{j,S}\|_1 - 2\|\widehat{\Delta}_{j,S^c}\|_1 + 4\|\beta_{j,S^c}^*\|_1 \right) \\
 &\leq \frac{\lambda_n}{2} \left[3\|\widehat{\Delta}_{j,S}\|_1 - \|\widehat{\Delta}_{j,S^c}\|_1 + 4\|\beta_{j,S^c}^*\|_1 \right], \tag{39}
 \end{aligned}$$

and so

$$\|\widehat{\Delta}_j\|_1 \leq 4 \left(\|\widehat{\Delta}_{j,S}\|_1 + \|\beta_{j,S^c}^*\|_1 \right).$$

This inequality implies (by Cauchy-Schwarz inequality)

$$\|\widehat{\Delta}_j\|_1^2 \leq \left(4\|\widehat{\Delta}_{j,S}\|_1 + 4\|\beta_{j,S^c}^*\|_1 \right)^2 \leq 32 \left(|S| \|\widehat{\Delta}_j\|_2^2 + \|\beta_{j,S^c}^*\|_1^2 \right). \tag{40}$$

Next, from (29) and (40), we know,

$$\begin{aligned}
 \frac{\|\widehat{L}X\widehat{\Delta}_j\|_2^2}{n} &\geq \left(c_1 \bar{\kappa} - 32c_2 \rho^2 |S| \frac{\log p}{n} \right) \|\widehat{\Delta}_j\|_2^2 - 32c_2 \rho^2 \frac{\log p}{n} \|\beta_{j,S^c}^*\|_1^2 \\
 &\geq \frac{c_1 \bar{\kappa}}{2} \|\widehat{\Delta}_j\|_2^2 - 32c_2 \rho^2 \frac{\log p}{n} \|\beta_{j,S^c}^*\|_1^2, \tag{41}
 \end{aligned}$$

where the last inequality comes from the condition $|S| \leq \frac{c_1}{64c_2} \frac{\bar{\kappa}}{\rho^2(\Psi^*)} \frac{n}{\log p}$. Now we analyze the following two cases regarding (41):

Case 1 Suppose that $\frac{c_1 \bar{\kappa}}{4} \|\widehat{\Delta}_j\|_2^2 \geq 32c_2 \rho^2 \frac{\log p}{n} \|\beta_{j,S^c}^*\|_1^2$, then from (39) we can get

$$\frac{c_1 \bar{\kappa}}{4} \|\widehat{\Delta}_j\|_2^2 \leq \lambda_n \left(3\sqrt{|S|} \|\widehat{\Delta}_j\|_2 + 4\|\beta_{j,S^c}^*\|_1 \right).$$

Solving for the zeros of this quadratic form in $\|\widehat{\Delta}_j\|_2$ yields

$$\|\widehat{\Delta}_j\|_2^2 \leq \frac{144\lambda_n^2}{c_1^2 \bar{\kappa}^2} |S| + \frac{32\lambda_n \|\beta_{j,S^c}^*\|_1}{c_1 \bar{\kappa}}.$$

Case 2 Otherwise, we have $\frac{c_1 \bar{\kappa}}{4} \|\widehat{\Delta}_j\|_2^2 \leq 32c_2 \rho^2 \frac{\log p}{n} \|\beta_{j,S^c}^*\|_1^2$.

After combining the two cases, we obtain the claimed bound in (30). To get the prediction bound in (31), we first show $\|\widehat{\Delta}_j\|_1 \leq 4\|\beta_j^*\|_1$. From basic inequality, we have

$$0 \leq \frac{1}{2n} \|\widehat{L}X\widehat{\Delta}_j\|_2^2 \leq \frac{\varepsilon_j^\top \widehat{\Theta}X\widehat{\Delta}_j}{n} + \lambda_n \left(\|\beta_j^*\|_1 - \|\widehat{\beta}_j\|_1 \right).$$

By Hölder's inequality and Theorem 15, with high probability, we have that

$$\left| \frac{\varepsilon_j^\top \widehat{\Theta}X\widehat{\Delta}_j}{n} \right| \leq \left\| \frac{X^\top \widehat{\Theta}\varepsilon_j}{n} \right\|_\infty \|\widehat{\Delta}_j\|_1 \leq \frac{\lambda_n}{2} \left(\|\beta_j^*\|_1 + \|\widehat{\beta}_j\|_1 \right).$$

Combine the two inequalities above, we get

$$0 \leq \frac{3\lambda_n}{2} \|\beta_j^*\|_1 - \frac{\lambda_n}{2} \|\hat{\beta}_j\|_1,$$

which implies $\|\hat{\beta}_j\|_1 \leq 3\|\beta_j^*\|_1$. Consequently, we have

$$\|\hat{\Delta}_j\|_1 \leq \|\hat{\beta}_j\|_1 + \|\beta_j^*\|_1 \leq 4\|\beta_j^*\|_1.$$

Return to the basic inequality, we have

$$\begin{aligned} \frac{\|\hat{L}X\hat{\Delta}_j\|_2^2}{2n} &\leq \frac{\lambda_n}{2} \|\hat{\Delta}_j\|_1 + \lambda_n \left(\|\beta_j^*\|_1 - \|\beta_j^*\|_1 + \|\hat{\Delta}_j\|_1 \right) \\ &\leq \frac{3}{2} \lambda_n \|\hat{\Delta}_j\|_1 \leq 6\lambda_n \|\beta_j^*\|_1. \end{aligned}$$

