

High-Dimensional Poisson Structural Equation Model Learning via ℓ_1 -Regularized Regression

Gunwoong Park

*Department of Statistics
University of Seoul
Seoul, 02504, South Korea*

GW.PARK23@GMAIL.COM

Sion Park

*Department of Statistics
University of Seoul
Seoul, 02504, South Korea*

ROCKGOAT95@GMAIL.COM

Editor: Qiang Liu

Abstract

In this paper, we develop a new approach to learning high-dimensional Poisson structural equation models from only observational data without strong assumptions such as faithfulness and a sparse moralized graph. A key component of our method is to decouple the ordering estimation or parent search where the problems can be efficiently addressed using ℓ_1 -regularized regression and the moments relation. We show that sample size $n = \Omega(d^2 \log^9 p)$ is sufficient for our polynomial time Moments Ratio Scoring (MRS) algorithm to recover the true directed graph, where p is the number of nodes and d is the maximum indegree. We verify through simulations that our algorithm is statistically consistent in the high-dimensional $p > n$ setting, and performs well compared to state-of-the-art ODS, GES, and MMHC algorithms. We also demonstrate through multivariate real count data that our MRS algorithm is well-suited to estimating DAG models for multivariate count data in comparison to other methods used for discrete data.

Keywords: Bayesian Networks, Directed Acyclic Graph, Identifiability, Structure Learning, ℓ_1 -Regularization, Multivariate Count Distribution

1. Introduction

Directed acyclic graphical (DAG) models, also referred to as Bayesian networks, are popular probabilistic statistical models to analyze and visualize (functional) causal or directional dependence relationships among random variables. (see e.g., Kephart and White, 1991; Friedman et al., 2000; Doya, 2007; Peters and Bühlmann, 2014). However, learning DAG models from only observational data is a notoriously difficult problem due to non-identifiability and exponentially growing computational complexity. Prior works have addressed the question of identifiability for different classes of joint distribution $P(G)$. Frydenberg (1990) and Heckerman et al. (1995) show the Markov equivalence class (MEC) where graphs that belong to the same MEC have the same conditional independence relations. Spirtes et al. (2000), Chickering (2003), Tsamardinos and Aliferis (2003) and Zhang and Spirtes (2016) show that the underlying graph of a DAG model is recoverable up to the MEC under faithfulness or related assumptions that can be very restrictive (Uhler et al., 2013).

Also well studied is how learning a DAG model is computationally non-trivial due to the super-exponentially growing size of the space of DAGs in the number of nodes. Hence, it is NP-hard to search DAG space (Chickering et al., 1994; Chickering, 1996), and many existing algorithms such as PC (Spirtes et al., 2000), Greedy Equivalence Search (GES) (Chickering, 2003), Min-Max Hill Climbing (MMHC) (Tsamardinos et al., 2006) and Greedy DAG Search (GDS) (Peters and Bühlmann, 2014), take greedy search methods that may not guarantee to recover the true MEC.

Recently, a number of fully identifiable classes of DAG models have been introduced (Shimizu et al., 2006; Hoyer et al., 2009; Peters et al., 2011; Peters and Bühlmann, 2014; Park and Raskutti, 2015, 2018; Ghoshal and Honorio, 2018; Park and Park, 2019). In addition, some of these models can be successfully learned from high-dimensional data by decomposing the DAG learning problem into ordering estimation and skeleton estimation (Shimizu et al., 2011; Bühlmann et al., 2014; Ghoshal and Honorio, 2017b; Drton et al., 2018). The main reasoning is that if ordering is known or recoverable, learning a directed graphical model is as hard as learning an undirected graphical model or Markov random field (MRF). Meinshausen and Bühlmann (2006), Wainwright et al. (2006), Ravikumar et al. (2011) and Yang et al. (2015) show that sparse undirected graphs can be estimated via ℓ_1 -regularized regression in high-dimensional settings under suitable conditions.

In this paper, we focus on learning Poisson DAG models (Park and Raskutti, 2015, 2018) for multivariate count data in high-dimensional settings since large-scale multivariate *count data* frequently arises in many fields, such as high-throughput genomic sequencing data, spatial incidence data, sports science data, and disease incidence data. Like learning the Poisson undirected graphical model or MRF introduced in Yang et al. (2015), where the sample bound is $\Omega(d_m^2 \log^3 p)$, it is not surprising that Poisson DAG models can be learned in high dimensional settings when the indegree of the graph d is bounded. Park and Raskutti (2018) establishes the consistency of learning Poisson DAG models with the sample bound $n = \Omega(\max\{d_m^4 \log^{12} p, \log^{5+d} p\})$ where d_m is the maximum degree of the moralized graph and d is the maximum indegree of a graph. This huge sample complexity difference between directed and undirected graphical models is induced mainly for three reasons: (i) nonexistence ordering, (ii) the known parametric functional form (the standard log link) for the dependencies, and (iii) the restrictive non-positive parameter space in Poisson MRFs (see details in Yang et al., 2015).

The main objective of this work is to propose a new milder identifiability assumption for Poisson DAG models, and to develop a new polynomial time approach, called Moments Ratio Scoring (MRS), for learning a high-dimensional Poisson structural equation models (SEM), that is a Poisson DAG model where the parametric functional form for the dependencies is known while the parameters are unbounded and unknown. We address the question of learning high-dimensional Poisson SEMs under the causal sufficiency assumption that all relevant variables have been observed. However, we do not require the sparse moralized graph and faithfulness assumption that might be restrictive (Uhler et al., 2013).

The MRS algorithm combines the idea of the mean-variance (moments) relation for recovering an ordering, and the sparsity-encouraging ℓ_1 -regularized regression in finding the parents of each node. We provide its sufficient conditions and sample complexity $n = \Omega(d^2 \log^9 p)$ under which the MRS algorithm recovers the Poisson SEM with a high probability in the high-dimensional $p > n$ setting. The sample complexity of $n = \Omega(d^2 \log^9 p)$ is

close to the information-theoretic limit of $\Omega(d \log p)$ for learning sparse DAG models with any exponential family distributions (Ghoshal and Honorio, 2017a). We point out that the sample complexity does not depend on the maximum degree of the moralized graph, d_m , but on the indegree of a DAG, d . Since a sparse directed graph does not necessarily lead to the sparse moralized graph (e.g., a star graph in Fig. 2), to the best of our knowledge, the proposed algorithm is the most efficient and probable for learning sparse Poisson SEMs.

We demonstrate through simulations and a real baseball data application involving multivariate count data that our MRS algorithm performs better than state-of-the-art OverDispersion Scoring (ODS) (Park and Raskutti, 2015), GES (Chickering, 2003), MMHC (Tsamardinos et al., 2006), and Poisson MRF learning (PMRF) algorithms (Yang et al., 2015), on average, in terms of the both run-time and accuracy of recovering a graph structure and its MEC. In our simulation study, we consider both the extremely sparse ($d = 1$) and sparse ($d = 10$) high-dimensional settings. Our real data example involving MLB player statistics for 2003 season shows that our MRS algorithm is applicable to multivariate count data while the PMRF algorithm finds too many edges, and the MMHC algorithm tends to select very few edges when variables represent counts. We also investigate the accuracy of our MRS algorithm when samples are generated from general Poisson DAG models and (truncated) Poisson MRFs. The simulation results empirically verify that the MRS algorithm can consistently recover the true edges.

1.1. Our Contributions

We summarize the major contributions of the paper as follows:

- We introduce a milder identifiability condition for Poisson DAG models for multivariate count data.
- We develop the reliable and scalable lasso-based MRS algorithm which learns sparse high-dimensional Poisson SEMs.
- We provide the more realistic conditions for learning Poisson SEMs in Section 3.2.
- We also provide the sample complexity $n = \Omega(d^2 \log^9 p)$ under which the MRS algorithm recovers the Poisson SEM. We emphasize that our theoretical result does not depend on the degree of the moralized graph d_m , and hence, the MRS algorithm can recover a graph with hub nodes in the high dimensional setting.

To the best of our knowledge, our MRS algorithm is the only provable and realistic method that applies for the high-dimensional multivariate count data when samples are from Poisson SEMs with hub nodes. We must point out that such improved assumptions and sample complexity are not only from our new identifiability condition, but from the additional constraints on the standard log link function for the dependencies.

The remainder of this paper is structured as follows. Section 2.1 summarizes the necessary notations and problem settings, Section 2.2 discusses the Poisson DAG model and its new identifiability condition, and Section 2.3 provides a detailed comparison between Poisson DAG models and MRFs. In Section 3, we introduce our polynomial-time DAG learning algorithm, which we refer to as the Moments Ratio Scoring (MRS). Section 3.1 discusses

computational complexity of our algorithm, and Section 3.2 provides statistical guarantees for learning Poisson SEMs via the MRS algorithm. Section 4 empirically evaluates our methods, compared to state-of-the-art ODS, GES, and MMHC algorithms using synthetic data, and confirms that our algorithm is one of the few DAG-learning algorithms that performs well in terms of statistical and computational complexity in low and high-dimensional settings. In addition, we investigate how well the MRS algorithm learns general Poisson DAG models and (truncated) Poisson MRFs using synthetic data. Section 5 compares our MRS algorithm to the Poisson MRF and MMHC algorithm by analyzing a real 2003 season MLB multivariate count data. Lastly, Section 6 discusses some future works.

2. Poisson DAG Models

We first introduce some necessary notations and definitions for DAG models. Then, we give a detailed description of previous work on learning Poisson DAG models (Park and Raskutti, 2015), and we propose a strictly milder identifiability condition. Lastly, we discuss how Poisson DAG models and MRFs (Yang et al., 2015) are related.

2.1. Problem Set-up and Notations

A DAG $G = (V, E)$ consists of a set of nodes $V = \{1, 2, \dots, p\}$ and a set of directed edges $E \subset V \times V$ with no directed cycles. A directed edge from node j to k is denoted by (j, k) or $j \rightarrow k$. The set of *parents* of node k , denoted by $\text{Pa}(k)$, consists of all nodes j such that $(j, k) \in E$. If there is a directed path $j \rightarrow \dots \rightarrow k$, then k is called a *descendant* of j , and j is an *ancestor* of k . The set $\text{De}(k)$ denotes the set of all descendants of node k . The *non-descendants* of node k are $\text{Nd}(k) := V \setminus (\{k\} \cup \text{De}(k))$. An important property of DAGs is that there exists a (possibly non-unique) *ordering* $\pi = (\pi_1, \dots, \pi_p)$ of a directed graph that represents directions of edges such that for every directed edge $(j, k) \in E$, j comes before k in the ordering. Hence, learning a graph is equivalent to learning the ordering and the skeleton that is the set of directed edges without their directions.

We consider a set of random variables $X := (X_j)_{j \in V}$ with a probability distribution taking values in a sample space \mathcal{X}_V over the nodes in G . Suppose that a random vector X has a joint probability density function $P(G) = P(X_1, X_2, \dots, X_p)$. For any subset S of V , let $X_S := \{X_j : j \in S \subset V\}$ and $\mathcal{X}_S := \times_{j \in S} \mathcal{X}_j$ where \mathcal{X}_j is a sample space of X_j . For any node $j \in V$, $P(X_j | X_S)$ denotes the conditional distribution of a variable X_j given a random vector X_S . Then, a DAG model has the following factorization (Lauritzen, 1996):

$$P(G) = P(X_1, X_2, \dots, X_p) = \prod_{j=1}^p P(X_j | X_{\text{Pa}(j)}), \quad (1)$$

where $P(X_j | X_{\text{Pa}(j)})$ is the conditional distribution of X_j given its parents variables $X_{\text{Pa}(j)} := \{X_k : k \in \text{Pa}(j) \subset V\}$.

We suppose that there are n independent and identically distributed samples $X^{1:n} := (X^{(i)})_{i=1}^n$ from a given graphical model where $X^{(i)} := X_{1:p}^{(i)} = (X_1^{(i)}, X_2^{(i)}, \dots, X_p^{(i)})$ is a p -variate random vector. The notation $\hat{\cdot}$ denotes an estimate based on samples $X^{1:n}$. We also accept the causal sufficiency assumption that all important variables have been observed.

2.2. Poisson DAG Model and its Identifiability

The definition of Poisson DAG models in Park and Raskutti (2015) is that each conditional distribution given its parents $X_j | X_{\text{Pa}(j)}$ is Poisson such that

$$X_j | X_{\text{Pa}(j)} \sim \text{Poisson}(g_j(X_{\text{Pa}(j)})), \quad (2)$$

where for any arbitrary positive link function $g_j : \mathcal{X}_{\text{Pa}(j)} \rightarrow \mathbb{R}^+$. Hence using the factorization in Equation (1), the joint distribution is as follows:

$$f_G(X) = \prod_{j \in V} f_j(X_j | X_{\text{Pa}(j)}). \quad (3)$$

where f_j is the probability density function of Poisson.

A Poisson structural equation model (SEM) is a special case of a Poisson DAG model where the link functions g_j 's in Equation (2) are the standard log link function for Poisson generalized linear models (GLMs), i.e., $g_j(X_{\text{Pa}(j)}) = \exp(\theta_j + \sum_{k \in \text{Pa}(j)} \theta_{jk} X_k)$ where $(\theta_{jk})_{k \in \text{Pa}(j)}$ represents the linear weights. Using factorization (1), the joint distribution of a Poisson SEM can be written as:

$$f(X_1, X_2, \dots, X_p) = \exp\left(\sum_{j \in V} \theta_j X_j + \sum_{(k,j) \in E} \theta_{jk} X_j X_k - \sum_{j \in V} \log X_j! - \sum_{j \in V} e^{\theta_j + \sum_{k \in \text{Pa}(j)} \theta_{jk} X_k}\right). \quad (4)$$

Poisson DAG models have a useful moments relation for the identifiability:

Proposition 1 *Consider a Poisson DAG model (3) with non-degenerated rate parameter functions $(g_j(X_{\text{Pa}(j)}))_{j \in V}$. Then, for any node $j \in V$, and any set $S_j \subset \text{Nd}(j)$, the following moments relation holds:*

$$\frac{\mathbb{E}(X_j^2)}{\mathbb{E}[\mathbb{E}(X_j | X_{S_j}) + \mathbb{E}(X_j | X_{S_j})^2]} \geq 1 \quad (5)$$

Equivalently,

$$\mathbb{E}(\text{Var}(\mathbb{E}(X_j | X_{\text{Pa}(j)}) | X_{S_j})) \geq 0.$$

The equality only holds when S_j contains all parents of j , that is, $\text{Pa}(j) \subset S_j$.

We include the proof in Section A. Proposition 1 claims that when all parents of j , $\text{Pa}(j)$, contribute to its rate parameter, the moments ratio in Equation (5) is equal to 1 if a condition set S_j contains all parents of j , $\text{Pa}(j) \subset S_j$, otherwise greater than 1. In Poisson SEMs, it is clear that the non-degenerated rate parameter function assumptions are equivalent to the *non-zero coefficients* conditions, $|\theta_{jk}| > 0$ for all $k \in \text{Pa}(j)$ since $g_j(X_{\text{Pa}(j)}) = \exp(\theta_j + \sum_{k \in \text{Pa}(j)} \theta_{jk} X_k)$.

Now, we briefly explain how Poisson DAG models are identifiable from the moments ratio in Proposition 1 using the bivariate Poisson DAG models illustrated in Fig. 1: $G_1 : X_1 \sim \text{Poisson}(\lambda_1), X_2 \sim \text{Poisson}(\lambda_2)$, where X_1 and X_2 are independent; $G_2 : X_1 \sim \text{Poisson}(\lambda_1)$ and $X_2 | X_1 \sim \text{Poisson}(g_2(X_1))$; and $G_3 : X_2 \sim \text{Poisson}(\lambda_2)$ and $X_1 | X_2 \sim \text{Poisson}(g_1(X_2))$ for arbitrary non-degenerated positive functions $g_1, g_2 : \mathbb{N} \cup \{0\} \rightarrow \mathbb{R}^+$.

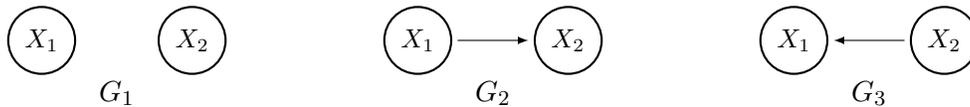


Figure 1: Bivariate directed acyclic graphs of G_1 , G_2 , and G_3 .

By Proposition 1, we can see that $\mathbb{E}(X_j^2) = \mathbb{E}(X_j) + \mathbb{E}(X_j)^2$ for all $j \in \{1, 2\}$ in G_1 . In G_2 , we can also see that

$$\mathbb{E}(X_1^2) = \mathbb{E}(X_1) + \mathbb{E}(X_1)^2, \quad \text{and} \quad \mathbb{E}(X_2^2) > \mathbb{E}(X_2) + \mathbb{E}(X_2)^2.$$

Similarly, in G_3 , we have $\mathbb{E}(X_1^2) > \mathbb{E}(X_1) + \mathbb{E}(X_1)^2$, while $\mathbb{E}(X_2^2) = \mathbb{E}(X_2) + \mathbb{E}(X_2)^2$. Hence, we can determine the true graph based on the moments ratio $\mathbb{E}(X_j^2)/(\mathbb{E}(X_j) + \mathbb{E}(X_j)^2)$.

This idea of a moments relation in Proposition 1 can easily apply to general p-variate Poisson DAG models, and hence, the models are identifiable by testing whether the moments ratio in Equation (5) is equal to 1 or greater than 1.

Theorem 2 *Consider a Poisson DAG model (3) with rate parameters $(g_j(X_{\text{Pa}(j)}))_{j \in V}$. If for any $j \in V$, rate parameter $g_j(\cdot)$ is non-degenerated, the Poisson DAG model is identifiable.*

We include the proof in Section 3.2. Theorem 2 claims that any Poisson DAG model is identifiable if all parents of node j contribute to its rate parameter. Hence, Theorem 2 shows that any Poisson SEM is identifiable under the non-zero coefficients condition, $|\theta_{jk}| > 0$ for all $k \in \text{Pa}(j)$. This condition is also commonly assumed in (Gaussian) linear structural equation models for the model identifiability (Spirtes, 1995; Ghoshal and Honorio, 2017b, 2018; Loh and Bühlmann, 2014; Peters and Bühlmann, 2014; Park and Park, 2019). We believe that it is a natural condition that is in accordance with the intuitive understanding of relationships among variables.

