

The Reduced PC-Algorithm: Improved Causal Structure Learning in Large Random Networks

Arjun Sondhi*

ASONDHI@UW.EDU

Ali Shojaie

ASHOJAIE@UW.EDU

*Department of Biostatistics
University of Washington
Seattle, WA 98195-4322, USA*

Editor: Edo Airoldi

Abstract

We consider the task of estimating a high-dimensional directed acyclic graph, given observations from a linear structural equation model with arbitrary noise distribution. By exploiting properties of common random graphs, we develop a new algorithm that requires conditioning only on small sets of variables. The proposed algorithm, which is essentially a modified version of the PC-Algorithm, offers significant gains in both computational complexity and estimation accuracy. In particular, it results in more efficient and accurate estimation in large networks containing hub nodes, which are common in biological systems. We prove the consistency of the proposed algorithm, and show that it also requires a less stringent faithfulness assumption than the PC-Algorithm. Simulations in low and high-dimensional settings are used to illustrate these findings. An application to gene expression data suggests that the proposed algorithm can identify a greater number of clinically relevant genes than current methods.

Keywords: causal discovery, directed acyclic graphs, faithfulness, high dimensions, random graphs

1. Introduction

Directed acyclic graphs, or DAGs, are commonly used to represent causal relationships in complex biological systems. For example, in gene regulatory networks, directed edges represent regulatory interactions among genes, which are represented as nodes of the graph. While causal effects in biological networks can be accurately inferred from perturbation experiments (Shojaie et al., 2014)—including single or double gene knockouts (Schadt et al., 2005; Yoo et al., 2002)—these are costly to run. Estimating DAGs from observational data is thus an important exploratory task for generating causal hypotheses (Friedman et al., 2000; Jansen et al., 2003), and designing more efficient experiments.

Since the number of possible directed graphs grows super-exponentially in the number of nodes, estimation of DAGs is an NP-hard problem (Chickering, 1996). Methods of estimating DAGs from observational data can be broadly categorized into three classes. The first class, score-based methods, search over the space of all possible graphs, and

*. Arjun Sondhi is currently a Quantitative Scientist at Flatiron Health.

attempt to maximize a goodness-of-fit score, generally using a greedy algorithm. Examples include the hill climbing and tabu search algorithms (Scutari, 2010), as well as Bayesian approaches (Eaton and Murphy, 2012). The second class, constraint-based methods, first estimate the graph skeleton by performing conditional independence tests; the skeleton of a directed acyclic graph is the undirected graph obtained by removing the direction of edges. Information from conditional independence relations is then used to partially orient the edges of the graph. The resulting completed partially directed acyclic graph (CPDAG) represents the class of all directed acyclic graphs that are Markov equivalent, and therefore not distinguishable from observational data. The most well-known constraint-based method is the PC-Algorithm (Spirtes et al., 2000), which was popularized by Kalisch and Bühlmann (2007) for high-dimensional estimation. Finally, hybrid methods combine score and constraint-based approaches. For example, the Max-Min Hill Climbing algorithm (Tsamardinos et al., 2006) estimates the skeleton using a constraint-based method, and then orients the edges by using a greedy search algorithm. Sparsity-inducing regularization approaches have also been used to develop efficient hybrid methods (Schmidt et al., 2007).

Estimating DAGs in high dimensions introduces new computational and statistical challenges. Until recently, graph recovery in high dimensions was only established for the PC-Algorithm (Kalisch and Bühlmann, 2007) and hybrid constraint-based methods (Ha et al., 2016). While the recent work of Nandy et al. (2015) extends these results to score-based algorithms and their hybrid extensions, the PC-Algorithm is still considered a gold standard in high-dimensional sparse settings, due to its polynomial time complexity (Kalisch and Bühlmann, 2007). Moreover, constraint-based methods are indeed the building blocks of various hybrid approaches. Therefore, we primarily focus on constraint-based methods in this paper.

Despite its appealing features, the PC-Algorithm entails several properties that do not scale well to common high-dimensional settings. Specifically, large real-world biological systems are known to commonly be sparse graphs containing a small number of highly connected *hub* nodes (Chen and Sharp, 2004; Jeong et al., 2000). In such graphs, the average node degree will be small, while the maximum tends to be much larger, and increases with the number of nodes. This is particularly problematic for the PC-Algorithm, whose computational and sample complexities scale with the *maximum node degree* in the graph. Moreover, the recent work by Uhler et al. (2013) shows that the distributional assumptions required for high-dimensional consistency of the PC-Algorithm are overly restrictive in practice, and that the class of graphs which do not satisfy these assumptions is large. Although work has been done by Peters et al. (2014) on estimating DAGs defined over a larger class of probability models, the resulting methods also do not scale to high dimensions.