■

C.4 Proof of Lemma 19

Proof The proof follows a similar approach as the proof for Lemma 1 in Ravikumar et al. (2011). ■

C.5 Proof of Corollary 20

Proof A little calculation using Lemma 19 and Definition 18 shows that the corresponding inverse functions for data from the Gaussian DAG model (2) are:

$$\bar{\delta}_f(p; r) = 40 \sqrt{\frac{2 \log(r/4)}{p}}, \quad \text{and} \quad \bar{p}_f(\delta; r) = \frac{3200 \log(r/4)}{\delta^2}.$$

Setting $r = n^\tau$ yields the desired result. ■

C.6 Proof of Lemma 22

Proof Let $\hat{S}_{ij}^{(t)}$ denote the (i, j) entry of the sample variance matrix $\hat{S}^{(t)}$ defined in (10). Let $X_{i \cdot}$ and $X_{\cdot j}$ denote the i th row and j th column of X , respectively. Let $\varepsilon_{ik}^* := X_{ik} - X_{i \cdot} \beta_k^*$ where β_k^* is the k th column of B^* , $\rho_k^* = 1/\omega_k^{*2}$, $\hat{\rho}_k = 1/\hat{\omega}_k^2$, $\hat{\Delta}_k^{(t)} := \hat{\beta}_k^{(t)} - \beta_k^*$, and $\hat{\delta}_k = \hat{\rho}_k - \rho_k^*$.

Then,

$$\begin{aligned}
 \widehat{S}_{ij}^{(1)} &= \frac{1}{p} \sum_{k=1}^p \widehat{\rho}_k \left(X_{ik} - X_{i \cdot} \widehat{\beta}_k^{(1)} \right) \left(X_{jk} - X_{j \cdot} \widehat{\beta}_k^{(1)} \right) \\
 &= \frac{1}{p} \sum_{k=1}^p \widehat{\rho}_k \left(\varepsilon_{ik}^* - X_{i \cdot} \widehat{\Delta}_k^{(1)} \right) \left(\varepsilon_{jk}^* - X_{j \cdot} \widehat{\Delta}_k^{(1)} \right) \\
 &= \widehat{\Sigma}_{ij} + \frac{1}{p} \sum_{k=1}^p \rho_k^* \left(-\varepsilon_{ik}^* X_{j \cdot} \widehat{\Delta}_k^{(1)} - \varepsilon_{jk}^* X_{i \cdot} \widehat{\Delta}_k^{(1)} + X_{i \cdot} \widehat{\Delta}_k^{(1)} X_{j \cdot} \widehat{\Delta}_k^{(1)} \right) \\
 &\quad + \frac{1}{p} \sum_{k=1}^p \widehat{\delta}_k \left(\varepsilon_{ik}^* - X_{i \cdot} \widehat{\Delta}_k^{(1)} \right) \left(\varepsilon_{jk}^* - X_{j \cdot} \widehat{\Delta}_k^{(1)} \right) \\
 &= \widehat{\Sigma}_{ij} + \frac{1}{p} \sum_{k=1}^p \widehat{\rho}_k \left(-\varepsilon_{ik}^* X_{j \cdot} \widehat{\Delta}_k^{(1)} - \varepsilon_{jk}^* X_{i \cdot} \widehat{\Delta}_k^{(1)} + X_{i \cdot} \widehat{\Delta}_k^{(1)} X_{j \cdot} \widehat{\Delta}_k^{(1)} \right) + \frac{1}{p} \sum_{k=1}^p \widehat{\delta}_k \varepsilon_{ik}^* \varepsilon_{jk}^*.
 \end{aligned}$$

If we let $R_{ij} = \frac{1}{p} \sum_{k=1}^p \widehat{\rho}_k \left(-\varepsilon_{ik}^* X_{j \cdot} \widehat{\Delta}_k^{(1)} - \varepsilon_{jk}^* X_{i \cdot} \widehat{\Delta}_k^{(1)} + X_{i \cdot} \widehat{\Delta}_k^{(1)} X_{j \cdot} \widehat{\Delta}_k^{(1)} \right) + \frac{1}{p} \sum_{k=1}^p \widehat{\delta}_k \varepsilon_{ik}^* \varepsilon_{jk}^*$, we can upper bound $|R_{ij}|$ by dividing it into three terms and controlling each term separately.

Part 1.

Define the following events:

$$\begin{aligned}
 \mathcal{B}_1 &= \bigcup_{i=1}^n \left\{ \|\varepsilon_{i \cdot}^*\|_\infty \geq 6\bar{\omega} \sqrt{\log \max\{n, p\}} \right\}, \\
 \mathcal{B}_2 &= \bigcup_{k=1}^n \left\{ \|\varepsilon_{k \cdot}^*\|_2 \geq 6\bar{\omega} \sqrt{p} \right\}, \\
 \mathcal{B}_3 &= \bigcup_{k=1}^n \left\{ \|X_{k \cdot}\|_\infty \geq 6\bar{\psi} \sqrt{\log \max\{n, p\}} \right\}.
 \end{aligned}$$