Our identifiability condition is strictly milder than the previous identifiability result in Park and Raskutti (2015) that is equivalent to $\text{Var}(\mathbb{E}(X_j | X_{\text{Pa}(j)} | X_{S_j} = x)) > 0$ for all $x \in \mathcal{X}_{S_j}$ when $\text{Pa}(j) \not\subset S_j$. For a better comparison, we consider a fully connected graph where $X_1 \sim \text{Poisson}(\lambda)$, $X_2 | X_1 \sim \text{Poisson}(\lambda + X_1)$, and $X_3 | X_1, X_2 \sim \text{Poisson}(\lambda + X_2 \mathbf{1}(X_1 \neq 0))$ where λ is a positive constant and $\mathbf{1}(\cdot)$ is an indicator function. In this case, we can see $\text{Var}(\mathbb{E}(X_3 | X_1, X_2) | X_1 = 0) = 0$, and hence, the identifiability condition in Park and Raskutti (2015) is not satisfied, while our condition is satisfied.

In a Poisson SEM, the identifiability assumption in Park and Raskutti (2015) is also satisfied under the non-zero coefficients condition. However, in the finite sample setting, the difference of both assumptions gets more crucial. For a positive constant c , Park and Raskutti (2015) requires $\min_{x \in \mathcal{X}_{S_j}} \text{Var}(\mathbb{E}(X_j | X_{\text{Pa}(j)} | X_{S_j} = x)) > c$, while we need $\mathbb{E}(\text{Var}(\mathbb{E}(X_j | X_{\text{Pa}(j)} | X_{S_j}))) > c$. Hence, our new identifiability assumption makes learning Poisson SEMs easier. We discuss this more in Section 3.2.

2.3. Comparison to Poisson MRF

In this section, we compare Poisson DAG models and MRFs where the conditional distributions of each node given its parents and neighbors are Poisson, respectively. To simplify

the comparison, we consider the joint distribution of a Poisson SEM in Equation (4). This is a form similar to the joint distribution of Poisson MRFs in Yang et al. (2015), where the joint distribution has the following form:

$$f(X_1, X_2, \dots, X_p) = \exp\left(\sum_{j \in V} \theta_j X_j + \sum_{(k,j) \in E} \theta_{jk} X_j X_k - \sum_{j \in V} \log X_j! - A(\theta)\right), \quad (6)$$

where $A(\theta)$ is the log of the normalization constant. The key difference between a Poisson SEM and a Poisson MRF is the normalization constant $A(\theta)$ in Equation (6), as opposed to the term $\sum_{j \in V} e^{\theta_j + \sum_{k \in \text{Pa}(j)} \theta_{jk} X_k}$ in Equation (4), which depends on variables.

Yang et al. (2015) proves that a Poisson MRF (6) is normalizable if and only if all (θ_{jk}) values are less than or equal to 0. This means Poisson MRFs only capture negative dependency relations. In addition, Yang et al. (2015) addresses the learning Poisson MRFs when the functional form of dependencies is $X_j | X_{V \setminus j} \sim \text{Poisson}(\exp(\theta_j + \sum_{k \in \mathcal{N}(j)} \theta_{jk} X_k))$ where $\mathcal{N}(j)$ denotes the neighbors of a node j in the graph.

While Poisson MRFs have strong restrictions on the functional form for dependencies and the parameter space, they can be successfully learned in the high-dimensional settings with less restrictive constraints of sparsity. Yang et al. (2015) shows that Poisson MRFs can be recovered via ℓ_1 -regularized regression if $n = \Omega(d_m^2 \log^3 p)$, where d_m is the degree of the undirected graph. In contrast, Park and Raskutti (2018) shows that Poisson DAG models can be learned via the ODS algorithm if $n = \Omega(\max\{d_m^4 \log^{12} p, \log^{5+d} p\})$ where d_m is obtained by the moralized graph and d is the maximum indegree of the graph. This big difference in the sample complexity primarily comes from the unknown functional form for the dependencies in Poisson DAG models. In the next section, we will show that a significant advantage can be achieved by assuming the parametric function for the dependencies in terms of recovering the graphs.

3. Algorithm

Here, we present our Moments Ratio Scoring (MRS) algorithm for learning the identifiable Poisson SEM (4). Our algorithm alternates between an element-wise ordering search using the (conditional) moments ratio, and a parent search using ℓ_1 -regularized GLM. Hence, the algorithm chooses a node for the first element of the ordering, and then determines its parents. The algorithm iterates this procedure until the last element of the ordering and its parents are determined.

Without loss of generality, assume that $\pi = (1, 2, \dots, p)$ is the true ordering. Then Poisson SEMs (4) have the conditional distribution of X_j given that all variables before j in the ordering are reduced to the following Poisson GLM:

$$P(X_j | X_{1:(j-1)}) = \exp\left\{\theta_j X_j + \sum_{k \in 1:(j-1)} \theta_{jk} X_k X_j + \log X_j! - \exp\left(\theta_j + \sum_{k \in 1:(j-1)} \theta_{jk} X_k\right)\right\}, \quad (7)$$

where $\theta_{jk} \in \mathbb{R}$ represents the influence of node k on node j . For ease of notation, let $\theta(j)$ be a set of parameters related to Poisson GLM (7). Then $\theta(j) = (\theta_j, \theta_{\setminus j}) \in \mathbb{R} \times \mathbb{R}^{j-1}$ where $\theta_{\setminus j} = (\theta_{jk})_{k \in \{1, 2, \dots, j-1\}}$ is a zero-padded vector with non-zero entries if $k \in \text{Pa}(j)$.

Our MRS (Algorithm 1) involves learning the ordering by comparing moments ratio scores of nodes using the following equations:

$$\widehat{S}(1, j) := \frac{\widehat{\mathbb{E}}(X_j^2)}{\widehat{\mathbb{E}}(X_j) + \widehat{\mathbb{E}}(X_j)^2} \quad \text{and} \quad \widehat{S}(m, j) := \frac{\widehat{\mathbb{E}}(X_j^2)}{\widehat{\mathbb{E}}(\widehat{\mathbb{E}}(X_j | X_{\widehat{\pi}_{1:(m-1)}}) + \widehat{\mathbb{E}}(X_j | X_{\widehat{\pi}_{1:(m-1)}})^2)}, \quad (8)$$

where $\widehat{\pi}_{1:m} = \{\widehat{\pi}_1, \dots, \widehat{\pi}_m\}$, $\widehat{\mathbb{E}}(X_j) = \frac{1}{n} \sum_{i=1}^n X_j^{(i)}$, and $\widehat{\mathbb{E}}(\widehat{\mathbb{E}}(X_j | X_S)) = \frac{1}{n} \sum_{i=1}^n \exp(\widehat{\theta}_j^S + \sum_{k \in S} \widehat{\theta}_{jk}^S X_k^{(i)})$, and $\widehat{\mathbb{E}}(\widehat{\mathbb{E}}(X_j | X_S)^2) = \frac{1}{n} \sum_{i=1}^n \exp(2\widehat{\theta}_j^S + 2 \sum_{k \in S} \widehat{\theta}_{jk}^S X_k^{(i)})$ where $\widehat{\theta}_S(j) = (\widehat{\theta}_j^S, \widehat{\theta}_{\setminus j}^S)$ is the solution of the following ℓ_1 -regularized GLM:

$$\widehat{\theta}_S(j) := \arg \min \frac{1}{n} \sum_{i=1}^n \left[-X_j^{(i)} \left(\theta_j + \sum_{k \in S} \theta_{jk} X_k^{(i)} \right) + \exp \left(\theta_j + \sum_{k \in S} \theta_{jk} X_k^{(i)} \right) \right] + \lambda_j \sum_{k \in S} |\theta_{jk}|. \quad (9)$$

This score is an estimator of the moments ratio relation in Equation (5). Hence, the correct element of the ordering has a score of 1, otherwise strictly greater than 1 in population. The ordering is determined one node at a time by selecting the node with the smallest score. Similar strategies of element-wise ordering learning can be found in many existing algorithms (e.g., Shimizu et al., 2011; Ghoshal and Honorio, 2017b, 2018; Drton et al., 2018).

The novelty of our algorithm is learning an ordering by testing which nodes have the smallest moments ratio in Equation (5) using the ℓ_1 -regularized GLM. By substituting the estimation of parameters $\theta(j)$ for an estimation of the conditional mean, we gain significant computational and statistical improvements compared to the previous works in Park and Raskutti (2015, 2018) where the method of moments is used for estimating the conditional mean and variance.

In principle, the number of conditional variances exponentially grows in the number of conditioning variables. Hence, if a conditioning set contains d -variables with 10 possible outcomes, then the number of possible computations is 10^d . In other words, the minimum sample size for the ODS algorithm to be implemented is possibly 10^d , otherwise, none of conditional variances can be estimated.

As we discussed, the problem of a learning directed graph structure is the same as the problem of an learning undirected graph structure if the ordering is known. Hence, given the estimated ordering, the parents of each node j can be learned via ℓ_1 -regularized GLM (see details in Meinshausen and Bühlmann, 2006; Wainwright et al., 2006; Ravikumar et al., 2011; Yang et al., 2015). Therefore, we determine the estimated parents of a node j as $\widehat{\text{Pa}}(j) := \{k \in S : \widehat{\theta}_{jk}^S \neq 0\}$ where $S = \widehat{\pi}_{1:(j-1)}$ and $\widehat{\theta}_S(j)$ is the solution to Equation (9).

3.1. Computational Complexity

The computational complexity for the MRS algorithm involves the ℓ_1 -regularized GLM algorithm (Friedman et al., 2009) where the worse-case complexity is $O(np)$ for a single ℓ_1 -regularized regression run. More precisely, the coordinate descent method updates each

Algorithm 1: Moments Ratio Scoring (MRS)

Input : n i.i.d. samples, $X^{1:n}$ **Output**: Estimated ordering $\hat{\pi} = (\hat{\pi}_1, \dots, \hat{\pi}_p)$ and an edge structure, $\hat{E} \subset V \times V$ Set $\hat{\pi}_0 = \emptyset$;**for** $m = \{1, 2, \dots, p\}$ **do** Set $S = \{\hat{\pi}_1, \dots, \hat{\pi}_{m-1}\}$; **for** $j \in \{1, 2, \dots, p\} \setminus S$ **do** Estimate $\hat{\theta}_S(j)$ for ℓ_1 -regularized generalized linear model (9); Calculate scores $\hat{S}(m, j)$ using Equation (8); **end** The m^{th} element of the ordering, $\hat{\pi}_m = \arg \min_j \hat{S}(m, j)$; The parents of the m^{th} element of the ordering, $\widehat{\text{Pa}}(\hat{\pi}_m) = \{k \in S \mid \hat{\theta}_{\hat{\pi}_m k}^S \neq 0\}$;**end****Return**: Estimate the edge set, $\hat{E} = \cup_{m \in V} \{(k, \hat{\pi}_m) \mid k \in \widehat{\text{Pa}}(\hat{\pi}_m)\}$

gradient in $O(p)$ operations. Hence, with d non-zero terms in the GLM, a complete cycle costs $O(pd)$ operations if no new variables become non-zero, and costs $O(np)$ for each new variable entered (see details in Friedman et al., 2010). Since our algorithm has p iterations and there are $p-j+1$ regressions with $j-1$ features for the j th iteration, the total worst-case complexity is $O(np^3)$.

The estimation of a Poisson MRF also involves a node-wise ℓ_1 -regularized GLM over all other variables, and hence the worse-case complexity is $O(np^2)$ if the coordinate descent method is exploited. The addition of estimation of ordering makes p times more computationally inefficient than the standard method for learning Poisson MRFs.

Learning a DAG model is NP-hard in general (Chickering et al., 1994). Hence, many state-of-the-art MEC and DAG learning algorithms, such as PC (Spirtes et al., 2000), GES (Chickering, 2003), and MMHC (Tsamardinos et al., 2006), are inherently greedy search algorithms. In the numerical experiments in Section 4, we compare MRS to greedy hill-climbing search-based GES and MMHC algorithms in terms of run time, and show that MRS has a significantly better computational complexity.

3.2. Theoretical Guarantees

In this section, we provide theoretical guarantees on the MRS algorithm for learning Poisson SEMs (4). The main result is expressed in terms of the triple (n, p, d) , where n is a sample size, p is a graph node size, and d is the indegree of a graph.

3.2.1. ASSUMPTIONS

We begin by discussing the assumptions we impose on Poisson SEMs. Since we apply ℓ_1 -regularized regression for the parent selection, most assumptions are similar to those imposed in Wainwright et al. (2006), Ravikumar et al. (2011), Yang et al. (2015) and Park and Raskutti (2018) where ℓ_1 -regularized regression was used for graphical model learning.

Important quantities are the Hessian matrices of the negative conditional log-likelihood of a node j given some subsets of the nodes in the ordering, $S_j \in \{\{\pi_1\}, \{\pi_1, \pi_2\}, \dots, \{\pi_1, \dots, \pi_{j-1}\}\}$. Let $Q^{j, S_j} := \nabla^2 \ell_j^{S_j}(\theta_S^*(j); X^{1:n})$ where

$$\ell_j^{S_j}(\theta_{S_j}(j), X^{1:n}) := \frac{1}{n} \sum_{i=1}^n \left[-X_j^{(i)} \left(\theta_j^{S_j} + \sum_{k \in S_j} \theta_{jk}^{S_j} X_k^{(i)} \right) + \exp \left(\theta_j^{S_j} + \sum_{k \in S_j} \theta_{jk}^{S_j} X_k^{(i)} \right) \right], \quad (10)$$

$$\theta_{S_j}^*(j) := \arg \min \mathbb{E} \left[-X_j \left(\theta_j^{S_j} + \sum_{k \in S_j} \theta_{jk}^{S_j} X_k \right) + \exp \left(\theta_j^{S_j} + \sum_{k \in S_j} \theta_{jk}^{S_j} X_k \right) \right]. \quad (11)$$

For ease of notation, we define a set for the non-zero elements of $\theta_{S_j}^*(j)$,

$$T_j := \{k \in S_j \mid \theta_{jk}^* \neq 0 \text{ where } \theta_{S_j}^*(j) = (\theta_j^*, \theta_{jk}^*)\}. \quad (12)$$

We note that if S_j contains all parents of j , $\text{Pa}(j) \subset S_j$, then $T_j = \text{Pa}(j)$. Lastly, for simplicity, we let A_{SS} denote the $|S| \times |S|$ sub-matrix of the matrix A corresponding to variables X_S .

Assumption 1 (Dependence Assumption) *For any $j \in V$ and any $S_j \in \{\{\pi_1\}, \{\pi_1, \pi_2\}, \dots, \{\pi_1, \dots, \pi_{j-1}\}\}$, there exist positive constants ρ_{\min} and ρ_{\max} such that*

$$\min_{j \in V} \lambda_{\min} \left(Q_{T_j T_j}^{j, S_j} \right) \geq \rho_{\min}, \quad \text{and} \quad \max_{j \in V} \lambda_{\max} \left(\frac{1}{n} \sum_{i=1}^n X_{\text{Pa}(j)}^{(i)} (X_{\text{Pa}(j)}^{(i)})^T \right) \leq \rho_{\max},$$

where T_j is in Equation (12), $\lambda_{\min}(A)$ and $\lambda_{\max}(A)$ are the smallest and largest eigenvalues of the matrix A , respectively.

Assumption 2 (Incoherence Assumption) *For any $j \in V$ and any $S_j \in \{\{\pi_1\}, \{\pi_1, \pi_2\}, \dots, \{\pi_1, \dots, \pi_{j-1}\}\}$, there exists a constant $\alpha \in (0, 1]$ such that*

$$\max_{j, S_j} \max_{t \in T_j^c} \|Q_{t T_j}^{j, S_j} (Q_{T_j T_j}^{j, S_j})^{-1}\|_1 \leq 1 - \alpha,$$

where T_j is in Equation (12).

Assumption 1 ensures that the parent variables are not too dependent. In addition, Assumption 2 ensures that parent and non-parent variables are not highly correlated. These two assumptions are standard in all neighborhood regression approaches to variable selection involving ℓ_1 -regularized based methods, and these conditions have imposed in proper works for both high-dimensional regression and graphical model learning.

To control the tail behavior of likelihood functions, we require a bounded sample assumption which is also imposed in the standard ℓ_1 -regularized Poisson regression (e.g., Jia et al., 2017).

Assumption 3 (Bounded Sample Assumption) For any $i \in \{1, 2, \dots, n\}$, $j \in V$, and for all $S_j \in \{\{\pi_1\}, \{\pi_1, \pi_2\}, \dots, \{\pi_1, \dots, \pi_{j-1}\}\}$, the samples are bounded:

$$\max_{i,j} \{X_j^{(i)}\} < C_x \log(\max\{n, p\}) \quad \text{and} \quad \max_{i,j} \{\exp(\theta_j^* + \sum_{k \in S_j} \theta_{jk}^* X_k^{(i)})\} < C_x \log(\max\{n, p\}).$$

where $C_x > 2$ is a positive constant.

Assumption 3 is closely related to the rate parameters. For instance, the rate parameter of $X_j^{(i)}$ is $\exp(\theta_j^* + \sum_{k \in \text{Pa}(j)} \theta_{jk}^* X_k^{(i)})$ by the definition of Poisson SEMs. Hence, Assumption 3 can be understood that too large rate parameters, that leads to a large value of a sample, are not allowed for all conditional distributions.