A common limitation of existing methods for estimating DAGs is that they do not account for structural properties of large networks. For instance, the PC-Algorithm only incorporates the sparsity of the network, by assuming that the maximum node degree in the graph skeleton is small relative to the sample size. However, real-world networks, particularly those observed in biology, are known to possess a number of other important properties. Of particular interest in estimating DAGs is the so-called *local separation property* of large networks (Anandkumar et al., 2012), which implies that the number of short paths between any two nodes is bounded. This property is observed in many large sparse networks, including polytrees, Erdős-Rényi, power law, and small world graphs (Durrett, 2007). Power

law graphs are of particular interest in many biological applications, as they allow for the presence of hub nodes.

In this paper, we propose a low-complexity constraint-based method for estimating high-dimensional sparse DAGs. The new method, termed *reduced PC* (rPC), exploits the local separation property of large random networks, which was used by Anandkumar et al. (2012) in estimation of undirected graphical models. We show that rPC can consistently estimate the skeleton of high-dimensional DAGs by conditioning only on sets of small cardinality. This is in contrast to previous heuristic DAG learning approaches that set an upper bound on the number of parents for each node (Acid and de Campos, 1994; Huete and Campos, 1993), which is an assumption that cannot be justified in many real-world networks. We also show that computational and sample complexities of rPC only depend on average sparsity of the graph—a notion that is made more precise in Sections 3 and 4. This leads to considerable advantages over the PC-Algorithm, whose computational and sample complexities scale with the maximal node degree. Moreover, these properties hold for linear structural equation models (Shojaie and Michailidis, 2010) with arbitrary noise distributions, and require a weaker faithfulness conditions on the underlying probability distributions than the PC-Algorithm. We present two version of the rPC algorithm: a “full” version for which we provide theoretical guarantees, and an approximate version which is much faster and performs almost identically in practice.

The rest of the paper is organized as follows. In Section 2 we review basic properties of graphical models over DAGs, and give a short overview of the PC-Algorithm. Our new algorithm is presented in Section 3 and its properties, including consistency in high dimensions are discussed in Section 4. Results of simulation studies and a real data example concerning the estimation of gene regulatory networks are presented in Sections 5 and 6, respectively. We end with a brief discussion in Section 7. Technical proofs and additional simulations are presented in the appendices.

2. Preliminaries

In this section, we review relevant properties of graphical models defined over DAGs, and briefly describe the theory and implementation of the PC-Algorithm.

2.1. Background

For p random variables X_1, \dots, X_p , we define a graph $G = (V, E)$ with vertices, or nodes, $V = \{1, \dots, p\}$ such that variable X_j corresponds to node j . The edge set $E \subset V \times V$ contains directed edges; that is, $(j, k) \in E$ implies $(k, j) \notin E$. Furthermore, there are no directed cycles in G . We denote an edge from j to k as $j \rightarrow k$ and call j a parent of k , and k a child of j . The set of parents of node k is denoted $pa(k)$, while the set of nodes adjacent to it, or all of k ’s parents and children, is denoted $adj(k)$. These notations are also used for the corresponding random variable X_k . We assume there are no hidden common parents of node pairs (that is, no unmeasured confounders). The degree of node k is defined as the number of nodes which are adjacent to it, $|adj(k)|$; we denote the maximal degree in the graph as d_{max} . A triplet of nodes (i, j, k) is called an *unshielded triple* if i and j are adjacent to k but i and j are not adjacent. An unshielded triple is called a *v-structure* if $i \rightarrow j \leftarrow k$.

We assume random variables follow a linear structural equation model (SEM),

$$X_k = \sum_{j \in \text{pa}(k)} \rho_{jk} X_j + \epsilon_k, \quad (1)$$

where for $k = 1, \dots, p$, ϵ_k are independent random variables with finite variance, and ρ_{jk} are fixed unknown constants. The directed Markov property, stated below, is usually assumed in order to connect the joint probability distribution of X_1, \dots, X_p to the structure of the graph G .