Under event \mathcal{B}_1 , \mathcal{B}_2 , and \mathcal{B}_3 ,

$$\begin{aligned}
 \left| \frac{1}{p} \sum_{k=1}^p \widehat{\rho}_k \varepsilon_{ik}^* X_{j \cdot} \widehat{\Delta}_k^{(1)} \right| &\leq \frac{1}{b} \sup_k |\varepsilon_{ik}^* X_{j \cdot} \widehat{\Delta}_k^{(1)}| \quad (\text{By Lemma 7}) \\
 &\leq \frac{6\bar{\omega}}{b} \sqrt{\log \max\{n, p\}} \sup_k \|X_{j \cdot}\|_\infty \|\widehat{\Delta}_k^{(1)}\|_1 \quad (\text{By Hölder's Inequality and } \mathcal{B}_1) \\
 &\leq \frac{6\bar{\omega} s \sqrt{c \log p \log \max\{n, p\}}}{b \sqrt{n}} \|X_{j \cdot}\|_\infty \quad (\text{From } \|\widehat{\Delta}_k^{(1)}\|_1 \leq 4\sqrt{2s} \|\widehat{\Delta}_k^{(1)}\|_2) \\
 &\leq \frac{36\bar{\omega} \bar{\psi} s}{b} \sqrt{\frac{c \log p \log^2 \max\{n, p\}}{n}} \quad (\text{By event } \mathcal{B}_3). \\
 &= \frac{36\bar{\omega} \bar{\psi}}{b} \cdot \log \max\{n, p\} \cdot \sqrt{\frac{c s^2 \log p}{n}}. \tag{42}
 \end{aligned}$$

where $b > 0$ is the same constant defined in Lemma 7. The second last inequality comes from $\|\widehat{\Delta}^{(1)}\|_1 \leq 4\sqrt{2s} \|\widehat{\Delta}^{(1)}\|_2$ in the proof of Theorem 17.

Part 2

Notice that

$$\begin{aligned}
 \left| \frac{1}{p} \sum_{k=1}^p \hat{\rho}_k X_{i \cdot} \hat{\Delta}_k^{(1)} X_{j \cdot} \hat{\Delta}_k^{(1)} \right| &\leq \frac{2}{b} \sup_k |X_{j \cdot} \hat{\Delta}_k^{(1)}| |X_{i \cdot} \hat{\Delta}_k^{(1)}| \\
 &\leq \frac{2s^2 \log p}{bn} \|X_{j \cdot}\|_\infty \|X_{i \cdot}\|_\infty \\
 &\leq \frac{12\bar{\psi}s^2}{b} \sqrt{\frac{c \log^2 p \log^2 \max\{n, p\}}{n^2}} \\
 &= \frac{12\bar{\psi}}{b} \cdot \log \max\{n, p\} \cdot \sqrt{c \frac{s^2 \log p}{n}}. \tag{43}
 \end{aligned}$$

Since $\frac{s^2 \log p}{n} \lesssim 1$ from the assumption, E.q. (42) dominates E.q. (43).

Part 3

By definition, $\|\hat{\delta}\|_\infty = \sup_k |\hat{\rho}_k - \rho_k^*| = \sup_k |1/\hat{\omega}_k^2 - 1/\omega_k^{*2}| = r(\hat{\Omega})$. Combining with \mathcal{B}_2 , we have

$$\left| \frac{1}{p} \sum_{k=1}^p \hat{\delta}_k \varepsilon_{ik}^* \varepsilon_{jk}^* \right| \leq \frac{1}{p} \|\hat{\delta}\|_\infty \sum_{k=1}^p |\varepsilon_{ik}^* \varepsilon_{jk}^*| \leq \frac{1}{p} \|\delta\|_\infty \|\varepsilon_{i \cdot}\|_2 \|\varepsilon_{j \cdot}\|_2 \leq 6\bar{\omega} \|\hat{\delta}\|_\infty = 6\bar{\omega} r(\hat{\Omega}).$$

Combine all three parts, we have

$$|R_{ij}| \leq \max \left\{ 6\bar{\omega} r(\hat{\Omega}), \frac{144\bar{\omega}\bar{\psi}s}{b} \sqrt{\frac{c \log p \log^2 \max\{n, p\}}{n}} \right\}.$$

Using Lemma 14 and Lemma 12, we can derive the upper bound for the probabilities of $\mathcal{B}_1, \mathcal{B}_2, \mathcal{B}_3$:

$$\begin{aligned}
 \mathbb{P}(\mathcal{B}_1) &\leq 2/\max\{n, p\}^4, \\
 \mathbb{P}(\mathcal{B}_2) &\leq 1/\max\{n, p\}^4 \quad \text{if } n > 2\sqrt{10} \log \max\{n, p\}, \\
 \mathbb{P}(\mathcal{B}_3) &\leq 2/\max\{n, p\}^4, \\
 \mathbb{P}\left(\bigcup_{i=1}^3 \mathcal{B}_i\right) &\leq 5/\max\{n, p\}^4.
 \end{aligned}$$

Applying union bound,

$$\|R\|_\infty \leq \max \left\{ 6\bar{\omega} r(\hat{\Omega}), \frac{144\bar{\omega}\bar{\psi}s}{b} \sqrt{\frac{c \log p \log^2 \max\{n, p\}}{n}} \right\},$$

with probability at least $1 - \frac{5n^2}{\max\{n, p\}^4}$. Take event \mathcal{A} from Lemma 21 into account and apply union bound one more time, we arrive at the desired conclusion. \blacksquare