In fact, Assumption 3 is satisfied with a high probability when (θ_{jk}^*) are negative. Since the second condition in Assumption 3 is directly satisfied with negative (θ_{jk}^*) , we discuss the first condition: Using the union bound,

$$P\left(\max_{i,j} X_j^{(i)} \geq C_x \log(\max\{n, p\})\right) \leq n \cdot p \max_{i,j} \frac{\mathbb{E}(\exp(X_j^{(i)}))}{(\max\{n, p\})^{C_x}} \leq \max_{i,j} \frac{\mathbb{E}(\exp(X_j^{(i)}))}{(\max\{n, p\})^{C_x - 2}}.$$

In addition, the moment generating function is bounded when (θ_{jk}^*) are negative.

$$\mathbb{E}(\exp(X_j)) \leq \mathbb{E}(\mathbb{E}(\exp(X_j) \mid X_{\text{Pa}(j)})) \leq \mathbb{E}(\exp(\theta_j^* + \sum \theta_{jk}^* X_k)) \leq \exp(\theta_j^*).$$

Hence, given the negative (θ_{jk}^*) assumption, Assumption 3 is satisfied with probability at least $1 - \max_j \exp(\theta_j^*) / (\max\{n, p\})^{C_x - 2}$.

Lastly, we require a stronger version of the moments ratio relation in Equation (5), because we move from the population to the finite samples. This assumption only involves learning the ordering of a graph.

Assumption 4 For all $j \in V$ and $S_j \in \{\{\pi_1\}, \{\pi_1, \pi_2\}, \dots, \{\pi_1, \dots, \pi_{j-1}\}\}$, there exists an $M_{\min} > 0$ such that

$$\mathbb{E}(X_j^2) > (1 + M_{\min})\mathbb{E}[\mathbb{E}(X_j \mid X_{S_j}) + \mathbb{E}(X_j \mid X_{S_j})^2].$$

Now, we compare Assumptions 1, 2, 3, and 4 to the assumptions for learning Poisson MRFs and DAG models. As discussed, our assumptions are similar to the assumptions in Yang et al. (2015) and Park and Raskutti (2018) since all methods exploit the ℓ_1 -regularized GLM. However, the assumptions in Yang et al. (2015) only involve neighbors of node j , that is, $S_j = V \setminus j$. While our assumptions involve some subsets of parents, that is, $S_j \in \{\{\pi_1\}, \{\pi_1, \pi_2\}, \dots, \{\pi_1, \dots, \pi_{j-1}\}\}$ due to the unknown ordering. In addition, they do not assume the bounded sample assumption. However, they assume the restricted negative parameter space $\theta_{jk} < 0$ due to the normalizability issue. As we explained, if all parameters are negative in a Poisson SEM, the moment generating function is bounded, and hence, the bounded sample assumption is satisfied with a high probability. Lastly, Yang et al. (2015) does not have the moments ratio assumption, since it is only used for recovering the ordering.

We compare the required assumptions for the MRS and ODS algorithms in Park and Raskutti (2018). A major difference is that the MRS algorithm directly estimates the

graph, while the ODS algorithm estimates the moralized graph to reduce the search space of DAGs, and then, estimates the graph. Hence, our assumptions involve some parents of node j , while their assumptions involve not only parents, but neighbors of node j , that is, $S_j = \{\{\pi_1, \dots, \pi_{j-1}\}, V \setminus j\}$. In addition, they require a sparse moralized graph and adjacent faithfulness that are also known to be restrictive. We note that the sparse moralized graph assumption can be very strong since a sparse moralized graph is not implied by a sparse graph. For instance, consider a star graph where $X_1 \rightarrow X_j$ for all $j \in \{2, 3, \dots, p\}$ in Fig. 2. This star graph has the maximum degree of the moralized graph is $p-1$, while the maximum indegree is 1.

Another major difference is in the moments ratio assumption. More precisely, Park and Raskutti (2015, 2018) assume $\text{Var}(\mathbb{E}(X_j | X_S = x)) > c$ for all $x \in \mathcal{X}_S$ when $\text{Pa}(j) \not\subset S$, while we require $\mathbb{E}(\text{Var}(\mathbb{E}(X_j | X_S = x))) > c$. To emphasize the difference, we consider a 3-node graph $X_1 \rightarrow X_2 \rightarrow X_3$ where $X_1 \sim \text{Poisson}(\lambda)$, $X_2 | X_1 \sim \text{Poisson}(\exp(\theta_1 X_1))$, and $X_3 | X_2 \sim \text{Poisson}(\exp(\theta_2 X_2))$. Then, for $j = 3$ and $S = 1$, we have

$$\text{Var}(\mathbb{E}(X_3 | X_2) | X_1) = \text{Var}(\exp(\theta_2 X_2) | X_1) < \mathbb{E}(\exp(2\theta_2 X_2) | X_1) = \exp(e^{\theta_1 X_1} (e^{2\theta_2} - 1)).$$

Hence, for some constants θ_1, θ_2 and c , if $X_1 < \frac{1}{\theta_1}(\log \log c - \log(e^{2\theta_2} - 1))$, their assumption is not satisfied, while Assumption 4 holds.

Lastly, the ODS algorithm requires at least two distinct element of $X_{\text{Pa}(j)}^{(i)}$ for a conditional variance estimation, $\text{Var}(X_j | X_{\text{Pa}(j)})$. In principle, it can be 2^d by assuming all variables are binary. Hence when d is not so sparse, the ODS algorithm often fails to be implemented. In Section 4, we empirically verify that it can be a critical issue for the ODS algorithm when a graph is not so sparse ($d = 5$). Therefore, we believe that the assumptions for the MRS algorithm are more realistic.

Although our assumptions are standard in the previous works of ℓ_1 -regularized Poisson regressions, we have to note that the assumptions cannot be confirmed from data and they could be restrictive. However, they are not strong for ℓ_1 -regularized regression when samples are from Gaussian SEMs (see e.g., Ravikumar et al., 2011). Hence, we conjecture that our assumptions can be satisfied with a high probability under mild conditions, and leave this to future study.

3.2.2. MAIN RESULT

Putting together Assumptions 1, 2, 3, and 4, we have the following main result that a Poisson SEM can be recovered via our MRS algorithm in high-dimensional settings. The theorem provides not only sufficient conditions, but also the probability that our method recovers the true graph structure.

Theorem 3 *Consider a Poisson SEM (4) with parameter vector $(\theta(j))_{j \in V}$ and the maximum indegree of the graph d . Suppose that the regularization parameter (9) is chosen, such that*

$$\frac{4C_x^2 \sqrt{2}(2 - \alpha) \log^2(\max\{n, p\})}{\alpha \kappa_1(n, p)} \leq \lambda_j \leq \frac{\alpha \rho_{\min}^2}{10^2 C_x^2 (2 - \alpha) \rho_{\max} d \log^2(\max\{n, p\})},$$

for any $\alpha = (0, 1]$, and $\kappa_1(n, p) \geq \frac{4\sqrt{2} \cdot 10^2 C_x^4 \cdot (2 - \alpha)^2 \rho_{\max}}{\alpha^2 \rho_{\min}^2} d \log^4(\max\{n, p\})$. Suppose also that Assumptions 1, 2, 3 and 4 are satisfied and the values of the parameters in Equation (4) are

sufficiently large such that $\min_{(j,k) \in E} |\theta_{jk}| \geq \frac{10}{\rho_{\min}} \sqrt{d} \lambda_j$. Then, for any $\epsilon > 0$, there exists a positive constant C_ϵ such that if the sample size is sufficiently large $n > C_\epsilon (\kappa_1(n, p))^2 \log p$, then the MRS algorithm uniquely recovers the graph with a high probability:

$$P(\widehat{G} = G) \geq 1 - \epsilon.$$

Detailed proof is provided in Appendices C and D. Appendix C provides the error probability that ℓ_1 -regularized regression recovers the true parents of each node given the true ordering, and Appendix D provides the error probability that ℓ_1 -regularized regression recovers the ordering. The key technique for the proof is that the *primal-dual witness* method used in sparse regularized regressions and related techniques (Meinshausen and Bühlmann, 2006; Wainwright et al., 2006; Ravikumar et al., 2011; Yang et al., 2015). Theorem 3 intuitively makes sense because neighborhood selection via the ℓ_1 -regularized regression is a well-studied problem, and its bias can be controlled by choosing the appropriate regularization parameter λ_j . Hence, our moments ratio scores can be sufficiently close to the true scores to recover the true ordering.

Theorem 3 claims that if $n = \Omega(d^2 \log^9 p)$, our MRS algorithm recovers an underlying graph with a high probability. Hence, our MRS algorithm works in a high-dimensional setting, provided that the indegree of a graph d is bounded. This sample bound result shows that our method has much more relaxed constraints on the sparsity of the graph than the previous work in Park and Raskutti (2018), where the sample bound is $n = \Omega(\max\{d_m^4 \log^{12} p, \log^{5+d} p\})$. Moreover, it also shows that learning Poisson DAG models may require more samples than the learning Poisson MRFs in Yang et al. (2015), where the sample bound is $n = \Omega(d_m^2 \log^3 p)$ due to the existence of the ordering and the unrestricted parameter space.

3.2.3. POISSON SEM WITH A STAR GRAPH EXAMPLE

In this section, we discuss the validity of our assumptions using a special Poisson SEM with the star graph in Fig. 2 where $X_1 \sim \text{Poisson}(\lambda)$, $X_j | X_1 \sim \text{Poisson}(\exp(\theta X_1))$, for $j \in \{2, 3, \dots, p\}$. This consists of a single hub node connected to the rest of nodes. With this star graph, we show that our assumptions can be satisfied with positive (θ_{jk}) .

In order to discuss the validity of Assumptions 1, 2, 3, and 4 in this particular example, we first calculate the expectation of the Hessian matrix of Equation (10): For any $j \in \{2, 3, \dots, p\}$,

$$\begin{aligned} \mathbb{E}(X_1^2 \exp(\theta X_1)) &= \frac{\partial^2}{\partial \theta^2} \mathbb{E}(\exp(\theta X_1)) = \lambda \exp(\lambda(\exp(\theta) - 1) + \theta) (\lambda \exp(\theta) + 1), \\ \mathbb{E}(X_1 X_j \exp(\theta X_1)) &= \frac{\partial}{\partial \theta} \mathbb{E}(\exp(\theta X_1) X_j) = \frac{\partial}{\partial \theta} \mathbb{E}(\exp(\theta X_1) E(X_j | X_1)) \\ &= \frac{\partial}{\partial \theta} \mathbb{E}(\exp(2\theta X_1)) = 2\lambda \exp(\lambda(\exp(2\theta) - 1) + 2\theta). \end{aligned}$$

Hence, the population version of Assumption 1 is reduced to

$$\rho_{\min} < \lambda \exp(\lambda(\exp(\theta) - 1) + \theta) (\lambda \exp(\theta) + 1) \quad \text{and} \quad \lambda + \lambda^2 < \rho_{\max}.$$

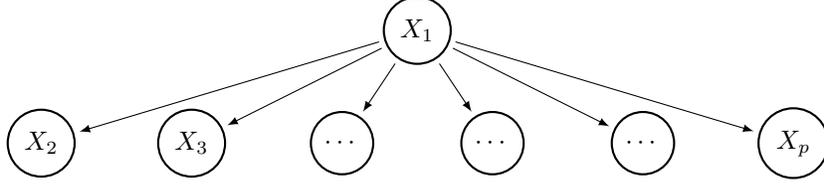


Figure 2: Star graph example

It can be satisfied with some positive values of θ . For $\lambda = 2$, $\rho_{\min} = 0.01$ and $\rho_{\max} = 10$, Assumption 1 is satisfied if $\theta > -3.426$. In addition, for $\lambda = 5$, $\rho_{\min} = 0.01$ and $\rho_{\max} = 50$, it is also satisfied if $\theta > -2.205$.

In addition, the population version of Assumption 2 can be written as

$$\max_{j \in V \setminus \{1\}} \max_{t \in V \setminus \{1, j\}} |\mathbb{E}(Q_{t1}^{j,1}) \mathbb{E}((Q_{11}^{j,1})^{-1})| = \frac{2 \cdot \exp(\lambda \exp(\theta)(\exp(\theta) - 1) + \theta)}{\lambda \exp(\theta) + 1} \leq 1 - \alpha.$$

This condition is also satisfied with positive values of θ . For $\lambda = 2$ and $\alpha = 0.01$, a simple algebra yields that Assumption 2 is satisfied if $\theta < 0.141$. In addition, for $\lambda = 5$ and $\alpha = 0.01$, the assumption is satisfied if $\theta < 0.165$.

In terms of Assumption 3, we also claim that it can be satisfied with some positive θ . Since the moment generating function of X_1 is $\exp(\lambda(e - 1))$, we have,

$$P(X_1^{(i)} > C_x \log(\max\{n, p\})) < \frac{\mathbb{E}(\exp(X_1^{(i)}))}{\max\{n, p\}^{C_x}} = \frac{\exp(\lambda(e - 1))}{\max\{n, p\}^{C_x}}.$$

where $C_x > 2$ is a positive constant in Assumption 3.

For other nodes $j \in \{2, 3, \dots, p\}$, we have,

$$P(X_j^{(i)} > C_x \log(\max\{n, p\})) \leq \frac{\mathbb{E}(\mathbb{E}(\exp(X_j^{(i)}) | X_1^{(i)}))}{\max\{n, p\}^{C_x}} = \frac{\mathbb{E}(\exp(\exp(\theta X_1^{(i)})(e - 1)))}{\max\{n, p\}^{C_x}}.$$

For $\theta < \log X_1^{(i)} / X_1^{(i)}$, we have,

$$P(X_j^{(i)} > C_x \log(\max\{n, p\})) \leq \frac{\mathbb{E}(\exp(X_1^{(i)}(e - 1)))}{\max\{n, p\}^{C_x}} = \frac{\exp(\lambda(e^{e-1} - 1))}{\max\{n, p\}^{C_x}}.$$

Hence, for $\theta < \log(C_x \log(\max\{n, p\})) / C_x \log(\max\{n, p\})$ that is the lower bound of $\log X_1^{(i)} / X_1^{(i)}$ given $X_1^{(i)} < C_x \log(\max\{n, p\})$, Assumption 3 is satisfied with a high probability:

$$P\left(\max_{i,j} X_j^{(i)} > C_x \log(\max\{n, p\})\right) \leq \frac{\exp(\lambda(e^{e-1} - 1))}{\max\{n, p\}^{C_x - 2}}.$$

Now, we discuss Assumption 4. A simple calculation shows that, for any $j \in \{2, 3, \dots, p\}$,

$$\mathbb{E}(X_j) = \exp(\lambda(\exp(\theta) - 1)), \text{ and } \mathbb{E}(X_j^2) = \exp(\lambda(\exp(2\theta) - 1)) + \exp(\lambda(\exp(\theta) - 1)).$$

Hence, Assumption 4 is equivalent to the constraint,

$$\exp(\lambda(\exp(2\theta) - 1)) > M_{\min} \exp(\lambda(\exp(\theta) - 1)) + (1 + M_{\min}) \exp(2\lambda(\exp(\theta) - 1)).$$

This condition is also satisfied with some positive θ . For $\lambda = 1$ and $M_{min} = 0$, as we discussed in Proposition 1, Assumption 2 is always satisfied with any value of $\theta \neq 0$. For $\lambda = 2$ and $M_{min} = 0.001$, Assumption 2 is satisfied if $|\theta| > 0.033$. Lastly, for $\lambda = 5$ and $M_{min} = 0.001$, Assumption 2 is satisfied if $|\theta| > 0.021$. Therefore, we show that for this particular star graph, Assumption 1, 2, 3, and 4 can be satisfied with a high probability by allowing positive θ .

Finally, we emphasize that the sample complexity of the MRS algorithm, $n = \Omega(d^2 \log^9 p)$, does not rely on the maximum degree of the moralized graph, d_m , while many DAG learning algorithms using the sparsity of the moralized graph or Markov blanket inevitably depend on d_m . For the star graph with $d = 1$ and $d_m = p - 1$, the MRS algorithm requires $n = \Omega(\log^9 p)$ to recover the graph in high dimensional settings, while the ODS algorithm may fail since its sample complexity is $\Omega(d_m^4 \log^{12} p)$. This fact implies that, unlike the ODS algorithm, the MRS algorithm can recover a sparse graph containing hub nodes in high dimensional settings.

4. Numerical Experiments

In this section, we provide simulation results to support our main theoretical results of Theorem 3 and the computational complexity in Section 3.1: (i) the MRS algorithm recovers the ordering and edges more accurately as sample size increases; (ii) the required sample size $n = \Omega(d^2 \log^9 p)$ depends on the number of nodes p and the complexity of the graph d ; (iii) the MRS algorithm accurately learns the graphs in high-dimensional settings ($p > n$); and (iv) the computational complexity is $O(np^3)$ at worst. We also show that the MRS algorithm performs favorably compared to the ODS (Park and Raskutti, 2015), GES (Chickering, 2003), and MMHC (Tsamardinos et al., 2006) algorithms. In addition, we investigate how sensitive our MRS algorithm is to deviations from the assumption about the link functions by using the identity link function in Equation (3). Lastly, we also investigate how well the MRS algorithm recovers undirected edges when samples are generated by Poisson and truncated Poisson MRFs (Yang et al., 2013, 2015; Inouye et al., 2017).

4.1. Random Poisson SEMs

We conducted simulations using 200 realizations of p -node Poisson SEMs (4) with the randomly generated underlying DAG structures while respecting the indegree constraints $d \in \{1, 5, 10\}$. A graph with $d = 1$ is a special case where there is no v-structure, and therefore, the corresponding MEC is completely undirected. The set of non-zero parameters $\theta_j, \theta_{jk} \in \mathbb{R}$ in Equation (4) was generated uniformly at random in the range $\theta_j \in [1, 3]$, $\theta_{jk} \in [-1.5, -0.5] \cup [0.5, 1.5]$ for $d = 1$, and $\theta_{jk} \in [-1, -0.1] \cup [0.1, 1]$ for $d = 5, 10$, which helps the generated values of samples to avoid either all zeros or from going beyond the maximum possible value of the R program ($> 10^{309}$). Nevertheless, if some samples were beyond the maximum possible value, we regenerated the parameters and samples.