Definition 1 *A probability distribution is Markov on a DAG $G = (V, E)$ if, conditional on its parents, every random variable X_k is independent of its non-descendants; that is, $X_k \perp\!\!\!\perp X_j \mid \text{pa}(X_k)$ for all $j \in V$ which are non-descendants of k .*

Although this assumption allows us to connect conditional independence relationships to the DAG structure, there are generally multiple graphs that generate the same distribution under the Markov property. More concretely, DAGs are Markov equivalent if they have the same skeleton and the same set of v-structures. Therefore, constraint-based methods focus primarily on estimating the skeleton of the DAGs from observational data. Conditional independence relations identified when learning the skeleton are then used to orient some of the edges to obtain the CPDAG, which represents the Markov equivalence class of directed graphs (Kalisch and Bühlmann, 2007).

We next define d-separation, a graphical property which is used to read conditional independence relationships from the DAG structure.

Definition 2 *In a DAG G , two nodes k_1 and k_2 are d-separated by a set S if and only if, for all paths π between k_1 and k_2 :*

- (i) π contains a chain $i \rightarrow m \rightarrow j$ or a fork $i \leftarrow m \rightarrow j$ such that the middle node m is in S , or
- (ii) π contains an inverted fork (or collider) $i \rightarrow m \leftarrow j$ such that the middle node m is not in S and no descendant of m is in S .

A path π which does not satisfy the requirements of this definition is known as a *d-connecting path*. Using observed data, d-separations in a graph G can be identified based on conditional independence relationships. To this end, we require the following assumption, known as *faithfulness*, on the probability distribution of random variable on G .

Definition 3 *A probability distribution is faithful to a DAG G if $X_i \perp\!\!\!\perp X_j \mid X_S$ whenever i and j are d-separated by S .*

2.2. The PC-Algorithm

Together, d-separation and faithfulness suggest a simple algorithm for recovering the DAG skeleton. If we discover that $X_i \perp\!\!\!\perp X_j \mid S$ for some set S , then there cannot be an edge $(i, j) \in E$. Conversely, if we discover $X_i \not\perp\!\!\!\perp X_j \mid S$ for all possible sets S , then there must be an edge $(i, j) \in E$. Therefore, under faithfulness, an obvious strategy for skeleton estimation would be to test all possible conditional independence relations for each pair of variables; that is, test whether $X_i \perp\!\!\!\perp X_j \mid S$ for any $S \subset V \setminus \{i, j\}$. While this strategy is

computationally infeasible for large p , and statistically problematic when $p > n$, it forms the basis of the PC-Algorithm. The PC-Algorithm starts with a complete undirected graph and deletes edges (i, j) if a set S can be found such that $X_i \perp\!\!\!\perp X_j \mid S$. The algorithm also uses the fact that if such an S exists, then there exists a set S' such that all nodes in S' are adjacent to i or j and $X_i \perp\!\!\!\perp X_j \mid S'$. Thus, at each step of the algorithm, only local neighbourhoods need to be examined in order to find the separating sets.

Although consistent for sufficiently sparse high-dimensional DAGs, the PC-Algorithm's computational and sample complexity scale with the maximal degree of the graph, d_{max} . Specifically, the algorithm's computational complexity is $O(p^{d_{max}})$ and its sample complexity is $\Omega\{\max(\log p, d_{max}^{1/b})\}$ for some $b \in (0, 1]$. This is problematic for graphs with highly connected hub nodes, which are common in real-world networks (Chen and Sharp, 2004; Jeong et al., 2000). In such graphs, d_{max} typically grows with the number of nodes p , leading to poor accuracy and runtime for the PC-Algorithm.

Another limitation of the PC-Algorithm is that it requires partial correlations between adjacent nodes to be bounded away from 0; this requirement, which needs to hold for all conditioning sets S such that $|S| \leq d_{max}$, is known as *restricted strong faithfulness* (Uhler et al., 2013), and is defined next.

Definition 4 *Given $\lambda \in (0, 1)$, a distribution P is said to be restricted λ -strong-faithful to a DAG $G = (V, E)$ if the following conditions are satisfied:*

- (i) $\min\{|\rho(X_i, X_j \mid X_S)| : (i, j) \in E, S \subset V \setminus \{i, j\}, |S| \leq d_{max}\} > \lambda$, and
- (ii) $\min\{|\rho(X_i, X_j \mid X_S)| : (i, j, S) \in N_G\} > \lambda$, where N_G is the set of triples (i, j, S) such that i, j are not adjacent, but there exists $k \in V$ making (i, j, k) an unshielded triple, and i, j are not d -separated given S .