C.7 Proof of Theorem 23

Proof Let $\bar{R}(s, p, n)$ and $\bar{\delta}_f(p; n^\tau)$ be defined as stated, then the monotonicity of the inverse tail function (32) and condition (36) on (n, p) implies that $\bar{\delta}_f(p; n^\tau) \leq 1/40$. Lemma 21 and Lemma 22 imply that the event \mathcal{A} defined in (34) and the events $\mathcal{B}_1, \mathcal{B}_2, \mathcal{B}_3$ defined in the proof of Lemma 22 hold with high probability. Assume $\mathcal{A}, \mathcal{B}_1, \mathcal{B}_2, \mathcal{B}_3$ hold. Then,

$$\|\widehat{W}^{(1)}\|_\infty \leq \bar{\delta}_f(p; n^\tau) + \bar{R}(s, p, n).$$

Choose $\lambda_p = \bar{\delta}_f(p; n^\tau) + \bar{R}$. By Lemma 22, condition (36), and the monotonicity:

$$p > \bar{p}_f(\delta, r) \implies \bar{\delta}_f(p, r) \leq \delta \text{ for } \delta > 0,$$

we have that

$$2\kappa_{\Gamma^*} \left(\|\widehat{W}^{(1)}\|_\infty + \lambda_p \right) \leq 4\kappa_{\Gamma^*} (\bar{\delta}_f(p; n^\tau) + \bar{R}) \leq \min \left\{ \frac{1}{3\kappa_{\Sigma^*} m}, \frac{1}{3\kappa_{\Sigma^*}^3 2\kappa_{\Gamma^*} m} \right\}.$$

Applying Lemma 6 in Ravikumar et al. (2011) we obtain

$$\begin{aligned} \|\widehat{\Theta}^{(1)} - \Theta^*\|_\infty &\leq 4\kappa_{\Gamma^*} (\bar{\delta}_f(p; n^\tau) + \bar{R}), \\ \|\widehat{\Theta}^{(1)} - \Theta^*\|_2 &\leq \|A\|_2 \|\widehat{\Theta}^{(1)} - \Theta^*\|_\infty \leq (m+1) \|\widehat{\Theta}^{(1)} - \Theta^*\|_\infty. \end{aligned}$$

■

Appendix D. Proofs of Other Auxiliary Results

This section includes the proofs for Lemma 6 and 7 as well as the four lemmas introduced in Section A.1.

D.1 Proof of Lemma 6

Proof Notice that

$$\begin{aligned} \sup_j \left| \hat{\omega}_j^2 - \omega_j^{*2} \right| &= \sup_j \left| \hat{\omega}_j^2 - \frac{\|\varepsilon_j^{(B)}\|_2^2}{N} + \frac{\|\varepsilon_j^{(B)}\|_2^2}{N} - \omega_j^{*2} \right| \\ &\leq \sup_j \left| \hat{\omega}_j^2 - \frac{\|\varepsilon_j^{(B)}\|_2^2}{N} \right| + \sup_j \left| \frac{\|\varepsilon_j^{(B)}\|_2^2}{N} - \omega_j^{*2} \right|. \end{aligned}$$

We will show separately that with the desired probability:

$$\sup_j \left| \hat{\omega}_j^2 - \frac{\|\varepsilon_j^{(B)}\|_2^2}{N} \right| \leq \lambda_N s \bar{\beta}, \tag{44}$$

$$\sup_j \left| \frac{\|\varepsilon_j^{(B)}\|_2^2}{N} - \omega_j^{*2} \right| \leq 5\bar{\omega} \sqrt{\frac{\log 2 + \log p}{N}}. \tag{45}$$

From Lemma 1 in Yu and Bien (2019) we know that if $\lambda_N \geq N^{-1} \|X^{(B)\top} \varepsilon_j^{(B)}\|_\infty$, then

$$\left| \hat{\omega}_j^2 - N^{-1} \|\varepsilon_j^{(B)}\|_2^2 \right| \leq 2\lambda_N \|\beta_j^*\|_1 \leq \lambda_N s \bar{\beta}.$$

Pick $\lambda_N = 12\bar{\psi}\bar{\omega} \left(\sqrt{\frac{2\log p}{N}} + \sqrt{\frac{2\log 2+6\log p}{N}} \right)$, then applying the tail bound for maxima of sub-Gaussian random variables (e.g., see proof of Lemma 13) we can show that

$$\lambda_N \geq \sup_j N^{-1} \|X^{(B)\top} \varepsilon_j^{(B)}\|_\infty$$

holds with probability at least $(1 - \frac{1}{p})^2$. This proves the inequality in (44). To prove the second inequality, notice that, from χ^2 concentration inequality (e.g. Wainwright 2019 Example 2.11),

$$\begin{aligned} \left| \omega_j^{*2} - \frac{1}{N} \|\varepsilon_j^{(B)}\|_2^2 \right| &\geq 2\sqrt{2}\bar{\omega} \sqrt{\frac{\log 2 + 3\log p}{N}} \quad \text{with probability at most } 1/p^3, \\ \sup_j \left| \omega_j^{*2} - \frac{1}{N} \|\varepsilon_j^{(B)}\|_2^2 \right| &\geq 2\sqrt{2}\bar{\omega} \sqrt{\frac{\log 2 + 3\log p}{N}} \quad \text{with probability at most } 1/p^2. \end{aligned}$$