The MRS and ODS algorithms were implemented using ℓ_1 -regularized likelihood where we used five-fold cross validation to choose the regularization parameters. Where mean squared error was within two standard error of the minimum mean squared error, we chose the minimum value for the moments ratio scores and the largest value for parent selection. That was because a less biased estimator is preferred for the score calculation, and we

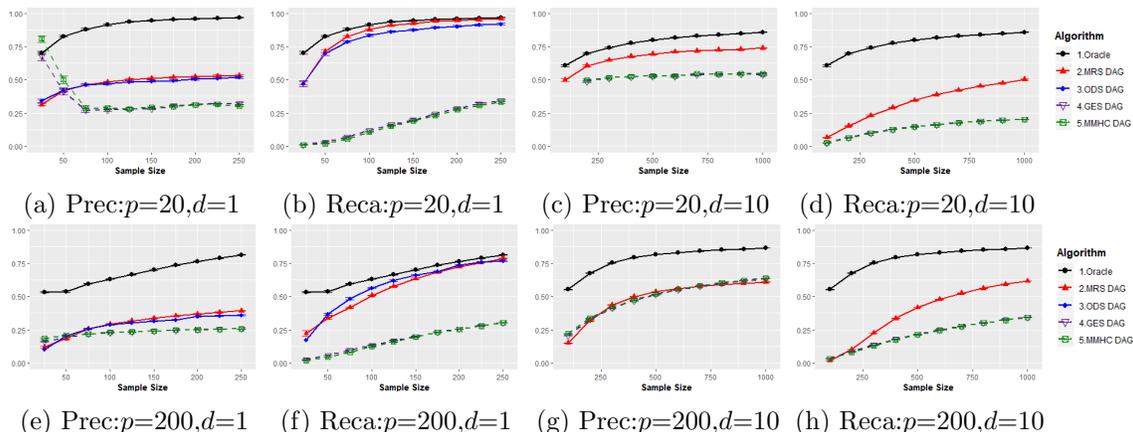


Figure 3: Comparison of the MRS algorithm to the oracle, ODS, GES and MMHC algorithms in terms of precision and recall for Poisson SEMs with $p \in \{20, 200\}$ and $d \in \{1, 10\}$.

preferred a sparse graph containing only legitimate edges. We acknowledge that the level of sparsity can be adjusted according to the importance of precision or recall.

In Fig. 3, we compare the MRS algorithm to state-of-the-art ODS, GES and MMHC algorithms for graph node size $p = \{20, 200\}$, varying sample size $n \in \{25, 50, \dots, 250\}$ for $d = 1$ and $n = \{100, 200, \dots, 1000\}$ for $d = 10$, and provide two results: (i) the average precision ($\frac{\# \text{ of correctly estimated edges}}{\# \text{ of estimated edges}}$); (ii) the average recall ($\frac{\# \text{ of correctly estimated edges}}{\# \text{ of true edges}}$). As discussed, the both GES and MMHC algorithms only recover the partial graph by leaving some arrows undirected. Therefore, we also provide average precision and recall for the estimated MECs in Fig. 4. Lastly, we provide an oracle, where the true parents of each node are used, while the ordering is estimated via ℓ_1 -regularized GLM. Hence, we can see where the errors come from between the ordering estimation or parent selection. We considered more parameters (θ_{jk}, n, p, d) , but for brevity, we focus on these settings.

As we can see in Fig. 3, the MRS algorithm more accurately recovers the true directed edges as sample size increases. In addition, the MRS algorithm is more precise for small sparse graphs than for large-scale or dense graphs, given the same sample size. Hence it confirms that the MRS algorithm is consistent, and the sample bound $n = \Omega(d^2 \log^9 p)$ depends on p and d .

The MRS algorithm significantly outperforms state-of-the-art GES and MMHC algorithms in terms of both precision and recall, on average, except for cases $p = 20, d = 1, n \leq 50$. It is worth noting that the GES and MMHC algorithms are not consistent, because the recall for any tree graph must be zero in population, whereas the recall from GES and MMHC increases as sample size increases. Hence, we can conclude that the GES and MMHC algorithms find correct directed edges by finding incorrect v-structures. It is an expected result because the comparison methods only work with a non-faithful distribution, which rarely arises in finite sample settings (Uhler et al., 2013).

Fig. 3 shows that the MRS and ODS algorithms have similar performance in identifying directed edges when the indegree is a small $d = 1$. It makes sense because the ODS

n	100	200	300	400	500	600	700	800	900	1000
p = 20	199	175	107	64	1	0	0	0	0	0
p = 50	200	200	200	199	192	179	151	140	99	86

Table 1: Number of failures in ODS algorithm implementations from among 200 sets of samples for different node sizes $p \in \{20, 50\}$, and sample sizes $n \in \{100, 200, \dots, 1000\}$, when the indegree is $d = 5$.

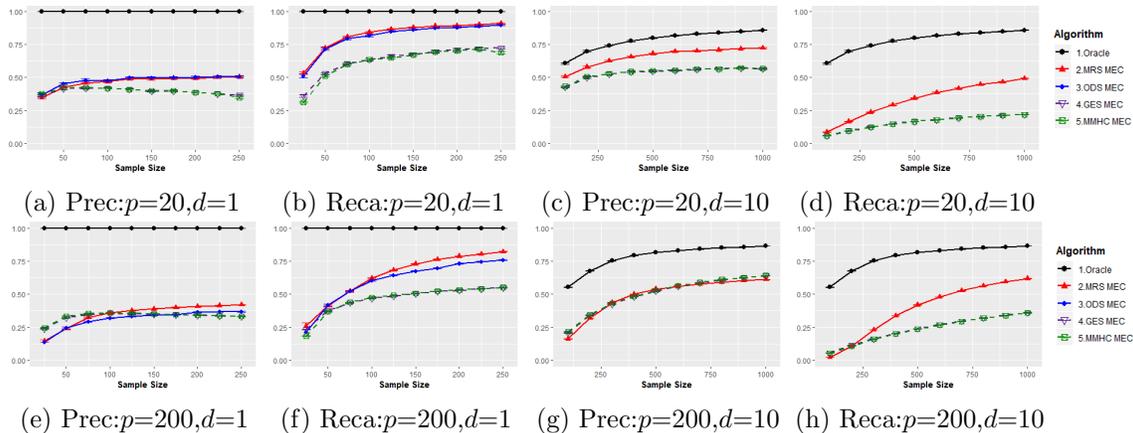


Figure 4: Comparison of the MRS algorithm to the oracle, ODS, GES, and MMHC algorithms in terms of the precision and recall for the MECs of Poisson SEMs with $p \in \{20, 100\}$ and $d \in \{1, 10\}$.

algorithm recovers any Poisson DAG models if the moralized graph is sparse. In other words, the accuracy of the ODS algorithm may be poor for the non-sparse graph. Moreover, the ODS algorithm often fails to be implemented due to a lack of samples for the estimation of conditional variance, that is, $\sum_{i=1}^n \mathbf{1}(X_S^{(i)} = x) < 2$ for all $x \in \mathcal{X}_S$. Table 1 shows the number of failures in the ODS algorithm implementations for node size $p \in \{20, 50\}$ and sample size $n \in \{100, 200, \dots, 1000\}$ when the indegree is $d = 5$, and the degree of the moralized graph is at most $d_m = p - 1$. It empirically confirms that the ODS algorithm requires a huge number of samples to be implemented when a true graph is not sparse. Hence, we do not apply the ODS algorithm for the graphs with $d = 10$. It is consistent with our main result that our method can learn the Poisson SEMs with some hub nodes while the ODS algorithm might not.

Fig. 4 shows the analogous results for the recovery of MECs, in which the MRS and all comparison algorithms consistently learn the true MECs. The performance of the MRS algorithm gets better as sample size increases or node size decreases. In addition, we can see that the MRS algorithm still recovers the MEC of the Poisson SEM better on average than the comparison methods. However, it must be pointed out that our MRS algorithm applies to Poisson SEMs (4), while the ODS algorithm accurately learns sparse Poisson

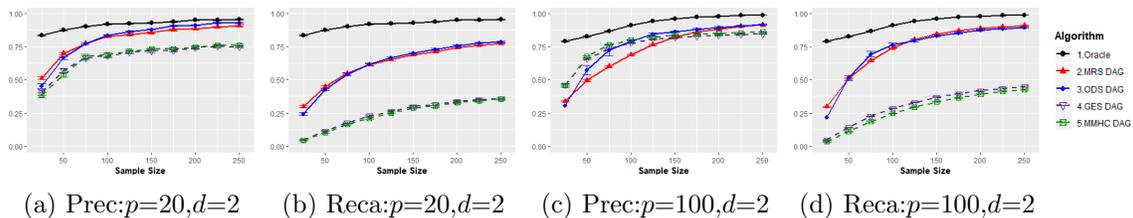


Figure 5: Comparison of the MRS algorithm to the oracle, ODS, GES and MMHC algorithms in terms of the precision and recall for Poisson DAG models with $p \in \{20, 100\}$, $d = 2$, and the identity link function.

DAG models where arbitrary link functions are allowed. In addition, the GES and MMHC algorithms apply to more general classes of DAG models.

4.2. Random Poisson DAG Models

When the data are generated by a random Poisson DAG model (2) where g_j is not the standard log link function, our MRS algorithm is not guaranteed to estimate the true directed acyclic graph and its ordering. Hence, an important question is how sensitive our method is to deviations from the link assumption. In this section, we empirically investigate this question.

We generated the 200 samples with the same procedure specified in Section 4.1, but with the indegree constraint $d = 2$, and except that identity link function $g_j(\eta) = \eta$ and the range of parameters was $\theta_{jk} \in [-1.5, -0.5] \cup [0.5, 1.5]$. We note that the link function must be positive, but we allow the negative value of θ_{jk} by randomly choosing $\theta_j \in [1, 10]$. If any Poisson rate parameter is negative, we regenerated the parameters.

In Fig. 5, we compare the MRS to state-of-the-art ODS, GES and MMHC algorithms for varying sample size $n \in \{25, 50, \dots, 250\}$, and node size $p \in \{20, 100\}$. Fig. 5 shows that the MRS algorithm consistently recovers the true graph, and hence, we can see that the MRS algorithm is not so sensitive to deviations from the link assumption. Comparing it to the ODS algorithm, the MRS algorithm shows slightly worse performance because the ODS algorithm is designed to learn general Poisson DAG models with any type of link functions. However, we can see that the MRS algorithm still performs better than the greedy search-based methods in both average precision and recall.

4.3. Random Poisson and Truncated Poisson Markov Random Fields

When samples are generated by a Poisson or truncated Poisson MRF, our MRS algorithm is not guaranteed to find the true dependence relationships of variables. Hence, it is also important to investigate how well our algorithm recovers undirected edges when multivariate count data is from an MRF. In this section, we compare our MRS algorithm to state-of-the-art Poisson MRF (PMRF) and truncated Poisson MRF learning (TMRF) algorithms (Yang et al., 2013, 2015; Inouye et al., 2017) when multivariate count data is from Poisson MRFs and truncated Poisson MRFs, respectively. We used the R package XMRF (Wan et al., 2016) for truncated Poisson MRFs.

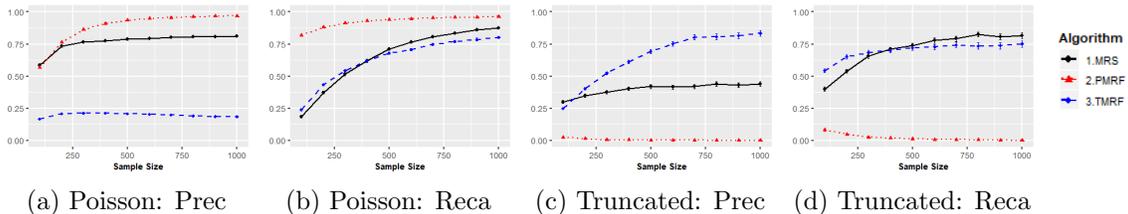


Figure 6: Comparison of the MRS algorithm to the Poisson MRF learning (PMRF) and truncated Poisson MRF learning (TMRF) algorithms in terms of the precision and recall for undirected edges of random 20-nodes Poisson MRFs and truncated Poisson MRFs with $d_m = 5$, and $R = 100$.

We generated 100 samples of 20-nodes random Poisson MRF and truncated Poisson MRF with the randomly generated underlying undirected graphs, respectively. For Poisson MRFs, we set the maximum Markov blanket $d_m = 5$ and the non-zero parameters in Equation (6) was generated uniformly at random in the range $\theta_j \in [1, 2]$, but we fixed $\theta_{jk} = -0.1$ for all $j \in V$. This is a similar setting used in Yang et al. (2015). For truncated Poisson MRFs, we set $d_m = 5$, $\theta_j = 0$, $\theta_{jk} = 0.1$, and the truncation level is $R = 100$, meaning that all samples are less than 100 (see details in Equation 3 of Yang et al., 2013). In terms of the choice of regularization parameters for the MRS and PMRF algorithms, we used five-fold cross validation as we used in Section 4.1. For the TMRF algorithm, we set the regularization parameters to 0.1 since this value seems to work well.

Fig. 6 compares the MRS algorithm to state-of-the-art PMRF and TMRF algorithms in terms of recovering undirected edges by varying sample size $n \in \{100, 200, \dots, 1000\}$. For a fair comparison, we used the skeleton of the estimated MEC via the MRS algorithm, because our algorithm returns a DAG. As we can see in Fig. 6, the MRS algorithm consistently finds the true edges from both Poisson MRF and truncated Poisson MRF samples. Hence, we empirically verify that the MRS algorithm can recover some dependence relationships of variables even if samples are from Poisson or truncated Poisson MRFs.

Fig. 6 also shows that the MRS algorithm performs significantly worse than the comparison PMRF and TMRF algorithm, on average, when samples are from Poisson MRFs and truncated Poisson MRFs, respectively. It is an expected result because the PMRF and TMRF algorithms are for learning Poisson MRFs and truncated MRFs, while our algorithm is for Poisson SEMs. However, it is worth noting that the TMRF algorithm seems not to work on average when samples are from a Poisson MRF in our setting. It is mainly because the TMRF algorithm is for learning truncated Poisson MRFs, not Poisson MRFs. We emphasize that, in another setting where θ_j is fixed to 1, the TMRF algorithm works much better. It is also worth noting that the PMRF algorithm seems not to recover any undirected edges when samples are from a truncated Poisson MRF. It can be clearly explained by the fact that the PMRF algorithm cannot capture the positive dependencies, however all parameters are positive in our setting.

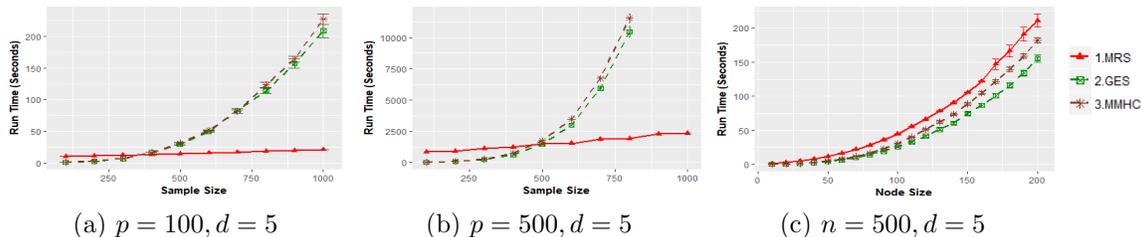


Figure 7: Comparison of the MRS algorithm to the GES and MMHC algorithms in terms of the running time with respect to node size p and sample size n

4.4. Computational Complexity

Fig. 7 compares the run-time of the MRS, GES, and MMHC algorithms for learning Poisson SEMs with indegree $d = 5$ by varying sample size $n \in \{100, 200, \dots, 1000\}$ with fixed node size $p \in \{100, 500\}$, and varying node size $p \in \{10, 20, \dots, 200\}$ with fixed sample size $n = 500$. Fig. 7 supports the worst case computational complexity $O(np^3)$ discussed in Section 3.1. In addition, it shows that the MRS algorithm is significantly faster than the greedy search-based GES and MMHC algorithms when a sample size is large ($n > 500$).

5. Real Multivariate Count Data: MLB Statistics

We now apply the MRS algorithm and state-of-the-art ODS and MMHC algorithms to a simple data set that involves multivariate count data that models baseball statistics for Major League Baseball (MLB) players during the 2003 season. To the best of our knowledge, our MRS algorithm is the only algorithm that provides a reliable and scalable approach to non-sparse DAG learning with multivariate count data although it is under strong assumptions. In particular, other approaches, such as PC, MMHC, and approaches based on conditional independence testing, suffer severely from the fact that we are dealing with count variables where the number of discrete states is potentially infinite. In addition, ODS algorithm cannot deal with a non-sparse graph such as a graph containing hub nodes. Lastly, both Poisson MRF and truncated Poisson MRF may provide an extremely complicated graph because it connects all pairs of nodes having a common child like a moralized graph.

Our original data set consists of 800 MLB player salary and batting statistics from the 2003 season (see R package Lahman in Friendly, 2017 for detailed information). The data set contains 23 covariates: Salary, Number of: Games Played (G), At Bats (AB), Runs (R), Hits (H), Doubles (X2B), Triples (X3B), Home Runs (HR), Runs Batted In (RBI), Stolen Bases (SB), times Caught Stealing (CS), Bases on Balls (BB), Strikeouts (SO), Intentional Walks (IBB), times Hit by Pitch (HBP), Sacrifice Hits (SH), Sacrifice Flies (SF), and times Grounded into Double Plays (GIDP), plus Player ID, Year ID, Stint, Team ID, and League ID. However, we eliminated Player ID, Year ID, Stint, Team ID, and League ID because our focus is to find the directional or causal relationships between salary and batting statistics. In addition, we only considered players in the top 25% in terms of the number of games played, because the baseball statistics relationships from players who played only a few

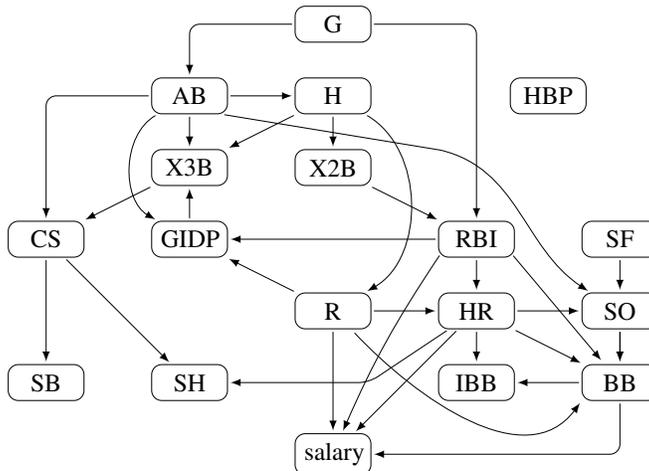


Figure 8: MLB player statistics directed graph estimated by the MRS algorithm for Poisson DAG models.

games could be uncertain. Therefore, the data set we considered contained 18 variables and 200 observations.