Kalisch and Bühlmann (2007) assume that $\lambda = \Omega(n^{-w})$ for $w \in (0, b/2)$ where $b \in (0, 1]$ relates to the scaling of d_{max} . In the low-dimensional setting, it has been shown that the PC-Algorithm achieves uniform consistency with λ converging to zero at rate $n^{1/2}$, which then gives the same condition as ordinary faithfulness (Zhang and Spirtes, 2002). However, in the high-dimensional setting, the set of distributions which are not restricted strong faithful has nonzero measure. In fact, Uhler et al. (2013) showed that this assumption is overly restrictive and that the measure of unfaithful distributions converges to 1 exponentially in p . We will revisit the faithfulness assumption in Section 4.2.

To address the limitations of the PC-Algorithm, we next propose a new algorithm that takes advantage of the structure of large networks from common random graph families. By doing so, we obtain improved computational and sample complexity; as we will show, these complexities are unaffected by the increase in the maximal degree as the graph becomes larger. We also prove consistency under a weaker faithfulness assumption than that needed for the PC-Algorithm.

3. The Reduced PC-Algorithm

As with the PC-Algorithm, our strategy for estimating the graph skeleton is to start with a complete graph, and then delete edges by discovering separating sets. Under a faithfulness assumption on a linear structural equation model, we do so by computing partial correlations

and declaring X_i and X_j d-separated by S if $\rho(X_i, X_j \mid S)$ is smaller than some threshold α . Aside from thresholding, the key difference between our proposal and the PC-Algorithm is that we only consider partial correlations conditional on sets S with small cardinality.

We justify our method using two key observations. Our first key observation is based on the decomposition of covariances over *treks*, which are special types of paths in directed graphs.

Definition 5 *A trek between two nodes i and j in a DAG G is either a path from i to j , a path from j to i , or a pair of paths from a third node k to i and j such that the two paths only have k in common.*

In a linear SEM (1), the covariance between two random variables is characterized by the treks between them. Denoting a trek between nodes i and j as $\pi : i \leftrightarrow j$ with common node or source k , the covariance is given by

$$\text{cov}(X_i, X_j) = \sum_{\pi: i \leftrightarrow j} \sigma_k \prod_{e \in \pi} \rho_e, \quad (2)$$

where σ_k is the variance of ϵ_k from Equation (1) and ρ_e denotes the weight of an edge along the trek, which is the corresponding coefficient in the SEM. This is shown in full detail by Sullivant et al. (2010), who also show that the covariance conditional on a d-separating set S leaves out treks which include any nodes in S . This conditioning effect is illustrated in Figure 1, where conditioning on an appropriate set S blocks the treks between non-adjacent nodes, without resulting in any additional d-connecting paths. If all edge weights are bounded by 1 in absolute value, then the contribution of each trek to the covariance decays exponentially in trek length. In practice, we scale the data matrix X so that each column has unit standard deviation. Under some conditions, most of the edge weights in the linear SEM (1) then satisfy $|\rho_{ij}| < 1$ (this is discussed and shown empirically in Appendix B). Hence, the contribution of long treks to the conditional covariance among non-adjacent nodes is negligible. This decay motivates the thresholding of partial correlations in rPC, and is further discussed in Section 4.

The above observation suggests a new strategy for learning DAG structures by only considering short treks. Suppose S is the set that blocks all short treks between two nodes j and k . If the correlation over all remaining d-connecting paths (after conditioning on S) between j and k is negligible, then the partial correlation given S , $\rho(X_j, X_k \mid S)$ can be used to determine whether j and k are adjacent.

To determine the size of the conditioning set, S , we need to determine the number of short treks between any two nodes j and k . Our second key observation addresses this question, by utilizing properties of large random graphs. More specifically, motivated by Anandkumar et al.’s proposal for estimating undirected graphs, we consider a key feature of large random networks, known as the *local separation property*.