Combining all the inequalities, we can show that the second inequality holds with probability at least $(1 - 1/p)^2 - 1/p$. \blacksquare

D.2 Proof of Lemma 7

Proof Simply notice that

$$\sup_{1 \leq j \leq p} \left| \frac{1}{\hat{\omega}_j^2} - \frac{1}{\omega_j^{*2}} \right| = \sup_{1 \leq j \leq p} \left| \frac{1}{\hat{\omega}_j^2 \omega_j^{*2}} \right| \sup_{1 \leq j \leq p} \left| \omega_j^{*2} - \hat{\omega}_j^2 \right| \leq \frac{1}{b^4} \sup_{1 \leq j \leq p} |\omega_j^{*2} - \hat{\omega}_j^2|. \quad \blacksquare$$

D.3 Proof of Lemma 11

Proof The claim follows since

$$\begin{aligned} \|\hat{\Delta}_{prec}\|_2 &= \|\hat{\Theta} - \Theta^*\|_2 = \left\| \left(L^* + \hat{\Delta}_{chol} \right)^\top \left(L^* + \hat{\Delta}_{chol} \right) - L^{*\top} L^* \right\|_2 \\ &= \|L^{*\top} \hat{\Delta}_{chol} + \hat{\Delta}_{chol}^\top L^* + \hat{\Delta}_{chol}^\top \hat{\Delta}_{chol}\|_2 \\ &\geq \max_{x \in \mathbb{S}^{n-1}} x^\top \left(L^{*\top} \hat{\Delta}_{chol} + \hat{\Delta}_{chol}^\top L^* + \hat{\Delta}_{chol}^\top \hat{\Delta}_{chol} \right) x \\ &\geq \max_{x \in \mathbb{S}^{n-1}} x^\top \left(L^{*\top} \hat{\Delta}_{chol} + \hat{\Delta}_{chol}^\top L^* \right) x \quad (\text{as } \hat{\Delta}_{chol}^\top \hat{\Delta}_{chol} \geq 0.) \\ &= \|L^{*\top} \hat{\Delta}_{chol} + \hat{\Delta}_{chol}^\top L^*\|_2 \geq 2\sigma_{\min}(L^*) \|\hat{\Delta}_{chol}\|_2. \end{aligned} \quad \blacksquare$$

D.4 Proof of Lemma 12

Proof Notice that $\tilde{X}_k \in \mathbb{R}^n$ is a sub-Gaussian random vector with variance smaller than $\bar{\psi}$. By Theorem 1.19 in Rigollet (2015), we have that

$$\mathbb{P} \left(\|\tilde{X}_k\|_2 > 4\bar{\psi}\sqrt{n} + 2\bar{\psi}\sqrt{2\log(1/\delta)} \right) \leq \delta \quad \forall \delta > 0.$$

Setting $\delta = 1/p^\alpha$ and using union bound we obtain the desired conclusion. \blacksquare

D.5 Proof of Lemma 13

Proof Lemma 12 implies that with probability at least $1 - 1/p$,

$$\|\tilde{X}_k\|_2 \leq 4\bar{\psi}\sqrt{n} + 2\bar{\psi}\sqrt{2\log p} \leq 6\bar{\psi}\sqrt{n}, \quad (46)$$

for all k . Under the event \mathcal{E} defined in (25), $\|\tilde{X}_{[j-1]}^\top \tilde{\varepsilon}_j\|_\infty/n$ corresponds to the absolute maximum of $j-1$ zero-mean Gaussian variables, each with variance at most $36\bar{\psi}^2\bar{\omega}^2/n$. Next, we calculate the probability of the event $\mathcal{T} \cap \mathcal{E}$, where $\delta = 1/p^2$. We also let

$$\begin{aligned} t &= \sqrt{\frac{2\log 2 + 4\log p}{n}}, \\ \lambda_n &= 12\bar{\psi}\bar{\omega} \left(\sqrt{\frac{2\log p}{n}} + t \right). \end{aligned}$$

Because both \tilde{X} and $\tilde{\varepsilon}$ are random, we use the equivalence: $p(y) = \mathbb{E}_{p(x)}[p(y|x)]$ to apply the properties of fixed-design Lasso: Let $X_{[j-1]}$ denote the first $j-1$ columns in X ,

$$\begin{aligned} 1 - \mathbb{P}(\mathcal{T}_j \mid \mathcal{E}) &= \mathbb{E}_{X_{[j-1]}} \mathbb{P} \left\{ 2\|\tilde{X}_{[j-1]}^\top \tilde{\varepsilon}_j\|_\infty/n > \lambda_n \mid X_{[j-1]}, \mathcal{E} \right\} \\ &= \mathbb{E}_{X_{[j-1]}} \mathbb{P} \left\{ \|\tilde{X}_{[j-1]}^\top \tilde{\varepsilon}_j\|_\infty/n > 6\bar{\psi}\bar{\omega} \left(\sqrt{\frac{2\log p}{n}} + t \right) \mid X_{[j-1]}, \mathcal{E} \right\} \\ &\leq 2 \exp \left\{ -\frac{nt^2}{2} \right\} = 1/p^2, \end{aligned}$$

where in the last inequality we apply the tail bound on the maxima of sub-Gaussian random variables (Duchi, 2017): If $X_j \sim \text{subGaussian}(\sigma^2)$ are i.i.d., then

$$\mathbb{P} \left(\max_{1 \leq j \leq n} X_j \geq \sqrt{2\sigma^2 \log n} + t \right) \leq \exp \left(-\frac{t^2}{2\sigma^2} \right).$$