We assumed each node to a conditional distribution given its parents is Poisson because most MLB statistics, except for salary, reflect the number of successes or attempts that were counted during the season. Hence, we applied the MRS algorithm for Poisson DAG models with leave-one-out cross validation to choose the tuning parameters, and we chose the largest value where mean squared error is within 2.5 standard error of the minimum mean squared error, because we prefer a sparse graph containing only legitimate edges.

Fig. 8 shows the directed graph estimated by our MRS algorithm. The estimated graph reveals clear causal/directional relationships between batting statistics. This makes sense, because players with larger numbers of HR, BB, RBI, and/or R have a better salary. The more games played, or the more batting chances, the higher H, BB, SO, RBI, and other statistics. Moreover, the higher the total number of hits, the more X2Bs, X3Bs, Rs and the fewer SOs. Players with more home runs and base on balls get intentional walks more frequently. Lastly, the more stolen bases are attempted, the more they are caught stealing, because there is no success without failure.

We acknowledge that our proposed DAG model returns many errors due to restrictive assumptions that are not completely satisfied by the real data. However, the benefit is best seen by comparing MRS to other DAG learning approaches and an undirected graphical model for multivariate count data. In particular, we applied Poisson undirected graphical models (Yang et al., 2015) in which ℓ_1 -regularized Poisson regressions are applied. We provide the estimated undirected graph with the largest tuning parameter where mean squared of error is within 2.5 standard error of the minimum mean squared error. The estimated undirected graph in Fig. 9 (left side) shows that a lot of nodes are connected by edges, that many edges are unexplainable, and that some legitimate edges are missing (e.g., [H, X3B], [SB, CS] are not connected), because the Poisson undirected graphical model only

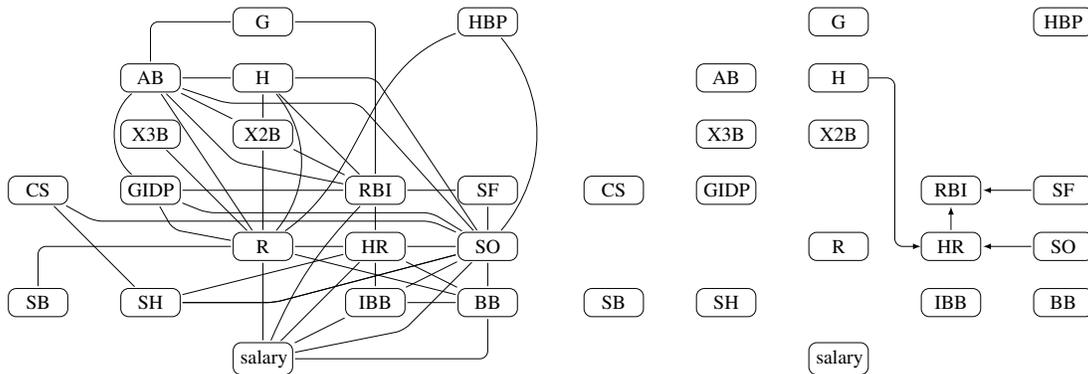


Figure 9: MLB player statistics undirected graph estimated by ℓ_1 -penalized likelihood regression (left) and a directed acyclic graph estimated by the MMHC algorithm (right).

permits negative conditional relationships, whereas most variables are positively correlated. Hence, it may not be useful to understand the relationships between MLB statistics.

We also compared the MMHC algorithm. As discussed, the MMHC algorithm does not guarantee to find a complete directed graph, and prefers a sparser graph when the faithfulness assumption is violated, which often arises in finite sample settings (Uhler et al., 2013). Hence, the estimated directed graph in Fig. 9 (right side) is extremely sparse, with only four directed edges: $[H, HR]$, $[SO, HR]$, $[HR, RBI]$, and $[SF, RBI]$. Lastly, ODS algorithm failed to be implemented as expected because of some hub nodes such as the number of games, at bats, and runs batted in.

Since our method is the first identifiability result for the strongly correlated count data when variables are directional/causal relationships and there exist hub variables, to the best of our knowledge, our method better identifies the directional/causal relationships between MLB statistics. However, we acknowledge that, like most other DAG-learning approaches, very strong assumptions, such as dependency, incoherence, are required for reliable recovery.

6. Future Works

Several topics remain for future works. Although our assumptions are similar to the assumptions in the previous works of ℓ_1 -regularized Poisson regression, our assumptions could be very restrictive. In addition, they cannot be confirmed from data. However, we conjecture that the assumptions are satisfied with a high probability under mild conditions, and one may be able to prove this. In addition, it is an important problem of finding the minimax rate of the Poisson DAG models, and it should be investigated in the future. Lastly, it would be also interesting to explore if our idea can be applied to other structural equation models with Binomial, Negative Binomial, Exponential, and Gamma distributions. We believe that our node-wise ℓ_1 -regularized based approach can be extended to the identifiable linear SEMs under some suitable conditions.

Acknowledgments

This work was supported under the framework of international cooperation program managed by National Research Foundation of Korea(NRF-2018R1C1B5085420)

Appendix A. Proof for Proposition 1

Proof For a notational simplicity, we define a moments related function for Poisson, $f(\mu) = \mu + \mu^2$ for $\mu > 0$. Then, for any node $j \in V$, any non-empty set $S_j \subset \text{Nd}(j)$,

$$\mathbb{E}(X_j^2 | S_j) = \mathbb{E}(\mathbb{E}(X_j^2 | X_{\text{Pa}(j)}) | S_j) = \mathbb{E}(f(\mathbb{E}(X_j | X_{\text{Pa}(j)})) | S_j).$$

Using the Jensens inequality and $f(\cdot)$ is convex, we have,

$$\mathbb{E}(f(\mathbb{E}(X_j | X_{\text{Pa}(j)})) | S_j) \geq f(\mathbb{E}(\mathbb{E}(X_j | X_{\text{Pa}(j)}) | S_j)) = f(\mathbb{E}(X_j | S_j)).$$

Using the fact that $\mathbb{E}(X_j | X_{\text{Pa}(j)}) = g_j(X_{\text{Pa}(j)})$ and it is non-degenerated by definition, the equality only holds when S_j contains all parents of j , $\text{Pa}(j) \subset S_j \subset \text{Nd}(j)$.

By restating the above inequality, we have,

$$\mathbb{E}(X_j^2 | S_j) - \mathbb{E}(X_j | S_j) - \mathbb{E}(X_j | S_j)^2 \geq 0.$$

In addition, by taking the expectations, we have,

$$\mathbb{E}(X_j^2) - \mathbb{E}(\mathbb{E}(X_j | X_{S_j}) + \mathbb{E}(X_j | X_{S_j})^2) \geq 0.$$

Since j and S_j are arbitrary, we complete the first part of the proof.

Now, we prove that $\mathbb{E}(X_j^2) \geq \mathbb{E}(\mathbb{E}(X_j | X_{S_j}) + \mathbb{E}(X_j | X_{S_j})^2)$ is equivalent to $\mathbb{E}(\text{Var}(\mathbb{E}(X_j | \text{Pa}(j)) | X_{S_j})) \geq 0$. Using the total variance decomposition, we have,

$$\mathbb{E}(\text{Var}(X_j | X_{S_j})) = \mathbb{E}(\mathbb{E}(\text{Var}(X_j | X_{\text{Pa}(j)}) | X_{S_j})) + \mathbb{E}(\text{Var}(\mathbb{E}(X_j | X_{\text{Pa}(j)}) | X_{S_j})).$$

Using the fact that the conditional distribution, $X_j | X_{\text{Pa}(j)}$, is Poisson where its mean and variance are equal, we have,

$$\mathbb{E}(\text{Var}(\mathbb{E}(X_j | X_{\text{Pa}(j)}) | X_{S_j})) = \mathbb{E}(\text{Var}(X_j | X_{S_j})) - \mathbb{E}(X_j).$$

Using the definition of the conditional variance, we have,

$$\mathbb{E}(\text{Var}(X_j | X_{S_j})) - \mathbb{E}(X_j) = \mathbb{E}(X_j^2) - \mathbb{E}(\mathbb{E}(X_j | X_{S_j}) + \mathbb{E}(X_j | X_{S_j})^2).$$

Therefore, we complete the proof. ■

Appendix B. Proof for Theorem 2

Proof Without loss of generality, we assume the true ordering is unique, and $\pi = (\pi_1, \dots, \pi_p)$. For simplicity, we define $X_{1:j} = (X_{\pi_1}, X_{\pi_2}, \dots, X_{\pi_j})$ and $X_{1:0} = \emptyset$. In addition, we define a moments related function, $f(\mu) = \mu + \mu^2$.

We now prove identifiability of Poisson DAG models using mathematical induction:

Step (1) For the first step π_1 , using Proposition 1, we have $\mathbb{E}(X_{\pi_1}^2) = \mathbb{E}(f(\mathbb{E}(X_{\pi_1})))$, while for any node $j \in V \setminus \{\pi_1\}$: $\mathbb{E}(X_j^2) > \mathbb{E}(f(\mathbb{E}(X_j)))$.

Hence, we can determine π_1 as the first element of the causal ordering.

Step (m-1) For the $(m-1)^{th}$ element of the ordering, assume that the first $m-1$ elements of the ordering and their parents are correctly estimated.

Step (m) Now, we consider the m^{th} element of the causal ordering and its parents. It is clear that π_m achieves $\mathbb{E}(X_{\pi_m}^2) = \mathbb{E}(f(\mathbb{E}(X_{\pi_m} | X_{1:(m-1)})))$. However, for $j \in \{\pi_{m+1}, \dots, \pi_p\}$, $\mathbb{E}(X_j^2) > \mathbb{E}(f(\mathbb{E}(X_j | X_{1:(m-1)})))$ by Proposition 1. Hence, we can estimate a true m^{th} component of the ordering π_m .

In terms of the parent search, it is clear that by conditional independence relations naturally encoded by factorization (1) $\mathbb{E}(X_{\pi_m}^2) = \mathbb{E}(f(\mathbb{E}(X_{\pi_m} | X_{1:(m-1)}))) = \mathbb{E}(f(\mathbb{E}(X_{\pi_m} | X_{\text{Pa}(\pi_m)})))$. Hence, we can also choose the minimum conditioning set from among $X_{1:(m-1)}$ as the parents of π_m such that the above moments relation holds. By mathematical induction, this completes the proof. ■

Appendix C. Proof for Theorem 3: Parents Recovery

Proof We provide the proof for Theorem 3 using the *primal-dual witness method* that is also used many other works (Meinshausen and Bühlmann, 2006; Wainwright et al., 2006; Ravikumar et al., 2011; Yang et al., 2015). In this proof, we show in Appendix C, the error probability for the recovery of the parents of a node π_j from among all the nodes given the partial ordering $(\pi_1, \pi_2, \dots, \pi_{j-1})$ via ℓ_1 -regularized regression. In Appendix D, the error bounds for the recovery of the ordering both via ℓ_1 -regularized regression.

Without loss of generality, let the true ordering be $\pi = (1, 2, \dots, p)$, and hence, $\pi_{1:j} = (\pi_1, \pi_2, \dots, \pi_j) = (1, 2, \dots, j)$. For ease of notation, $[\cdot]_k$ and $[\cdot]_S$ denote parameters corresponding to variable X_k and random vector X_S , respectively. In order to make the arguments easier to understand, we restate the negative log likelihood (10) and related arguments.

First, we define a new parameter vector $\theta_{S_j} \in \mathbb{R}^{|S_j|}$ without parameter θ_j corresponding to the node j since the node j is not penalized in regression problem (9). Then, the conditional negative log-likelihood of the GLM for X_j given X_{S_j} can be written as:

$$\ell_j^{S_j}(\theta_{S_j}; X^{1:n}) := \frac{1}{n} \sum_{i=1}^n \left(-X_j^{(i)} \langle \theta_{S_j}, X_{S_j}^{(i)} \rangle + \exp(\langle \theta_{S_j}, X_{S_j}^{(i)} \rangle) \right), \quad (13)$$

where $\langle \cdot, \cdot \rangle$ is an inner product.

We also define $\theta_{S_j}^* \in \mathbb{R}^{|S_j|}$ for Equation (11):

$$\theta_{S_j}^* := \arg \min_{\theta \in \mathbb{R}^{|S_j|}} \mathbb{E} \left(-X_j \langle \theta, X_{S_j} \rangle + \exp(\langle \theta, X_{S_j} \rangle) \right). \quad (14)$$

We define a set non-zero elements index of $\theta_{S_j}^*$ as in Equation (12), $T_j := \{k \in S_j \mid [\theta_{S_j}^*]_k \neq 0\}$ where $\theta_{S_j}^*$ is in Equation (14).

The main goal of the proof is to find the unique minimizer of the following convex problem:

$$\hat{\theta}_{S_j} := \arg \min_{\theta \in \mathbb{R}^{|S_j|}} \mathcal{L}_j(\theta, \lambda_j) = \arg \min_{\theta \in \mathbb{R}^{|S_j|}} \{ \ell_j^{S_j}(\theta; X^{1:n}) + \lambda_j \|\theta\|_1 \}. \quad (15)$$

By setting the *sub-differential* to 0, $\widehat{\theta}_{S_j}$ satisfies the following condition:

$$\nabla_{\theta} \mathcal{L}_j^{S_j}(\widehat{\theta}_{S_j}, \lambda_j) = \nabla_{\theta} \ell_j^{S_j}(\widehat{\theta}_{S_j}; X^{1:n}) + \lambda_j \widehat{Z}_j^{S_j} = 0 \quad (16)$$

where $\widehat{Z}_j^{S_j} \in \mathbb{R}^{|S_j|}$ and $[\widehat{Z}_j^{S_j}]_t = \text{sign}([\widehat{\theta}_{S_j}]_t)$ if $t \in T_j$, otherwise $[\widehat{Z}_j^{S_j}]_t < 1$.

Lemma 4 directly follows from the prior work (Yang et al., 2015), where each node's conditional distribution is in the form of a generalized linear model.

Lemma 4 (Uniqueness of Solution, Lemma 8 in Yang et al., 2015) *Suppose that $|[\widehat{Z}_j^{S_j}]_t| < 1$ for $t \notin T_j$ in Equation (16). Then, the solution $\widehat{\theta}_{S_j}$ of Equation (15) satisfies $[\widehat{\theta}_{S_j}]_t = 0$ for all $t \notin T_j$. Furthermore, if the sub-matrix of Hessian matrix $Q_{T_j T_j}^{S_j}$ is invertible, then $\widehat{\theta}_{S_j}$ is unique.*

The remainder of the proof is to show $|[\widehat{Z}_j^{S_j}]_t| < 1$ for all $t \notin T_j$. Note that the restricted solution in Equation (22) is $(\widetilde{\theta}_{S_j}, \widetilde{Z}_j^{S_j})$ and the unrestricted solution in Equation (15) is $(\widehat{\theta}_{S_j}, \widehat{Z}_j^{S_j})$. Equation (16) with the dual solution can be represented by

$$\nabla^2 \ell_j^{S_j}(\theta_{S_j}^*; X^{1:n})(\widetilde{\theta}_{S_j} - \theta_{S_j}^*) = -\lambda_j \widetilde{Z}_j^{S_j} - W_j^{S_j} + R_j^{S_j} \quad (17)$$

where

(a) $W_j^{S_j}$ is the sample score function:

$$W_j^{S_j} := -\nabla \ell_j(\theta_{S_j}^*; X^{1:n}). \quad (18)$$

(b) $R_j^{S_j} = (R_{jk}^{S_j})_{k \in S_j}$ and $R_{jk}^{S_j}$ is the remainder term by applying the coordinate-wise mean value theorem:

$$R_{jk}^{S_j} := [\nabla^2 \ell_j^{S_j}(\theta_{S_j}^*; X^{1:n}) - \nabla^2 \ell_j^{S_j}(\bar{\theta}_{S_j}; X^{1:n})]_k^T (\widetilde{\theta}_{S_j} - \theta_{S_j}^*). \quad (19)$$

Here $\bar{\theta}_{S_j}$ is a vector on the line between $\widetilde{\theta}_{S_j}$ and $\theta_{S_j}^*$, and $[\cdot]_k^T$ is the row of a matrix corresponding to variable X_k .

Then, the following proposition provides a sufficient condition to control $\widetilde{Z}_j^{S_j}$.

Proposition 5 *If $\max(\|W_j^{S_j}\|_{\infty}, \|R_j^{S_j}\|_{\infty}) \leq \frac{\lambda_j \alpha}{4(2-\alpha)}$, then $|[\widetilde{Z}_j^{S_j}]_t| < 1$ for all $t \notin T_j$.*

Next, we introduce the following three lemmas under Assumptions 1, 2, and 3 to show that conditions in Proposition 5 hold. For ease of notation, let $\eta = \max\{n, p\}$, $\widetilde{\theta}_S = [\widetilde{\theta}_{S_j}]_{T_j}$, $\widetilde{Z}_S = [\widetilde{Z}_j^{S_j}]_{T_j}$, $\widetilde{\theta}_{S^c} = [\widetilde{\theta}_{S_j}]_{S_j \setminus T_j}$, and $\widetilde{Z}_{S^c} = [\widetilde{Z}_j^{S_j}]_{S_j \setminus T_j}$.

Lemma 6 *For any $S_j \in \{\pi_1, \pi_{1:2}, \dots, \pi_{1:j-1}\}$ and $\lambda_j \geq \frac{4C_x^2 \sqrt{2}(2-\alpha)}{\alpha} \frac{\log^2 \eta}{\kappa_1(n, p)}$ for some $\alpha \in (0, 1]$,*

$$P \left(\frac{\|W_j^{S_j}\|_{\infty}}{\lambda_j} \leq \frac{\alpha}{4(2-\alpha)} \right) \geq 1 - 2d \cdot \exp \left(-\frac{n}{\kappa_1(n, p)^2} \right).$$

where $\kappa_1(n, p)$ is an arbitrary function of n and p .