Definition 6 *Given a graph G , a γ -local separator $S_\gamma(i, j) \subset V$ between non-neighbours i and j minimally separates i and j with respect to paths of length at most γ .*

Definition 7 *A family of graphs \mathcal{G} satisfies the (η, γ) -local separation property if, as $p \rightarrow \infty$, $Pr(\exists G \in \mathcal{G} : \exists(i, j) \notin E_G, |S_\gamma(i, j)| > \eta) \rightarrow 0$.*

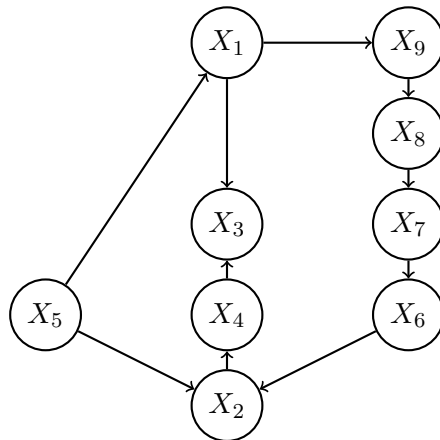


Figure 1: Illustration of treks between nodes X_1 and X_2 within a DAG. The middle path involves a collider, X_3 , so does not contribute to $\text{cov}(X_1, X_2)$, which is $\rho_{51}\rho_{52} + \rho_{19}\rho_{98}\rho_{87}\rho_{76}\rho_{62}$. The conditional covariance $\text{cov}(X_1, X_2 \mid X_5)$ excludes treks that involve X_5 , and is thus $\rho_{19}\rho_{98}\rho_{87}\rho_{76}\rho_{62}$. Here, $S = \{X_5, X_7\}$ is a d-separating set, as it blocks both treks, giving $\text{cov}(X_1, X_2 \mid S) = 0$.

Intuitively, under (η, γ) -local separation, with high probability, the number of short treks—of length at most γ —between any two non-neighbouring nodes is bounded above by η . In fact, as the local separation property refers to any type of path, there are likely even fewer than η short treks between any two neighbouring nodes. Therefore, we only need to consider conditioning on sets S of size at most η in order to remove the correlation induced by short treks. Combining this with our first insight, we ignore treks (and other possible d-connecting paths) of length longer than γ , which, for appropriate probability distributions, have a negligible impact on partial correlations. The resulting procedure, called the full reduced PC-Algorithm (rPC-full), is presented in Algorithm 1.

We note a key implementation difference between rPC-full and the ordinary PC-Algorithm: when searching for a separating set, rPC-full considers all $S \subset V \setminus \{i, j\}$, while PC-Algorithm only searches over the local neighbourhoods of i and j . Recall that if a set S d-separates nodes i and j , then there exists a d-separating set S' such that all nodes in S' are adjacent to i or j . However, $|S'|$ may be larger than $|S|$; because rPC only considers sets of size up to η , a full search is needed to ensure discovery of a d-separating set. Motivated by the PC-Algorithm, we also propose an approximate reduced PC-Algorithm (rPC-approx), which uses the same local neighbourhood search, i.e. $S \subset \text{adj}(i) \cup \text{adj}(j) \setminus \{i, j\}$. In practice, we show that rPC-approx performs almost identically to rPC-full.

To recap, our proposal in Algorithm 1 hinges on two important properties of probability models on large DAGs: (P1) boundedness of the number of short treks between any two nodes, and (P2) negligibility of correlation over the remaining (long) d-connecting paths. The first property, which is characterized by local separation, concerns solely the DAG structure. Anandkumar et al. (2012) show that many common graph families satisfy the local separation property with small η . Specifically, sparse, large binary trees, Erdős-Rényi

graphs, and graphs with power law degree distributions all satisfy this property with $\eta \leq 2$. Moreover, the sparsity requirement for these graph families is in terms of average node degree, and not the maximum node degree. For these graph families, our algorithm only needs to consider separating sets of size 0, 1, and 2, irrespective of the maximum node degree. Small-world graphs, as generated by the Watts-Strogatz algorithm (Kleinberg, 2000), also satisfy this property, but with $\eta > 2$. In addition, the γ parameter increases with p for these families; thus, as graphs get larger, the local separation property applies to a larger set of paths.

By only considering a bounded number of short paths, our algorithm has computational complexity $O(p^{\eta+2})$, and thus avoids the exponential scaling in d_{max} that the PC-Algorithm suffers from. This is particularly significant in the case of power-law graphs, where $d_{max} = O(p^a)$ for $a > 0$ (Molloy and Reed, 1995); in this case, PC-Algorithm has a computational complexity of $O(p^{p^a})$, which is significantly worse than rPC's complexity of $O(p^4)$. While rPC-full might not be faster in practice due to the larger search space, rPC-approx would result in significant speedup.