Hence,

$$1 - \mathbb{P}(\mathcal{T} \mid \mathcal{E}) = 1 - \mathbb{P} \left(\bigcap_{j=1}^p \mathcal{T}_j \mid \mathcal{E} \right) = \mathbb{P} \left(\bigcup_{j=1}^p \mathcal{T}_j^c \mid \mathcal{E} \right) \leq \frac{1}{p}. \quad (47)$$

Finally, by Lemma 12 with $\alpha = 2$ we get

$$\mathbb{P}(\mathcal{T}) \geq \mathbb{P}(\mathcal{E}) \mathbb{P}(\mathcal{T} \mid \mathcal{E}) \geq \left(1 - \frac{1}{p} \right)^2. \quad (48)$$

\blacksquare

D.6 Proof of Lemma 14

Proof By the sub-Gaussian maximal inequality (e.g., Theorem 1.14 in Rigollet 2015), we know that if X_1, \dots, X_N are random variables such that $X_i \sim$ sub-Gaussian with parameter σ^2 , then for any $t > 0$,

$$\mathbb{P} \left(\max_{1 \leq i \leq N} |X_i| \geq t \right) \leq 2N \exp \left(-\frac{t^2}{2\sigma^2} \right).$$

Letting $t = \sqrt{4\bar{\psi}^2 \log p}$ and taking $\sigma^2 = \bar{\psi}^2$, we arrive at the desired result. \blacksquare

References

Genevera I Allen and Robert Tibshirani. Transposable regularized covariance models with an application to missing data imputation. *The Annals of Applied Statistics*, 4(2):764, 2010.

Bryon Aragam and Qing Zhou. Concave penalized estimation of sparse gaussian Bayesian networks. *Journal of Machine Learning Research*, 16:2273–2328, 2015. URL <http://jmlr.org/papers/v16/aragam15a.html>.

Bryon Aragam, Arash Amini, and Qing Zhou. Globally optimal score-based learning of directed acyclic graphs in high-dimensions. In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019a.

Bryon Aragam, Jiaying Gu, and Qing Zhou. Learning large-scale Bayesian networks with the sparsebn package. *Journal of Statistical Software, Articles*, 91(11):1–38, 2019b. ISSN 1548-7660. doi: 10.18637/jss.v091.i11. URL <https://www.jstatsoft.org/v091/i11>.

Mohsen Bayati, Murat A Erdogdu, and Andrea Montanari. Estimating lasso risk and noise level. In C. J. C. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems*, volume 26. Curran Associates, Inc., 2013. URL <https://proceedings.neurips.cc/paper/2013/file/2b8a61594b1f4c4db0902a8a395ced93-Paper.pdf>.

A. Belloni, V. Chernozhukov, and L. Wang. Square-root lasso: pivotal recovery of sparse signals via conic programming. *Biometrika*, 98(4):791–806, 2011. ISSN 00063444, 14643510. URL <http://www.jstor.org/stable/23076172>.

Rohit Bhattacharya, Daniel Malinsky, and Ilya Shpitser. Causal inference under interference and network uncertainty. In *Uncertainty in Artificial Intelligence*, pages 1028–1038. PMLR, 2020.

Peter Buhlmann and Sara van de Geer. *Statistics for High-Dimensional Data: Methods, Theory and Applications*. Springer Publishing Company, Incorporated, 1st edition, 2011. ISBN 3642201911.

Xi Chen, Han Xu, Ping Yuan, Fang Fang, Mikael Huss, Vinsensius B Vega, Eleanor Wong, Yuriy L Orlov, Weiwei Zhang, Jianming Jiang, et al. Integration of external signaling pathways with the core transcriptional network in embryonic stem cells. *Cell*, 133(6):1106–1117, 2008.

David Maxwell Chickering. Optimal structure identification with greedy search. *Journal of Machine Learning Research*, 3:507–554, March 2003. ISSN 1532-4435. doi: 10.1162/153244303321897717. URL <https://doi.org/10.1162/153244303321897717>.

David Maxwell Chickering, David Heckerman, and Christopher Meek. Large-sample learning of Bayesian networks is np-hard. *Journal of Machine Learning Research*, 5(Oct):1287–1330, 2004.

Li-Fang Chu, Ning Leng, Jue Zhang, Zhonggang Hou, Daniel Mamott, David T. Vereide, Jeea Choi, Christina Kendziorski, Ron Stewart, and James A. Thomson. Single-cell rna-seq reveals novel regulators of human embryonic stem cell differentiation to definitive endoderm. *Genome Biology*, 17(1):173, 2016. doi: 10.1186/s13059-016-1033-x. URL <https://doi.org/10.1186/s13059-016-1033-x>.

Gregory F. Cooper and Edward Herskovits. A Bayesian method for the induction of probabilistic networks from data. *Machine Learning*, 9(4):309–347, Oct 1992. ISSN 1573-0565. doi: 10.1007/BF00994110. URL <https://doi.org/10.1007/BF00994110>.

John Duchi. Concentration inequalities and tail bounds. <https://web.stanford.edu/class/cs229t/2017/Lectures/concentration-slides.pdf>, 2017. [Online; accessed 04-August-2023].