Lemma 7 Suppose that for all $S_j \in \{\pi_1, \pi_{1:2}, \dots, \pi_{1:j-1}\}$, $\|W_j^{S_j}\|_\infty \leq \frac{\lambda_j}{4}$. Then, for $\lambda_j \leq \frac{\rho_{\min}^2}{10C_x^2\rho_{\max}d\log^2\eta}$,

$$\|\tilde{\theta}_S - \theta_S^*\|_2 \leq \frac{5}{\rho_{\min}}\sqrt{d}\lambda_j$$

Lemma 8 Suppose that for all $S_j \in \{\pi_1, \pi_{1:2}, \dots, \pi_{1:j-1}\}$, $\|W_j^{S_j}\|_\infty \leq \frac{\lambda_j}{4}$. Then, for $\lambda_j \leq \frac{\alpha\rho_{\min}^2}{100C_x^2(2-\alpha)\rho_{\max}d\log^2\eta}$ and $\alpha \in (0, 1]$,

$$\frac{\|R_j^{S_j}\|_\infty}{\lambda_j} \leq \frac{\alpha}{4(2-\alpha)}$$

The rest of the proof is straightforward using Lemmas 6, 7, and 8. Consider the choice of regularization parameter $\lambda_{j0} = \frac{4\sqrt{2}C_x^2(2-\alpha)}{\alpha} \frac{\log^2\eta}{\kappa_1(n,p)}$, where $\kappa_1(n,p) \geq \frac{4\sqrt{2}C_x^4 \cdot 10^2(2-\alpha)^2}{\alpha^2} \frac{\rho_{\max}}{\rho_{\min}^2} d\log^4\eta$ ensuring that $\frac{4C_x^2\sqrt{2}(2-\alpha)}{\alpha} \frac{\log^2\eta}{\kappa_1(n,p)} \leq \lambda_{j0} \leq \frac{\alpha\rho_{\min}^2}{10^2C_x^2(2-\alpha)\rho_{\max}d\log^2\eta}$ for any $\alpha \in (0, 1]$. Hence, if we set $\kappa_1(n,p) = C_{\max}d\log^4\eta$ where $C_{\max} = \frac{4\sqrt{2} \cdot 10^2 C_x^4 \cdot (2-\alpha)^2}{\alpha^2} \frac{\rho_{\max}}{\rho_{\min}^2}$, then all conditions for Lemma 6, 7, and 8 are satisfied. Therefore,

$$\|\tilde{Z}_{S^c}\|_\infty \leq (1-\alpha) + (2-\alpha) \left[\frac{\|W_j^{S_j}\|_\infty}{\lambda_j} + \frac{\|R_j^{S_j}\|_\infty}{\lambda_j} \right] \leq (1-\alpha) + \frac{\alpha}{4} + \frac{\alpha}{4} < 1, \quad (20)$$

with a probability of at least $1 - 2d \cdot \exp\left(-\frac{n}{\kappa_1(n,p)^2}\right) = 1 - 2d \cdot \exp\left(-\frac{n}{C_{\max}^2 d^2 \log^8\eta}\right)$.

Proposition 9 Suppose that, for any $j \in V$, partial ordering (π_1, \dots, π_j) is correctly estimated. If $\min_{t \in S} [\theta_S^*]_t \geq \frac{10}{\rho_{\min}}\sqrt{d}\lambda_j$ for all $j \in V$,

$$\text{supp}(\hat{\theta}_{S_j}) = Pa(j).$$

Proposition 9 guarantees that ℓ_1 -regularized likelihood regression recovers the parents for each node with a high probability. Since there are p regression problems, for any $\epsilon > 0$, there exists a positive constant $C_\epsilon > 0$ such that if $n \geq C_\epsilon(\kappa_1(n,p)^2 \log p)$ for $\kappa_1(n,p) \geq C_{\max}d\log^4\eta$,

$$P(\hat{G} = G) \geq 2dp \cdot \exp\left(-\frac{n}{\kappa_1(n,p)^2}\right) \geq 1 - 2dp \cdot \exp(-C_\epsilon \log p) \geq 1 - \epsilon. \quad \blacksquare$$

Appendix D. Proof for Theorem 3: Ordering Recovery

Proof We begin by reintroducing some necessary notations and definitions to make the proof concise. Without loss of generality, assume that the true ordering is unique and $\pi = (\pi_1, \dots, \pi_p) = (1, 2, \dots, p)$. For notational convenience, we define $X_{1:j} = (X_{\pi_1}, X_{\pi_2}, \dots, X_{\pi_j})$

$= (X_1, X_2, \dots, X_j)$ and $X_{1:0} = \emptyset$. We restate the moments ratio scores for a node k and the j th element of the ordering:

$$\mathcal{S}(j, k) := \frac{\mathbb{E}(X_k^2)}{\mathbb{E}(f(\mathbb{E}(X_k | X_{1:(j-1)}))}) \quad \text{and} \quad \widehat{\mathcal{S}}(j, k) := \frac{\widehat{\mathbb{E}}(X_k^2)}{\widehat{\mathbb{E}}(f(\widehat{\mathbb{E}}(X_k | X_{\widehat{\pi}_{1:(j-1)}}))}),$$

where $f(\mu) := \mu + \mu^2$, $\mathbb{E}(X_k | X_{S_k}) = \exp(\theta_k^* + \sum_{t \in S_k} \theta_{kt}^* X_t)$, and $\widehat{\mathbb{E}}(X_k | X_{S_k}) = \exp(\widehat{\theta}_k + \sum_{t \in S_k} \widehat{\theta}_{kt} X_t)$ where $\theta_{S_k}^* = (\theta_k^*, \theta_{kt}^*)$ and $\widehat{\theta}_{S_k} = (\widehat{\theta}_k, \widehat{\theta}_{kt})$ are the solutions of the problem (11) and of the ℓ_1 -regularized GLM (9), respectively. In addition, we use the unbiased method-of-moment estimator for a marginal expectation, $\widehat{\mathbb{E}}(X_k^2) = \frac{1}{n} \sum_{i=1}^n (X_k^{(i)})^2$ and $\widehat{\mathbb{E}}(f(\widehat{\mathbb{E}}(X_k | X_S))) = \frac{1}{n} \sum_{i=1}^n f(\exp(\widehat{\theta}_k + \sum_{t \in S_k} \widehat{\theta}_{kt} X_t^{(i)}))$.

We define the following necessary events: For each node $j \in V$, $S_j \in \{\pi_1, \pi_{1:2}, \dots, \pi_{1:(j-1)}\}$ and any $\epsilon_1 > 0$;

$$\begin{aligned} \zeta_1 &:= \left\{ \max_{j=1, \dots, p-1} \max_{k=j, \dots, p} \left| \mathcal{S}(j, \pi_k) - \widehat{\mathcal{S}}(j, \pi_k) \right| > \frac{M_{\min}}{2} \right\}, \\ \zeta_2 &:= \left\{ \max_{j \in V} \left| \widehat{\mathbb{E}}(X_j^2) - \mathbb{E}(X_j^2) \right| < \epsilon_1 \right\}, \\ \zeta_3 &:= \left\{ \max_{j \in V} \left| \widehat{\mathbb{E}} \left(f \left(\widehat{\mathbb{E}}(X_j | X_{S_j}) \right) \right) - \mathbb{E} \left(f \left(\widehat{\mathbb{E}}(X_j | X_{S_j}) \right) \right) \right| < \epsilon_1 \right\}, \\ \zeta_4 &:= \left\{ \max_{j \in V} \left| \mathbb{E} \left(f \left(\widehat{\mathbb{E}}(X_j | X_{S_j}) \right) \right) - \mathbb{E} \left(f \left(\mathbb{E}(X_j | X_{S_j}) \right) \right) \right| < \epsilon_1 \right\}. \end{aligned}$$

We begin by proving that our algorithm recovers the ordering of a Poisson SEM in the high-dimensional setting. The probability that ordering is correctly estimated from our method can be written as

$$\begin{aligned} &P(\widehat{\pi} = \pi) \\ &= P\left(\widehat{\mathcal{S}}(1, \pi_1) < \min_{j=2, \dots, p} \widehat{\mathcal{S}}(1, \pi_j), \widehat{\mathcal{S}}(2, \pi_2) < \min_{j=3, \dots, p} \widehat{\mathcal{S}}(2, \pi_j), \dots, \widehat{\mathcal{S}}(p-1, \pi_{p-1}) < \widehat{\mathcal{S}}(p-1, \pi_p) \right) \\ &= P\left(\min_{j=1, \dots, p-1} \min_{k=j+1, \dots, p} \widehat{\mathcal{S}}(j, \pi_k) - \widehat{\mathcal{S}}(j, \pi_j) > 0 \right) \\ &= P\left(\min_{\substack{j=1, \dots, p-1 \\ k=j+1, \dots, p}} \left\{ \left(\mathcal{S}(j, \pi_k) - \mathcal{S}(j, \pi_j) \right) - \left(\mathcal{S}(j, \pi_k) - \widehat{\mathcal{S}}(j, \pi_k) \right) + \left(\mathcal{S}(j, \pi_j) - \widehat{\mathcal{S}}(j, \pi_j) \right) \right\} > 0 \right) \\ &\geq P\left(\min_{\substack{j=1, \dots, p-1 \\ k=j+1, \dots, p}} \left\{ \left(\mathcal{S}(j, \pi_k) - \mathcal{S}(j, \pi_j) \right) \right\} > M_{\min}, \text{ and } \max_{\substack{j=1, \dots, p-1 \\ k=j, \dots, p}} \left| \mathcal{S}(j, \pi_k) - \widehat{\mathcal{S}}(j, \pi_k) \right| < \frac{M_{\min}}{2} \right). \end{aligned}$$

The first term in the above probability is always satisfied because $\mathcal{S}(j, \pi_k) - \mathcal{S}(j, \pi_j) > (1 + M_{\min}) - 1 = M_{\min}$ from Assumption 4. Hence, the lower bound of the probability that

ordering is correctly estimated using our method is reduced to

$$\begin{aligned}
 P(\widehat{\pi} = \pi) &\geq P\left(\max_{j=1,\dots,p-1} \max_{k=j,\dots,p} \left| \mathcal{S}(j, \pi_k) - \widehat{\mathcal{S}}(j, \pi_k) \right| < \frac{M_{\min}}{2}\right) \\
 &= 1 - P(\zeta_1) \\
 &= 1 - P(\zeta_1 \mid \zeta_2, \zeta_3, \zeta_4) P(\zeta_2, \zeta_3, \zeta_4) - P(\zeta_1 \mid (\zeta_2, \zeta_3, \zeta_4)^c) P((\zeta_2, \zeta_3, \zeta_4)^c) \\
 &\geq 1 - P(\zeta_1 \mid \zeta_2, \zeta_3, \zeta_4) - P((\zeta_2, \zeta_3, \zeta_4)^c) \\
 &\geq 1 - \underbrace{P(\zeta_1 \mid \zeta_2, \zeta_3, \zeta_4)}_{\text{Lem 10}} - \underbrace{P(\zeta_2^c) - P(\zeta_3^c) - P(\zeta_4^c)}_{\text{Lem 11}}. \tag{21}
 \end{aligned}$$

Next, we introduce the following two lemmas to show the lower bound of the probability in (21) as a function of the triple (n, p, d) :

Lemma 10 *Given the sets $\zeta_2, \zeta_3, \zeta_4$ and under Assumption 4, $P(\zeta_1 \mid \zeta_2, \zeta_3, \zeta_4) = 0$ if for some small ϵ_1 such that for any $S_j \in \{\pi_1, \pi_{1:2}, \dots, \pi_{1:(j-1)}\}$,*

$$\epsilon_1 < \min \left\{ \frac{\mathbb{E}(X_j^2) M_{\min}}{2(M_{\min} + 3)(M_{\min} + 1)}, \frac{M_{\min} \mathbb{E}(f(\mathbb{E}(X_j \mid X_{S_j})))^2}{6 \mathbb{E}(X_j^2)} \right\},$$

where $f(\mu) = \mu + \mu^2$.

The condition in Lemma 10 implies that if ϵ_1 is sufficiently small, the estimated score is close to the true score value.

The second lemma shows the error bound for the consistency of the estimators.

Lemma 11 *For any $\epsilon_1 > 0$ and*

$$(i) \text{ For } \zeta_2, P(\zeta_2^c) \leq 1 - 2 \cdot p \cdot \exp \left\{ -\frac{n\epsilon_1^2}{2C_x^2 \log^2 \eta} \right\}.$$

(ii) *For ζ_3 , there exist some positive constants C_{\max} and D_{\max} such that*

$$P(\zeta_3^c) \leq 1 - 2 \cdot p \cdot d \cdot \exp \left(-\frac{n}{\kappa_1(n,p)^2} \right) - 2 \cdot p \cdot \exp \left\{ -\frac{n\epsilon_1^2}{D_{\max} \log^4 \eta} \right\}.$$

where $\kappa_1(n, p) \geq C_{\max} d \log^4 \eta$.

(iii) *For ζ_4 , $P(\zeta_4^c) = 0$.*

Therefore, we complete the proof: our method recovers the true ordering at least of

$$P(\widehat{\pi} = \pi) \geq 1 - C_1 p \cdot d \cdot \exp \left(-C_2 \frac{n}{\kappa_1(n, p)^2} \right).$$

for $\kappa_1(n, p) \geq C_{\max} d \log^4 \eta$, and some positive constants C_1 and C_2 ■

Appendix E. Proposition 12

We begin by introducing an important proposition to control the tail behavior for the distribution of each node, which are required to prove the lemmas.

Proposition 12 *For given $j \in V$ and $S_j \in \{\pi_1, \pi_{1:2}, \dots, \pi_{1:(j-1)}\}$, the solution $\hat{\theta}_{S_j}$ in Equation (15) satisfies*

$$\frac{1}{n} \sum_{i=1}^n \exp \left(\langle \hat{\theta}_{S_j}, X_{S_j}^{(i)} \rangle \right) < C_x \log \eta.$$

where $C_x > 2$ is a constant in Assumption 3.

Proof By the first-order optimality condition of $\mathcal{L}_j^{S_j}(\theta_{S_j}, X^{1:n})$ in Equation (15), we have

$$\begin{aligned} \sum_{i=1}^n X_j^{(i)} &= \sum_{i=1}^n \exp(\langle \hat{\theta}_{S_j}, X_{S_j}^{(i)} \rangle) \\ \sum_{i=1}^n X_j^{(i)} X_k^{(i)} &= \sum_{i=1}^n \exp \left(\langle \hat{\theta}_{S_j}, X_{S_j}^{(i)} \rangle X_k^{(i)} \right) + \lambda_j \text{sign}([\hat{\theta}_{S_j}]_k). \end{aligned}$$

By Assumption 3, we have

$$\frac{1}{n} \sum_{i=1}^n \exp \left(\langle \hat{\theta}_{S_j}, X_{S_j}^{(i)} \rangle \right) \leq C_x \log \eta \iff \frac{1}{n} \sum_{i=1}^n X_j^{(i)} \leq C_x \log \eta. \quad \blacksquare$$

Appendix F. Proof for Propositions 5 and 9

F.1. Proof for Proposition 5

Proof We note that $\tilde{\theta}_{S^c} = (0, 0, \dots, 0)^T \in \mathbb{R}^{|S^c|}$ in our primal-dual construction. To improve readability, we let $\theta_S = [\theta_{S_j}]_{T_j}$, $\theta_{S^c} = [\theta_{S_j}]_{S_j \setminus T_j}$, and $A_S = [A_j^{S_j}]_{T_j}$ and $A_{S^c} = [A_j^{S_j}]_{S_j \setminus T_j}$. With these notations, W_S and R_S are sub-vectors of $W_j^{S_j}$ and $R_j^{S_j}$ corresponding to variables X_S , respectively.

We can restate condition (17) in block form as follows:

$$\begin{aligned} Q_{S^c S}[\tilde{\theta}_S - \theta_S^*] &= W_{S^c} - \lambda_j \tilde{Z}_{S^c} + R_{S^c}, \\ Q_{SS}[\tilde{\theta}_S - \theta_S^*] &= W_S - \lambda_j \tilde{Z}_S + R_S. \end{aligned}$$

Since Q_{SS} is invertible, the above equations can be rewritten as

$$Q_{S^c S} Q_{SS}^{-1} [W_S - \lambda_j \tilde{Z}_S - R_S] = W_{S^c} - \lambda_j \tilde{Z}_{S^c} - R_{S^c}.$$

Therefore,

$$[W_{S^c} - R_{S^c}] - Q_{S^c S} Q_{SS}^{-1} [W_S - R_S] + \lambda_j Q_{S^c S} Q_{SS}^{-1} \tilde{Z}_S = \lambda_j \tilde{Z}_{S^c}.$$

Taking the ℓ_∞ norm of both sides yields

$$\|\tilde{Z}_{S^c}\|_\infty \leq \|Q_{S^c S} Q_{SS}^{-1}\|_\infty \left[\frac{\|W_S\|_\infty}{\lambda_j} + \frac{\|R_S\|_\infty}{\lambda_j} + 1 \right] + \frac{\|W_{S^c}\|_\infty}{\lambda_j} + \frac{\|R_{S^c}\|_\infty}{\lambda_j}.$$

Recalling Assumption (2), we obtain $\|Q_{S^c S} Q_{SS}^{-1}\|_\infty \leq (1 - \alpha)$, and hence, we have

$$\begin{aligned} \|\tilde{Z}_{S^c}\|_\infty &\leq (1 - \alpha) \left[\frac{\|W_S\|_\infty}{\lambda_j} + \frac{\|R_S\|_\infty}{\lambda_j} + 1 \right] + \frac{\|W_{S^c}\|_\infty}{\lambda_j} + \frac{\|R_{S^c}\|_\infty}{\lambda_j} \\ &\leq (1 - \alpha) + (2 - \alpha) \left[\frac{\|W_j^{S_j}\|_\infty}{\lambda_j} + \frac{\|R_j^{S_j}\|_\infty}{\lambda_j} \right]. \end{aligned}$$

If both $\|W_j^{S_j}\|_\infty$ and $\|R_j^{S_j}\|_\infty$ are less than $\frac{\lambda_j \alpha}{4(2-\alpha)}$, as assumed, then

$$\|\tilde{Z}_{S^c}\|_\infty \leq (1 - \alpha) + \frac{\alpha}{2} < 1. \quad \blacksquare$$

F.2. Proof for Proposition 9

Proof To prove the support of $\hat{\theta}_S$ is not strictly subset the true support X_S , it is sufficient to show that the maximum bias is bounded:

$$\|\hat{\theta}_S - \theta_S^*\|_\infty \leq \frac{\min_{t \in S} [\theta_S^*]_t}{2}.$$

From Lemma 7, we have, with a high probability,

$$\|\hat{\theta}_S - \theta_S^*\|_\infty \leq \|\hat{\theta}_S - \theta_S^*\|_2 \leq \frac{5}{\rho_{\min}} \sqrt{d} \lambda_j.$$

Therefore, if $\min_{t \in S} [\theta_S^*]_t \geq \frac{10}{\rho_{\min}} \sqrt{d} \lambda_j$,

$$\|\hat{\theta}_S - \theta_S^*\|_\infty \leq \frac{\min_{t \in S} [\theta_S^*]_t}{2}. \quad \blacksquare$$

Appendix G. Proof for Lemmas

G.1. Proof for Lemma 4

Proof This lemma can be proved by the same manner developed for the special cases (Wainwright et al., 2006; Ravikumar et al., 2011). In addition, this proof is directly from Lemma 8 in Yang et al. (2015). And, we restate the proof in our framework. The main idea of the proof is the *primal-dual-witness* method which asserts that there is a solution to the dual problem $\tilde{\theta}_{S_j} = \hat{\theta}_{S_j}$ if the following Karush-Kuhn-Tucker (KKT) conditions are satisfied.