Unlike the first property (P1), the second property needed for our algorithm, namely the negligibility of correlation over d -connecting paths, concerns both the structure of the DAG G , and the probability distribution P of variables on the graph. In the next section, we discuss two alternative sufficient conditions that guarantee this property, and allow us to consistently estimate the DAG skeleton.

INPUT: Observations from random variables X_1, X_2, \dots, X_p ; threshold level α ;
 maximum separating set size η .
 OUTPUT: Estimated skeleton C .

```

Set  $V = \{1, \dots, p\}$ .
Form the complete undirected graph  $\tilde{C}$  on the vertex set  $V$ .
Set  $l = -1$ ;  $C = \tilde{C}$ .
repeat
     $l = l + 1$ 
    repeat
        Select a (new) ordered pair of nodes  $i, j$  that are adjacent in  $C$ 
        repeat
            Choose (new)  $S \subset V \setminus \{i, j\}$  with  $|S| = l$ 
            if  $\rho(X_i, X_j \mid S) \leq \alpha$ 
                Delete edge  $(i, j)$ 
                Denote this new graph by  $C$ 
            end if
        until edge  $(i, j)$  deleted or all  $S \subset V \setminus \{i, j\}$  with  $|S| = l$  have been chosen
    until all ordered pairs of adjacent nodes  $i$  and  $j$  have been examined for  $\rho(X_i, X_j \mid S) \leq \alpha$ 
until  $l > \eta$ 
    
```

Algorithm 1: The full reduced PC-Algorithm (rPC-full)

4. Algorithm Analysis and Asymptotics

In this section, we describe in detail the assumptions required for the rPC-full algorithm to consistently recover the DAG skeleton. We also discuss its computational and statistical properties, particularly in comparison with the PC-Algorithm.

4.1. Consistency

As discussed in the previous section, to consistently recover the DAG skeleton, rPC requires that properties (P1) and (P2) hold; namely, that the graph under consideration has a bounded number of short paths and the correlation over d -connecting paths of length greater than γ decays sufficiently quickly. In fact, the trek decomposition (2) indicates that the correlation for each long trek is small when most of the edge weights are bounded by 1 in absolute value. However, condition (P2) requires the total correlation over all long treks (and other d -connecting paths) to decay sufficiently quickly. To this end, we consider two alternative sufficient conditions. The first condition is a direct assumption on the boundedness of the conditional correlation. The second is inspired by Anandkumar et al. (2012), and assumes the underlying probability model satisfies what we term *directed walk-summability*. This condition mirrors the walk-summability condition for undirected graphical models, which has been well-studied and shown to hold in a large class of models (Malioutov et al., 2006).

Definition 8 *A probability model is directed β -walk-summable on a DAG with weighted adjacency matrix A , if $\|A\| \leq \beta < 1$ where $\|\cdot\|$ denotes the spectral norm.*

As an alternative to this condition, we also present a direct assumption, Assumption 4. This condition is less restrictive than directed walk-summability, but has not been characterized in the literature. However, given that γ increases with p in the graph families we are considering, it is intuitive that the sum of edge weight products over treks longer than length γ will be decreasing and asymptotically small. This also holds for non-trek d -connecting paths. These two assumptions lead to two parallel proofs of the consistency of our algorithm, presented in Theorem 1. Before stating the theorem, we discuss our assumptions.

Similar to the PC-Algorithm, our method requires a faithfulness condition. As stated previously, our condition, which we term *path faithfulness* and is defined next, is weaker than PC-Algorithm’s λ -strong faithfulness stated in Definition 4 (see Section 4.2 for additional details).

Definition 9 *Given $\lambda \in (0, 1)$, a distribution P is said to be λ -path-faithful to a DAG $G = (V, E)$ if both of the following conditions hold:*

- (i) $\min \{|\rho(X_i, X_j | X_S)| : (i, j) \in E, S \subset V \setminus \{i, j\}, |S| \leq \eta\} > \lambda$, for some η , and
- (ii) $\min \{|\rho(X_i, X_j | X_S)| : (i, j, S) \in N_G\} > \lambda$, where N_G is the set of triples (i, j, S) such that i, j are not adjacent, but there exists $k \in V$ making (i, j, k) an unshielded triple, and i, j are not d -separated given S .