Pierre Dutilleul. The MLE algorithm for the matrix normal distribution. *Journal of Statistical Computation and Simulation*, 64(2):105–123, 1999. doi: 10.1080/00949659908811970. URL <http://dx.doi.org/10.1080/00949659908811970>.

Gideon E. Schwarz. Estimating the dimension of a model. *The Annals of Statistics*, 6, 03 1978. doi: 10.1214/aos/1176344136.

Jianqing Fan, Shaojun Guo, and Ning Hao. Variance estimation using refitted cross-validation in ultrahigh dimensional regression. *Journal of the Royal Statistical Society. Series B (Statistical Methodology)*, 74(1):37–65, 2012. ISSN 13697412, 14679868. URL <http://www.jstor.org/stable/41430928>.

Fei Fu and Qing Zhou. Learning sparse causal gaussian networks with experimental intervention: regularization and coordinate descent. *Journal of the American Statistical Association*, 108(501):288–300, 2013.

Maxime Gasse, Alex Aussem, and Haytham Elghazel. A hybrid algorithm for Bayesian network structure learning with application to multi-label learning. *Expert Syst. Appl.*, 41(15):6755–6772, November 2014. ISSN 0957-4174. doi: 10.1016/j.eswa.2014.04.032. URL <https://doi.org/10.1016/j.eswa.2014.04.032>.

Jiaying Gu, Fei Fu, and Qing Zhou. Penalized estimation of directed acyclic graphs from discrete data. *Statistics and Computing*, 29(1):161–176, 2019.

Trevor Hastie, Robert Tibshirani, and Martin Wainwright. *Statistical Learning with Sparsity: The Lasso and Generalizations*. Chapman & Hall/CRC, 2015. ISBN 1498712169, 9781498712163.

David Heckerman, Dan Geiger, and David M. Chickering. Learning Bayesian networks: The combination of knowledge and statistical data. *Machine Learning*, 20(3):197–243, Sep 1995. ISSN 1573-0565. doi: 10.1023/A:1022623210503. URL <https://doi.org/10.1023/A:1022623210503>.

Michael G. Hudgens and M. Elizabeth Halloran. Toward causal inference with interference. *Journal of the American Statistical Association*, Jun 2008. doi: 10.1198/016214508000000292.

Markus Kalisch, Martin Mächler, Diego Colombo, Marloes H. Maathuis, and Peter Bühlmann. Causal inference using graphical models with the R package pcalg. *Journal of Statistical Software*, 47(11):1–26, 2012. URL <http://www.jstatsoft.org/v47/i11/>.

Suprateek Kundu and Benjamin B. Risk. Scalable Bayesian matrix normal graphical models for brain functional networks. *Biometrics*, n/a(n/a), 2020. doi: <https://doi.org/10.1111/biom.13319>. URL <https://onlinelibrary.wiley.com/doi/abs/10.1111/biom.13319>.

Steffen L. Lauritzen. *Graphical Models*. Oxford University Press, 1996. ISBN 0-19-852219-3.

Steffen L. Lauritzen and TS Richardson. Chain graph models and their causal interpretations. *Journal of the Royal Statistical Society, Series B (Statistical Methodology)*, 64:321–348, 2002. ISSN 1369-7412. doi: 10.1111/1467-9868.00340.

Wei Vivian Li and Jingyi Jessica Li. An accurate and robust imputation method scimpute for single-cell rna-seq data. *Nature Communications*, 9(1):997, 2018. doi: 10.1038/s41467-018-03405-7. URL <https://doi.org/10.1038/s41467-018-03405-7>.

Teppo Niinimäki, Pekka Parviainen, and Mikko Koivisto. Structure discovery in Bayesian networks by sampling partial orders. *The Journal of Machine Learning Research*, 17(1):2002–2048, 2016.

Tore Opsahl. *Structure and Evolution of Weighted Networks*. University of London (Queen Mary College), London, UK, 2009. URL <http://toreopsahl.com/publications/thesis/>.

Tore Opsahl and Pietro Panzarasa. Clustering in weighted networks. *Social Networks*, 31(2):155 – 163, 2009. ISSN 0378-8733. doi: <https://doi.org/10.1016/j.socnet.2009.02.002>. URL <http://www.sciencedirect.com/science/article/pii/S0378873309000070>.

Joseph Ramsey, Madelyn Glymour, Ruben Sanchez-Romero, and Clark Glymour. A million variables and more: the fast greedy equivalence search algorithm for learning high-dimensional graphical causal models, with an application to functional magnetic resonance images. *International Journal of Data Science and Analytics*, 3(2):121–129, Mar 2017. ISSN 2364-4168. doi: 10.1007/s41060-016-0032-z. URL <https://doi.org/10.1007/s41060-016-0032-z>.

Pradeep Ravikumar, Martin J. Wainwright, Garvesh Raskutti, and Bin Yu. High-dimensional covariance estimation by minimizing l_1 -penalized log-determinant divergence. *Electronic Journal of Statistics*, 5(none):935 – 980, 2011. doi: 10.1214/11-EJS631. URL <https://doi.org/10.1214/11-EJS631>.

Philippe Rigollet. High dimensional statistics lecture notes. https://ocw.mit.edu/courses/mathematics/18-s997-high-dimensional-statistics-spring-2015/lecture-notes/MIT18_S997S15_CourseNotes.pdf, 2015. [Online; accessed 10-December-2020].