- (a) We define $\tilde{\theta}_{S_j} \in \Theta_{S_j}$, where $\Theta_{S_j} = \{\theta \in \mathbb{R}^{|S_j|} : \theta_{S^c} = 0\}$ is the solution to the following optimization problem:

$$\tilde{\theta}_{S_j} := \arg \min_{\theta \in \Theta_{S_j}} \mathcal{L}_j^{S_j}(\theta, \lambda_j) = \arg \min_{\theta \in \Theta_{S_j}} \{\ell_j^{S_j}(\theta; X^{1:n}) + \lambda_j \|\theta\|_1\}. \quad (22)$$

- (b) Define $\tilde{Z}_j^{S_j}$ to be a sub-differential for the regularizer $\|\cdot\|_1$ evaluated at $\tilde{\theta}_{S_j}$. For any $t \in T_j$ in Equation (12), $[\tilde{Z}_j^{S_j}]_t = \text{sign}([\tilde{\theta}_{S_j}]_t)$.

- (c) For any $t \notin T_j$, $|\tilde{Z}_j^{S_j}]_t| < 1$.

If conditions (a) to (c) are satisfied, $\tilde{\theta}_{S_j} = \hat{\theta}_{S_j}$ meaning that the solution to unrestricted problem (15) is the same as the solution to restricted problem (22) (See Ravikumar et al., 2011 for details).

In addition, if the sub-matrix of the Hessian $Q_{S_j}^{S_j}$ is invertible, restricted problem (22) is strictly convex, and hence, $\tilde{\theta}_{S_j}$ is unique. \blacksquare

G.2. Proof for Lemma 6

Proof In order to improve readability, we omit the superscript S_j if it is understood (i.e., $W_j = W_j^{S_j}$). Each entry of the sample score function W_j in Equation (18) has the form $W_{jt} = \frac{1}{n} \sum_{i=1}^n W_{jt}^{(i)}$ for any $t \in S := \{k \in S_j \mid [\theta_{S_j}^*]_k \neq 0\}$. In addition, $W_{jt} = 0$ for all $t \notin S$, since $[\theta_{S_j}^*]_t = 0$ by the definition of S .

Hence simple calculation yields that, for any $t \in S$ and $i \in \{1, 2, \dots, n\}$,

$$W_{jt}^{(i)} = X_t^{(i)} X_j^{(i)} - \exp(\langle \theta_S^*, X_S^{(i)} \rangle) X_t^{(i)},$$

and $(|W_{jt}^{(i)}|)_{i=1}^n$ has mean 0 by the first-order optimality condition, $\mathbb{E}(X_j) = \mathbb{E}(\exp(\langle \theta_S^*, X_S \rangle))$.

Now, we show that $W_{jt}^{(i)}$ is bounded with a high probability given Assumption 3 by using Hoeffding's inequality. The both terms are bounded above $C_x^2 \log^2 \eta$ by Assumption 3. Therefore, $|W_{jt}^{(i)}|$ is bounded by $2C_x^2 \log^2 \eta$.

Applying the union bound and Hoeffding's inequality, we have

$$P(\|W_j\|_\infty > \delta) \leq d \cdot \max_{t \in S} P(|W_{jt}| > \delta) \leq 2d \cdot \exp\left(-\frac{2n\delta^2}{4C_x^4 \log^4 \eta}\right).$$

Suppose that $\delta = \frac{\lambda_j \alpha}{4(2-\alpha)}$ and $\lambda_j \geq \frac{4(2-\alpha)}{\alpha} \frac{2C_x^2 \log^2 \eta}{\sqrt{2\kappa_1(n,p)}}$. Then, we complete the proof:

$$P\left(\frac{\|W_j\|_\infty}{\lambda_j} > \frac{\alpha}{4(2-\alpha)}\right) \leq 2d \cdot \exp\left(-\frac{\alpha^2}{16(2-\alpha)^2} \frac{2n\lambda_j^2}{4C_x^4 \log^4 \eta}\right) \leq 2d \cdot \exp\left(-\frac{n}{\kappa_1(n,p)^2}\right). \quad (23)$$

\blacksquare

G.3. Proof for Lemma 7

Proof In order to establish error bound $\|\tilde{\theta}_S - \theta_S^*\| \leq B$ for some radius B , several works (Meinshausen and Bühlmann, 2006; Wainwright et al., 2006; Ravikumar et al., 2011; Yang et al., 2015; Park and Raskutti, 2018) already proved that it suffices to show $F(u_S) > 0$ for all $u_S := \tilde{\theta}_S - \theta_S^*$ such that $\|u_S\|_2 = B$ where

$$F(a) := \ell_j(\theta_S^* + a; X^{1:n}) - \ell_j(\theta_S^*; X^{1:n}) + \lambda_j(\|\theta_S^* + a\|_1 - \|\theta_S^*\|_1). \quad (24)$$

More specifically, since u_S is the minimizer of F and $F(0) = 0$ by the construction of Equation (24), $F(u_S) \leq 0$. Note that F is convex, and therefore we have $F(u_S) < 0$.

Next we claim that $\|u_S\|_2 \leq B$. In fact, if u_S lies outside the ball of radius B , then there exists $v \in (0, 1)$ such that the convex combination $v \cdot u_S + (1 - v) \cdot 0$ would lie on the boundary of the ball. However it contradicts the assumed strict positivity of F on the boundary because, by convexity,

$$F(v \cdot u_S + (1 - v) \cdot 0) \leq v \cdot F(u_S) + (1 - v) \cdot 0 \leq 0. \quad (25)$$

Thus it suffices to establish strict positivity of F on the boundary of the ball with radius $B := M_B \lambda_j \sqrt{d}$ where $M_B > 0$ is a parameter to be chosen later in the proof. Let $u_S \in \mathbb{R}^{|S|}$ be an arbitrary vector with $\|u_S\|_2 = B$. By the Taylor series expansion of F in (24),

$$F(u_S) = (W_S)^T u_S + u_S^T [\nabla^2 \ell_j(\theta_S^* + v u_S; X^{1:n})] u_S + \lambda_j(\|\theta_S^* + u_S\|_1 - \|\theta_S^*\|_1), \quad (26)$$

for some $v \in [0, 1]$.

The first term in Equation (26) has the following bound: applying $\|W_S\|_\infty \leq \frac{\lambda_j}{4}$ by assumption and $\|u_S\|_1 \leq \sqrt{d} \|u_S\|_2 \leq \sqrt{d} \cdot B$,

$$|(W_S)^T u_S| \leq \|W_S\|_\infty \|u_S\|_1 \leq \|W_S\|_\infty \sqrt{d} \|u_S\|_2 \leq (\lambda_j \sqrt{d})^2 \frac{M_B}{4}.$$

The third term in Equation (26) has the following bound: Applying the triangle inequality,

$$\lambda_j(\|\theta_S^* + u_S\|_1 - \|\theta_S^*\|_1) \geq -\lambda_j \|u_S\|_1 \geq -\lambda_j \sqrt{d} \|u_S\|_2 = -M_B (\lambda_j \sqrt{d})^2.$$

Now we show the bound for the second term using the minimum eigenvalue of a matrix $\nabla^2 \ell_j(\theta_S^* + v u_S)$:

$$\begin{aligned} q^* &:= \lambda_{\min}(\nabla^2 \ell_j(\theta_S^* + v u_S)) \\ &\geq \min_{v \in [0, 1]} \lambda_{\min}(\nabla^2 \ell_j(\theta_S^* + v u_S)) \\ &\geq \lambda_{\min}(\nabla^2 \ell_j(\theta_S^*)) - \max_{v \in [0, 1]} \left\| \frac{1}{n} \sum_{i=1}^n \exp(\langle \theta_S^* + v u_S, X_S^{(i)} \rangle) u_S^T X_S^{(i)} X_S^{(i)} (X_S^{(i)})^T \right\|_2 \\ &\geq \rho_{\min} - \max_{v \in [0, 1]} \max_{y: \|y\|_2=1} \frac{1}{n} \sum_{i=1}^n \exp(\langle \theta_S^* + v u_S, X_S^{(i)} \rangle) \cdot (y^T X_S^{(i)})^2 \cdot |u_S^T X_S^{(i)}|. \end{aligned} \quad (27)$$

We first show the bound of the first term in Equation (27): Note that $\theta_S^* + vu_S$ is a linear (convex) combination of θ_S^* and $\tilde{\theta}_S$. Hence, by Assumption 3 and Proposition 12, we obtain

$$\frac{1}{n} \sum_{i=1}^n \exp(\langle \theta_S^* + vu_S, X_S^{(i)} \rangle) \leq C_x \log \eta.$$

Now, we bound the second term in Equation (27): Recall that $\|X_S^{(i)}\|_\infty \leq C_x \log \eta$ for all i by Assumption 3. Recall $[u_S]_t = 0$ for $t \notin S$ by the primal-dual construction of (17). Applying $\|u_S\|_1 \leq \sqrt{d}\|u_S\|_2 \leq \sqrt{d} \cdot B$,

$$|u_S^T X_S^{(i)}| \leq C_x \log(\eta) \sqrt{d} \|u_S\|_2 \leq C_x \log(\eta) \cdot M_B \lambda_j d.$$

Lastly, it is clear that $\max_{y: \|y\|_2=1} (y^T X_S^{(i)})^2 \leq \rho_{\max}$ by the definition of the maximum eigenvalue and Assumption 1. Together with the above bounds, we obtain

$$P(q^* \leq \rho_{\min} - C_x^2 M_B \rho_{\max} d \lambda_j \log^2 \eta) \leq M \eta^{-2}.$$

For $\lambda_j \leq \frac{\rho_{\min}}{2C_x^2 M_B \rho_{\max} d \log^2 \eta}$, we have $q^* \geq \frac{\rho_{\min}}{2}$ with a high probability. Therefore,

$$F(u) \geq (\lambda_j \sqrt{n})^2 \left\{ -\frac{1}{4} M_B + \frac{\rho_{\min}}{2} M_B^2 - M_B \right\},$$

which is strictly positive for $M_B = \frac{5}{\rho_{\min}}$. Therefore, for $\lambda_j \leq \frac{\rho_{\min}^2}{10C_x^2 \rho_{\max} d \log^2 \eta}$,

$$\|\tilde{\theta}_S - \theta_S^*\|_2 \leq \frac{5}{\rho_{\min}} \sqrt{d} \lambda_j.$$

■

G.4. Proof for Lemma 8

Proof To improve readability, we use $R_S = [R_j^{S_j}]_S$ where $S := \{k \in S_j \mid [\theta_{S_j}^*]_k \neq 0\}$.

Then, each entry of $R_j^{S_j}$ in Equation (19) has the form $R_{jk} = \frac{1}{n} \sum_{i=1}^n R_{jk}^{(i)}$ for any $k \in S_j$, and it can be expressed as

$$\begin{aligned} R_{jk} &= \frac{1}{n} \sum_{i=1}^n [\nabla^2 \ell_j(\theta_{S_j}^*; X^{1:n}) - \nabla^2 \ell_j(\tilde{\theta}_{S_j}; X^{1:n})]_k^T (\tilde{\theta}_{S_j} - \theta_{S_j}^*) \\ &= \frac{1}{n} \sum_{i=1}^n \left[\exp(\langle \theta_{S_j}^*, X_S^{(i)} \rangle) - \exp(\langle \tilde{\theta}_S, X_S^{(i)} \rangle) \right] \left[X_S^{(i)} (X_S^{(i)})^T \right]_k^T (\tilde{\theta}_S - \theta_{S_j}^*) \end{aligned}$$

for $\tilde{\theta}_S$, which is a point on the line between $\tilde{\theta}_S$ and θ_S^* (i.e., $\tilde{\theta}_S^{(t)} = v \cdot \tilde{\theta}_S + (1-v) \cdot \theta_S^*$ for some $v \in [0, 1]$). The second equality holds because $\theta_{S_j^c}^* = \tilde{\theta}_{S_j^c} = (0, 0, \dots, 0) \in \mathbb{R}^{|S^c|}$.

Applying the mean value theorem again, we have,

$$R_{jk} = \frac{1}{n} \sum_{i=1}^n \left\{ \exp(\langle \tilde{\theta}_S, X_S^{(i)} \rangle) X_k^{(i)} \right\} \left\{ v (\tilde{\theta}_S - \theta_S^*)^T X_S^{(i)} (X_S^{(i)})^T (\tilde{\theta}_S - \theta_S^*) \right\}$$

for $\bar{\theta}_{S_j}$ which is a point on the line between $\bar{\theta}_{S_j}$ and $\theta_{S_j}^*$ (i.e., $\bar{\theta}_{S_j} = v \cdot \bar{\theta}_{S_j} + (1-v) \cdot \theta_{S_j}^*$ for $v \in [0, 1]$).

Note that $\bar{\theta}_{S_j}$ is a linear (convex) combination of $\theta_{S_j}^*$ and $\tilde{\theta}_S$. Hence, from Assumption 3 and Proposition 12, we obtain,

$$\frac{1}{n} \sum_{i=1}^n \exp \left(\left\langle \bar{\theta}_{S_j}, X_{S_j}^{(i)} \right\rangle \right) \leq C_x \log \eta, \quad \text{and} \quad \max_{i,j} X_j^{(i)} < C_x \log \eta.$$

Therefore, we have $|R_{jk}| \leq \rho_{\max} C_x^2 \log^2 \eta \|\tilde{\theta}_S - \theta_S^*\|_2^2$ for all $j, k \in V$.

In Section G.3, we showed that $\|\tilde{\theta}_S - \theta_S^*\|_2 \leq \frac{5}{\rho_{\min}} \sqrt{d} \lambda_j$ for $\lambda_j \leq \frac{\rho_{\min}^2}{10C_x^2 \rho_{\max} d \log^2 \eta}$. Therefore, if $\lambda_j \leq \frac{\rho_{\min}^2}{25 \cdot C_x^2 \rho_{\max} d \log^2 \eta} \frac{\alpha}{4(2-\alpha)}$, we obtain,

$$P \left(\|R_j\|_{\infty} > \frac{\alpha}{4(2-\alpha)} \lambda_j \right) \leq P \left(\|R_j\|_{\infty} > 25 C_x^2 \lambda_j^2 \frac{\rho_{\max}}{\rho_{\min}^2} d \log^2 \eta \right) = 0.$$

Therefore, we have,

$$\|R_j\|_{\infty} \leq \frac{\alpha}{4(2-\alpha)} \lambda_j$$

■

G.5. Proof for Lemma 10

Proof Conditioning on the sets ζ_2, ζ_3 , and ζ_4 , we provide the following results for different two cases:

(i) For any $j \in \{1, 2, \dots, p-1\}$, and $X_S = X_{1:(j-1)}$, we have $\frac{\mathbb{E}(X_j^2)}{\mathbb{E}(f(\mathbb{E}(X_j | X_S)))} = 1$. Therefore, for $k = \pi_j$, we have the following probability bound:

$$\begin{aligned} & P \left(|\hat{\mathcal{S}}(j, k) - \mathcal{S}(j, k)| < \frac{M_{\min}}{2} \Big| \zeta_2, \zeta_3, \zeta_4 \right) \\ &= P \left(\left| \frac{\hat{\mathbb{E}}(X_k^2)}{\hat{\mathbb{E}}(f(\hat{\mathbb{E}}(X_k | X_S)))} - \frac{\mathbb{E}(X_k^2)}{\mathbb{E}(f(\mathbb{E}(X_k | X_S)))} \right| < \frac{M_{\min}}{2} \Big| \zeta_2, \zeta_3, \zeta_4 \right) \\ &\geq P \left(\frac{\mathbb{E}(X_k^2) + \epsilon_1}{\mathbb{E}(f(\mathbb{E}(X_k | X_S))) - 2\epsilon_1} - \frac{\mathbb{E}(X_k^2)}{\mathbb{E}(f(\mathbb{E}(X_k | X_S)))} < \frac{M_{\min}}{2} \text{ and} \right. \\ &\quad \left. \frac{\mathbb{E}(X_k^2)}{\mathbb{E}(f(\mathbb{E}(X_k | X_S)))} - \frac{\mathbb{E}(X_k^2) - \epsilon_1}{\mathbb{E}(f(\mathbb{E}(X_k | X_S))) + 2\epsilon_1} < \frac{M_{\min}}{2} \right) \\ &\geq P \left(\epsilon_1 < \frac{\mathbb{E}(f(\mathbb{E}(X_k | X_S))) M_{\min}}{2(M_{\min} + 3)} \right) \\ &\geq P \left(\epsilon_1 < \frac{\mathbb{E}(X_k^2) M_{\min}}{2(M_{\min} + 3)(M_{\min} + 1)} \right). \end{aligned}$$