Part (i) of the assumption only requires partial correlations between true edges conditioned on sets of size up to η to be bounded away from zero, while the PC-Algorithm requires this

for conditioning sets of size up to d_{max} . In Section 4.2, we discuss how this affects bounds on the true partial correlations, and also empirically show that the above path faithfulness assumption is less restrictive than corresponding assumption for the PC-Algorithm.

Assumption 1 (Path faithfulness and Markov property) *The probability distribution P of random variables corresponds to a linear SEM (1) with sub-Gaussian errors, and is λ -path-faithful to the DAG G , with $\lambda = \Omega(n^{-c})$ for $c \in (0, 1/2)$.*

Our second assumption ensures that the covariance matrix of the structural equation model and its inverse remain bounded as p grows.

Assumption 2 (Covariance and precision matrix boundedness) *The covariance matrix of the model Σ_G and its inverse Σ_G^{-1} are bounded in spectral norm, that is, $\max(\|\Sigma_G\|, \|\Sigma_G^{-1}\|) \leq M < \infty$ for all p .*

The last three assumptions characterize applicable graph families and probability distributions.

Assumption 3 ((η, γ) -local separation) *The DAG G belongs to a family of random graphs \mathcal{G} that satisfies the (η, γ) -local separation property with $\eta = O(1)$ and $\gamma = O(\log p)$.*

Assumption 4 (Bounded long path weight) *Let π denote a d -connecting path of length $l(\pi)$ between two non-adjacent vertices i and j in G . Then, there exists a conditioning set S such that the total edge weight over d -connecting paths longer than γ satisfies:*

$$\sum_{l=\gamma+1}^{p-1} \sum_{l(\pi)=l} |\rho_{\pi,1} \cdots \rho_{\pi,l}| = O(\beta^\gamma),$$

for some $\beta \in (0, 1)$.

Assumption 4 guarantees that the sum of weights over long treks between any two nodes i and j is bounded. For a single trek, a sufficient condition is that all edge weights are bounded by 1 in magnitude. In Appendix B, we provide further discussion on Assumption 4, and empirically investigate its plausibility. In particular, we show that if the data matrix X is scaled so that each column has unit standard deviation, then with high probability all edge weights are bounded by 1 in absolute value. To account for residual correlation induced through conditioning on common descendants of i and j , Assumption 4 also includes non-trek d -connecting paths.

Assumption 4' (Directed β -walk-summability) *The probability distribution P is directed β -walk-summable.*

We are now ready to state our main result. The result is proved in Appendix A, where the error probabilities are also analyzed.

Theorem 1 *Under Assumptions 1-3 and either Assumption 4 or 4', there exists a parameter α for thresholding partial correlations such that, as $n, p \rightarrow \infty$ with $n = \Omega\{(\log p)^{1/(1-2c)}\}$, the full reduced PC (rPC-full) procedure, as described in Algorithm 1, consistently learns the skeleton of the DAG G .*

Several theoretical features of our algorithm are attractive. As stated previously, our faithfulness condition is weaker than the corresponding assumption for the PC-Algorithm and related methods. Similar to its computational complexity, the sample complexity of our algorithm also does not scale with the maximal node degree, and is only dependent on the parameter η as p increases. For example, in a power law graph, the sample complexity of the PC-Algorithm is $\Omega\{\max(\log p, p^{ab})\}$ for $0 < a, b < 1$, compared to $\Omega\{(\log p)^{1/1-2c}\}$ with $c \in (0, 1/2)$ for our method. This gain in efficiency is due to fact that the maximum separating set size, η , remains constant in rPC. Finally, our algorithm does not require the data to be jointly Gaussian. The proof of the algorithm’s consistency only requires that the population covariance matrix can be well-approximated from the data; for simplicity, we assume a sub-Gaussian distribution.

4.2. On Faithfulness Assumption

As stated in Section 2.2, for large biological networks of interest, the maximum node degree, d_{max} , often grows with p . Therefore, the λ -restricted strong faithfulness condition of the PC-Algorithm—Definition 4—becomes exponentially harder to satisfy with increasing network size. A full discussion of this phenomenon can be found in Uhler et al. (2013), where it is shown that the measure of strong unfaithful distributions converges to 1 for various graph structures. Although this would also occur with path faithfulness (Definition 9), our condition allows a rate for λ that is independent of d_{max} and p .