Teemu Roos. *Minimum Description Length Principle*, pages 823–827. Springer US, Boston, MA, 2017. ISBN 978-1-4899-7687-1. doi: 10.1007/978-1-4899-7687-1_894. URL https://doi.org/10.1007/978-1-4899-7687-1_894.

Mauro Scanagatta, Giorgio Corani, Cassio P de Campos, and Marco Zaffalon. Learning treewidth-bounded Bayesian networks with thousands of variables. In D. Lee, M. Sugiyama, U. Luxburg, I. Guyon, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 29. Curran Associates, Inc., 2016. URL <https://proceedings.neurips.cc/paper/2016/file/e2a2dcc36a08a345332c751b2f2e476c-Paper.pdf>.

Mark Schmidt, Alexandru Niculescu-Mizil, and Kevin Murphy. Learning graphical model structure using l_1 -regularization paths. In *Proceedings of the 22Nd National Conference on Artificial Intelligence - Volume 2*, AAAI'07, pages 1278–1283. AAAI Press, 2007. ISBN 978-1-57735-323-2. URL <http://dl.acm.org/citation.cfm?id=1619797.1619850>.

Marco Scutari. Learning Bayesian networks with the bnlearn r package. *Journal of Statistical Software, Articles*, 35(3):1–22, 2010. ISSN 1548-7660. doi: 10.18637/jss.v035.i03. URL <https://www.jstatsoft.org/v035/i03>.

Ilya Shpitser. Segregated graphs and marginals of chain graph models. In C. Cortes, N. Lawrence, D. Lee, M. Sugiyama, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 28. Curran Associates, Inc., 2015. URL https://proceedings.neurips.cc/paper_files/paper/2015/file/9ac403da7947a183884c18a67d3aa8de-Paper.pdf.

Peter Spirtes, Clark N Glymour, Richard Scheines, David Heckerman, Christopher Meek, Gregory Cooper, and Thomas Richardson. *Causation, Prediction, and Search*. MIT press, 2000.

Tingni Sun and Cun-Hui Zhang. Scaled sparse linear regression. *Biometrika*, 99(4):879–898, 2012. ISSN 00063444, 14643510. URL <http://www.jstor.org/stable/41720740>.

Robert Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society: Series B (Methodological)*, 58(1):267–288, 1996.

Ioannis Tsamardinos, Laura E. Brown, and Constantin F. Aliferis. The max-min hill-climbing Bayesian network structure learning algorithm. *Mach. Learn.*, 65(1):31–78, October 2006. ISSN 0885-6125. doi: 10.1007/s10994-006-6889-7. URL <https://doi.org/10.1007/s10994-006-6889-7>.

P. Tseng. Convergence of a block coordinate descent method for nondifferentiable minimization. *Journal of Optimization Theory and Applications*, 109(3):475–494, Jun 2001. ISSN 1573-2878. doi: 10.1023/A:1017501703105. URL <https://doi.org/10.1023/A:1017501703105>.

Theodoros Tsiligkaridis, Alfred O Hero III, and Shuheng Zhou. On convergence of kronecker graphical lasso algorithms. *IEEE transactions on signal processing*, 61(7):1743–1755, 2013.

Sara van de Geer and Peter Bühlmann. ℓ_0 -penalized maximum likelihood for sparse directed acyclic graphs. *The Annals of Statistics*, 41(2):536 – 567, 2013. doi: 10.1214/13-AOS1085. URL <https://doi.org/10.1214/13-AOS1085>.

Sara A. van de Geer and Peter Bühlmann. On the conditions used to prove oracle results for the Lasso. *Electronic Journal of Statistics*, 3(none):1360 – 1392, 2009. doi: 10.1214/09-EJS506. URL <https://doi.org/10.1214/09-EJS506>.

Martin J. Wainwright. *High-Dimensional Statistics: A Non-Asymptotic Viewpoint*. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, 2019. doi: 10.1017/9781108627771.

Duncan J Watts and Steven H Strogatz. Collective dynamics of ‘small-world’networks. *nature*, 393(6684):440, 1998.

Qiaoling Ye, Arash Amini, and Qing Zhou. Optimizing regularized cholesky score for order-based learning of Bayesian networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 43:3555–3572, 2021. doi: 10.1109/TPAMI.2020.2990820.

Guo Yu and Jacob Bien. Estimating the error variance in a high-dimensional linear model. *Biometrika*, 106(3):533–546, 05 2019. ISSN 0006-3444. doi: 10.1093/biomet/asz017. URL <https://doi.org/10.1093/biomet/asz017>.

Xun Zheng, Bryon Aragam, Pradeep Ravikumar, and Eric P. Xing. DAGs with NO TEARS: Continuous Optimization for Structure Learning. In *Advances in Neural Information Processing Systems*, 2018.

Qing Zhou, Hiram Chipperfield, Douglas A. Melton, and Wing Hung Wong. A gene regulatory network in mouse embryonic stem cells. *Proceedings of the National Academy of Sciences*, 104(42):16438–16443, 2007. ISSN 0027-8424. doi: 10.1073/pnas.0701014104. URL <https://www.pnas.org/content/104/42/16438>.

Shuheng Zhou. Gemini: Graph estimation with matrix variate normal instances. *The Annals of Statistics*, 42(2):532–562, 2014.