(ii) For $j \in \{1, 2, \dots, p-1\}$, $k \in \{\pi_{j+1}, \dots, \pi_p\}$ having parent π_j , and $X_S = X_{1:(j-1)}$, we have $\mathbb{E}(X_k^2) > (1 + M_{\min}) \mathbb{E}(f(\mathbb{E}(X_k | X_S)))$ by Assumption 4. In addition, some elementary

but complicated computations yield

$$\begin{aligned}
 & P\left(\left|\widehat{\mathcal{S}}(j, k) - \mathcal{S}(j, k)\right| < \frac{M_{\min}}{2} \mid \zeta_2, \zeta_3, \zeta_4\right) \\
 & \geq P\left(\epsilon_1 < \frac{\mathbb{E}(f(\mathbb{E}(X_k \mid X_S)))^2 M_{\min}}{4\mathbb{E}(X_k^2) + 2\mathbb{E}(f(\mathbb{E}(X_k \mid X_S))) + 2\mathbb{E}(f(\mathbb{E}(X_k \mid X_S)))M_{\min}}\right) \\
 & \geq P\left(\epsilon_1 < \frac{\mathbb{E}(f(\mathbb{E}(X_k \mid X_S)))^2 M_{\min}(1 + M_{\min})}{4\mathbb{E}(X_k^2)(1 + M_{\min}) + 2\mathbb{E}(X_k^2) + 2M_{\min}\mathbb{E}(X_k^2)}\right) \\
 & \geq P\left(\epsilon_1 < \frac{\mathbb{E}(f(\mathbb{E}(X_k \mid X_S)))^2 M_{\min}(1 + M_{\min})}{6(1 + M_{\min})\mathbb{E}(X_k^2)}\right) \\
 & \geq P\left(\epsilon_1 < \frac{M_{\min}}{6} \frac{\mathbb{E}(f(\mathbb{E}(X_k \mid X_S)))^2}{\mathbb{E}(X_k^2)}\right).
 \end{aligned}$$

Therefore $P(\zeta_1 \mid \zeta_2, \zeta_3, \zeta_4) = 0$ if ϵ_1 is sufficiently small enough. For any node j , any set $S_j \in \{\pi_1, \pi_{1:2}, \dots, \pi_{1:(j-1)}\}$, and $k \in \{\pi_j, \pi_{j+1}, \dots, \pi_p\}$,

$$\epsilon_1 < \min \left\{ \frac{\mathbb{E}(X_k^2)M_{\min}}{2(M_{\min} + 3)(M_{\min} + 1)}, \frac{M_{\min}}{6} \frac{\mathbb{E}(f(\mathbb{E}(X_k \mid X_{S_j})))^2}{\mathbb{E}(X_k^2)} \right\}.$$

■

G.6. Proof for Lemma 11

The proof for Lemma 11 is closely related to the proof in Appendix C. Hence, for brevity, we do not present the details of the proof already shown in Appendix C.

$$(i) \quad P(\zeta_2^c) \leq 2p \cdot \exp \left\{ -\frac{n\epsilon_1^2}{2C_x^2 \log^2 \eta} \right\}.$$

Proof Using Hoeffding's inequality given Assumption 3, for any $\epsilon > 0$ and $j \in V$,

$$P\left(\left|\widehat{\mathbb{E}}(X_j^2) - \mathbb{E}(X_j^2)\right| > \epsilon_1\right) \leq 2 \cdot \exp \left\{ -\frac{n\epsilon_1^2}{2C_x^2 \log^2 \eta} \right\}. \quad (28)$$

Hence, using the union bound, we have

$$P\left(\max_{j \in V} \left|\widehat{\mathbb{E}}(X_j^2) - \mathbb{E}(X_j^2)\right| > \epsilon_1\right) \leq 2p \cdot \exp \left\{ -\frac{n\epsilon_1^2}{2C_x^2 \log^2 \eta} \right\}.$$

■

$$(ii) \quad P(\zeta_3^c) \leq 2p \cdot d \cdot \exp \left(-\frac{n}{\kappa_1(n, p)^2} \right) + 2p \cdot \exp \left\{ -\frac{n\epsilon_1^2}{D_{\max} \log^4 \eta} \right\} \text{ for some constants } D_{\max} > 0.$$

Proof We restate the condition in the set ζ_3 as

$$\left| \frac{1}{n} \sum_{i=1}^n f(\widehat{\mathbb{E}}(X_j \mid X_{S_j}^{(i)})) - \mathbb{E}(f(\mathbb{E}(X_j \mid X_{S_j}))) \right| < \epsilon_1.$$

In order to apply Hoeffding's inequality, we first show the bound for $\tilde{E}(X_j | X_{S_j})$. Recall that $[\theta^*]_{S^c}$ and $[\tilde{\theta}]_{S^c} = (0, 0, \dots, 0) \in \mathbb{R}^{|S^c|}$ by the definition of S , and $|S| \leq d$. In Appendix G.3, we showed that $\|\tilde{\theta}_S - \theta_S^*\|_2 \leq \frac{5}{\rho_{\min}} \sqrt{d} \lambda_j$ for $\lambda_j \leq \frac{\rho_{\min}^2}{10C_x^2 \rho_{\max} d \log^2 \eta}$ with a high probability. Therefore, given Assumption 3, for all $i \in \{1, 2, \dots, n\}$,

$$\begin{aligned} \exp(\langle \hat{\theta}_{S_j}, X_{S_j}^{(i)} \rangle) &= \exp(\langle \hat{\theta}_{S_j} - \theta_{S_j}^*, X_{S_j}^{(i)} \rangle) \cdot \exp(\langle \theta_{S_j}^*, X_{S_j}^{(i)} \rangle) \\ &\leq \exp(\|\hat{\theta}_S - \theta_S^*\|_2 \|X_S^{(i)}\|_2) \cdot \exp(\langle \theta_S^*, X_S^{(i)} \rangle) \\ &\leq \exp\left\{\frac{5C_x d \lambda_j}{\rho_{\min}} \|X_S^{(i)}\|_\infty\right\} \cdot \exp(\langle \theta_S^*, X_S^{(i)} \rangle) \\ &\leq \exp\left\{\frac{5C_x d \lambda_j}{\rho_{\min}} \log(\eta)\right\} \cdot C_x \log \eta \\ &\leq \exp\left\{\frac{\rho_{\min}}{2C_x \rho_{\max} \log \eta}\right\} \cdot C_x \log \eta. \end{aligned}$$

Therefore,

$$f(\hat{\mathbb{E}}(X_j^{(i)} | X_{S_j}^{(i)})) \leq C_x^2 \cdot \exp\left\{\frac{\rho_{\min}}{C_x \rho_{\max}}\right\} \log^2 \eta + C_x \cdot \exp\left\{\frac{\rho_{\min}}{2C_x \rho_{\max}}\right\} \log \eta.$$

Hence there exists a positive constant $D_1 > 0$ such that for all $i \in \{1, 2, \dots, n\}$,

$$f(\hat{\mathbb{E}}(X_j | X_{S_j}^{(i)})) \leq D_1 \log^2 \eta.$$

Applying Hoeffding's inequality, for any $\epsilon_1 > 0$ and any $j \in V$,

$$P\left(\left|\hat{\mathbb{E}}(f(\tilde{E}(X_j | X_{S_j}))) - \mathbb{E}(f(\tilde{E}(X_j | X_{S_j})))\right| > \epsilon_1\right) \leq 2 \cdot \exp\left\{-\frac{2n\epsilon_1^2}{D_1^2 \log^4 \eta}\right\}. \quad (29)$$

Hence, there exist some constants $D_{\max} > 0$ such that

$$P\left(\max_{j \in V} \zeta_3^c\right) \leq 1 - 2p \cdot d \cdot \exp\left(-\frac{n}{\kappa_1(n, p)^2}\right) - 2p \cdot \exp\left\{-\frac{n\epsilon_1^2}{D_{\max} \log^4 \eta}\right\}.$$

■

(iii) $P(\zeta_4^c) = 0$.

Proof

We restate the condition in the set ζ_4 as

$$\left|\mathbb{E}\left(f(\mathbb{E}(X_j | X_{S_j})) - f(\hat{\mathbb{E}}(X_j | X_{S_j}))\right)\right| < \epsilon_1.$$

By the mean-value theorem, for some $v \in [0, 1]$,

$$\begin{aligned} & f(\widehat{\mathbb{E}}(X_j | X_{S_j})) - f(\mathbb{E}(X_j | X_{S_j})) \\ &= f'(v\widehat{\mathbb{E}}(X_j | X_{S_j}) + (1-v)\mathbb{E}(X_j | X_{S_j}))(\widehat{\mathbb{E}}(X_j | X_{S_j}) - \mathbb{E}(X_j | X_{S_j})) \\ &= 2(v\widehat{\mathbb{E}}(X_j | X_{S_j}) + (1-v)\mathbb{E}(X_j | X_{S_j}) + 1/2)(\widehat{\mathbb{E}}(X_j | X_{S_j}) - \mathbb{E}(X_j | X_{S_j})). \end{aligned}$$

Therefore,

$$\begin{aligned} & \mathbb{E}(f(\widehat{\mathbb{E}}(X_j | X_{S_j})) - f(\mathbb{E}(X_j | X_{S_j}))) \\ &= f'(v\widehat{\mathbb{E}}(X_j | X_{S_j}) + (1-v)\mathbb{E}(X_j | X_{S_j}))(\widehat{\mathbb{E}}(X_j | X_{S_j}) - \mathbb{E}(X_j | X_{S_j})) \\ &= 2(v\widehat{\mathbb{E}}(X_j | X_{S_j}) + (1-v)\mathbb{E}(X_j | X_{S_j}) + 1/2)(\widehat{\mathbb{E}}(X_j | X_{S_j}) - \mathbb{E}(X_j | X_{S_j})) \\ &\leq \max |2(v\widehat{\mathbb{E}}(X_j | X_{S_j}) + (1-v)\mathbb{E}(X_j | X_{S_j}) + 1/2)| \cdot \mathbb{E}(\widehat{\mathbb{E}}(X_j | X_{S_j}) - \mathbb{E}(X_j | X_{S_j})) \\ &= 0 \end{aligned}$$

In the same manner, $\mathbb{E}(f(\mathbb{E}(X_j | X_{S_j})) - f(\widehat{\mathbb{E}}(X_j | X_{S_j}))) \leq 0$. This completes the proof. ■

References

- Peter Bühlmann, Jonas Peters, Jan Ernest, et al. Cam: Causal additive models, high-dimensional order search and penalized regression. *The Annals of Statistics*, 42(6):2526–2556, 2014.
- David M Chickering, Dan Geiger, David Heckerman, et al. Learning bayesian networks is np-hard. Technical report, Citeseer, 1994.
- David Maxwell Chickering. Learning bayesian networks is np-complete. In *Learning from data*, pages 121–130. Springer, 1996.
- David Maxwell Chickering. Optimal structure identification with greedy search. *The Journal of Machine Learning Research*, 3:507–554, 2003.
- Kenji Doya. *Bayesian brain: Probabilistic approaches to neural coding*. MIT press, 2007.
- Mathias Drton, Wenyu Chen, and Y Samuel Wang. On causal discovery with equal variance assumption. *arXiv preprint arXiv:1807.03419*, 2018.
- Jerome Friedman, Trevor Hastie, and Rob Tibshirani. glmnet: Lasso and elastic-net regularized generalized linear models. *R package version*, 1(4), 2009.
- Jerome Friedman, Trevor Hastie, and Rob Tibshirani. Regularization paths for generalized linear models via coordinate descent. *Journal of statistical software*, 33(1):1, 2010.

- Nir Friedman, Michal Linial, Iftach Nachman, and Dana Pe'er. Using bayesian networks to analyze expression data. *Journal of computational biology*, 7(3-4):601–620, 2000.
- Michael Friendly. *Lahman: Sean 'Lahman' Baseball Database*, 2017. URL <https://CRAN.R-project.org/package=Lahman>. R package version 6.0-0.
- Morten Frydenberg. The chain graph markov property. *Scandinavian Journal of Statistics*, pages 333–353, 1990.
- Asish Ghoshal and Jean Honorio. Information-theoretic limits of bayesian network structure learning. In *Artificial Intelligence and Statistics*, pages 767–775, 2017a.
- Asish Ghoshal and Jean Honorio. Learning identifiable gaussian bayesian networks in polynomial time and sample complexity. In *Advances in Neural Information Processing Systems*, pages 6457–6466, 2017b.
- Asish Ghoshal and Jean Honorio. Learning linear structural equation models in polynomial time and sample complexity. In *Proceedings of the Twenty-First International Conference on Artificial Intelligence and Statistics*, pages 1466–1475, 2018.
- David Heckerman, Dan Geiger, and David M Chickering. Learning Bayesian networks: The combination of knowledge and statistical data. *Machine learning*, 20(3):197–243, 1995.
- Patrik O Hoyer, Dominik Janzing, Joris M Mooij, Jonas Peters, and Bernhard Schölkopf. Nonlinear causal discovery with additive noise models. In *Advances in Neural Information Processing Systems*, pages 689–696, 2009.
- David I Inouye, Eunho Yang, Genevera I Allen, and Pradeep Ravikumar. A review of multivariate distributions for count data derived from the poisson distribution. *Wiley Interdisciplinary Reviews: Computational Statistics*, 9(3), 2017.
- Jin Zhu Jia, Fang Xie, and Lihu Xu. Sparse poisson regression with penalized weighted score function. *arXiv preprint arXiv:1703.03965*, 2017.
- Jeffrey O Kephart and Steve R White. Directed-graph epidemiological models of computer viruses. In *Research in Security and Privacy, 1991. Proceedings., 1991 IEEE Computer Society Symposium on*, pages 343–359. IEEE, 1991.
- Steffen L Lauritzen. *Graphical models*. Oxford University Press, 1996.
- Po-Ling Loh and Peter Bühlmann. High-dimensional learning of linear causal networks via inverse covariance estimation. *The Journal of Machine Learning Research*, 15(1):3065–3105, 2014.
- Nicolai Meinshausen and Peter Bühlmann. High-dimensional graphs and variable selection with the lasso. *The Annals of Statistics*, pages 1436–1462, 2006.
- Gunwoong Park and Hyewon Park. Identifiability of generalized hypergeometric distribution (ghd) directed acyclic graphical models. In *Proceedings of Machine Learning Research*, pages 158–166, 2019.

- Gunwoong Park and Garvesh Raskutti. Learning large-scale poisson dag models based on overdispersion scoring. In *Advances in Neural Information Processing Systems*, pages 631–639, 2015.
- Gunwoong Park and Garvesh Raskutti. Learning quadratic variance function (qvf) dag models via overdispersion scoring (ods). *Journal of Machine Learning Research*, 18(224): 1–44, 2018.
- Jonas Peters and Peter Bühlmann. Identifiability of gaussian structural equation models with equal error variances. *Biometrika*, 101(1):219–228, 2014.
- Jonas Peters, Joris Mooij, Dominik Janzing, and Bernhard Schölkopf. Identifiability of causal graphs using functional models. pages 589–598, Corvallis, OR, USA, July 2011. AUAI Press.
- Pradeep Ravikumar, Martin J Wainwright, Garvesh Raskutti, Bin Yu, et al. High-dimensional covariance estimation by minimizing ℓ_1 -penalized log-determinant divergence. *Electronic Journal of Statistics*, 5:935–980, 2011.
- Shohei Shimizu, Patrik O Hoyer, Aapo Hyvärinen, and Antti Kerminen. A linear non-Gaussian acyclic model for causal discovery. *The Journal of Machine Learning Research*, 7:2003–2030, 2006.
- Shohei Shimizu, Takanori Inazumi, Yasuhiro Sogawa, Aapo Hyvärinen, Yoshinobu Kawahara, Takashi Washio, Patrik O Hoyer, and Kenneth Bollen. Directlingam: A direct method for learning a linear non-gaussian structural equation model. *Journal of Machine Learning Research*, 12(Apr):1225–1248, 2011.
- Peter Spirtes. Directed cyclic graphical representations of feedback models. In *Proceedings of the Eleventh conference on Uncertainty in artificial intelligence*, pages 491–498. Morgan Kaufmann Publishers Inc., 1995.
- Peter Spirtes, Clark N Glymour, and Richard Scheines. *Causation, prediction, and search*. MIT press, 2000.
- Ioannis Tsamardinos and Constantin F Aliferis. Towards principled feature selection: Relevancy, filters and wrappers. In *Proceedings of the ninth international workshop on Artificial Intelligence and Statistics*. Morgan Kaufmann Publishers: Key West, FL, USA, 2003.
- Ioannis Tsamardinos, Laura E Brown, and Constantin F Aliferis. The max-min hill-climbing bayesian network structure learning algorithm. *Machine learning*, 65(1):31–78, 2006.
- Caroline Uhler, Garvesh Raskutti, Peter Bühlmann, and Bin Yu. Geometry of the faithfulness assumption in causal inference. *The Annals of Statistics*, pages 436–463, 2013.
- Martin J Wainwright, John D Lafferty, and Pradeep K Ravikumar. High-dimensional graphical model selection using ℓ_1 -regularized logistic regression. In *Advances in Neural Information Processing Systems*, pages 1465–1472, 2006.

- Ying-Wooi Wan, Genevera I Allen, Yulia Baker, Eunho Yang, Pradeep Ravikumar, Matthew Anderson, and Zhandong Liu. Xmrf: an r package to fit markov networks to high-throughput genetics data. *BMC systems biology*, 10(3):69, 2016.
- Eunho Yang, Pradeep K Ravikumar, Genevera I Allen, and Zhandong Liu. On poisson graphical models. In *Advances in Neural Information Processing Systems*, pages 1718–1726, 2013.
- Eunho Yang, Pradeep Ravikumar, Genevera I Allen, and Zhandong Liu. Graphical models via univariate exponential family distributions. *Journal of Machine Learning Research*, 16(1):3813–3847, 2015.
- Jiji Zhang and Peter Spirtes. The three faces of faithfulness. *Synthese*, 193(4):1011–1027, 2016.