The rate for λ in the PC-Algorithm is $\lambda = \Omega(n^{-w})$ for $w \in (0, b/2)$, where $d_{max} = O(n^{1-b})$ for $b \in (0, 1]$. For $b = 1$, or constant d_{max} , the PC-Algorithm’s required scaling for λ is identical to that for our method in Assumption 1. This makes sense intuitively, since our method is not affected by the increase in d_{max} . For other values of b , the scaling of λ becomes more restricted for the PC-Algorithm; for example, if $b = 1/2$, then $\lambda = \Omega(n^{-w})$ for $w \in (0, 1/4)$. However, under path faithfulness (Definition 9), we can still achieve a rate of $\lambda = \Omega(n^{-1/2})$; that is, the partial correlations are allowed to be smaller and the condition is weaker.

We report the findings of a simulation study, similar to that in Uhler et al. (2013), which examines how often randomly generated DAGs satisfy part (i) of the path faithfulness assumption compared to restricted strong faithfulness. We are primarily interested in part (i), as this part is needed for consistent skeleton estimation; part (ii), on the other hand, is needed to obtain correct separating sets in order to obtain partial orientation of edges. In this simulation, 1000 random DAGs were generated from Erdős-Rényi and power law families, with edge weights drawn independently from a Uniform($-1, 1$) distribution. Each DAG had $p = 20$ nodes, with varying expected degrees per node. For each simulation setting, we computed the proportion of DAGs that satisfied part (i) of the λ -restricted-strong-faithfulness and λ -path-faithfulness conditions with $\lambda = 0.001$ and $\eta = 2$. The results are shown in Table 1. We see that path faithfulness is much more likely to be satisfied than restricted strong faithfulness, especially for power law graphs. This is to be expected, as the number of constraints required for restricted strong faithfulness grows with d_{max} , but remains constant for path faithfulness. It is, however, difficult for dense graphs to satisfy either condition, although there is a mild advantage for path faithfulness. In Appendix B, we provide the results of further simulation studies with $p = 10$ and $p = 30$;

Graph family	Expected degree	$Pr(\text{RSF})$	$Pr(\text{PF})$
Erdős-Rényi	2	0.77	0.92
Erdős-Rényi	5	0	0.08
Power law	2	0.54	0.85
Power law	6	0.003	0.08

Table 1: Empirical probabilities of random DAGs of size $p = 20$ satisfying faithfulness conditions; RSF refers to restricted strong faithfulness of the PC-Algorithm, and PF refers to path faithfulness of reduced PC (rPC).

both of these give similar conclusions and indicate that path faithfulness remains easier to satisfy with increasing network size.

4.3. Tuning Parameter Selection

Our algorithm requires two tuning parameters: the maximum separating set size η , and the threshold level for partial correlations α . The parameter η varies based on the underlying graph family. Thus, given knowledge of a plausible graph structure, η can be pre-specified. Alternatively, η can be selected by maximizing a goodness-of-fit score over a parameter grid, along with α . This may be preferable as the local separation results consider all short paths, not just treks, so better performance may be obtained by specifying a smaller η . Likewise, when using the rPC-approx algorithm, a larger η could be needed to discover appropriate separating sets.

For jointly Gaussian data, we can obtain a modified version of the Bayesian information criterion by fitting the likelihood to the CPDAG obtained based on the estimated DAG skeleton (Foygel and Drton, 2010). Following Anandkumar et al. (2012), and denoting by X_{obs} the observed data, we use:

$$\text{BIC}(X_{obs}; \hat{G}) = \log f(X_{obs}; \hat{\theta}) - 0.5|E| \log(n) - 2|E| \log(p), \quad (3)$$

where \hat{G} denotes one of the DAGs obtained from the estimated CPDAG containing $|E|$ edges. The CPDAG represents the Markov equivalence class of DAGs, so all possible graphs will result in the same fitted Gaussian model with parameters $\hat{\theta}$. We use this BIC for tuning parameter selection; higher scores imply a better fit. For linear SEMs with non-Gaussian noise distributions, the Gaussian likelihood serves as a surrogate goodness-of-fit measure, and BIC can still be used to select the tuning parameters.

5. Simulation Studies

In this section, we compare the performance of rPC-approx, rPC-full, and the standard PC-Algorithm in multiple simulation settings. We consider both setting a constant value for η in rPC, and tuning it to maximize the BIC.

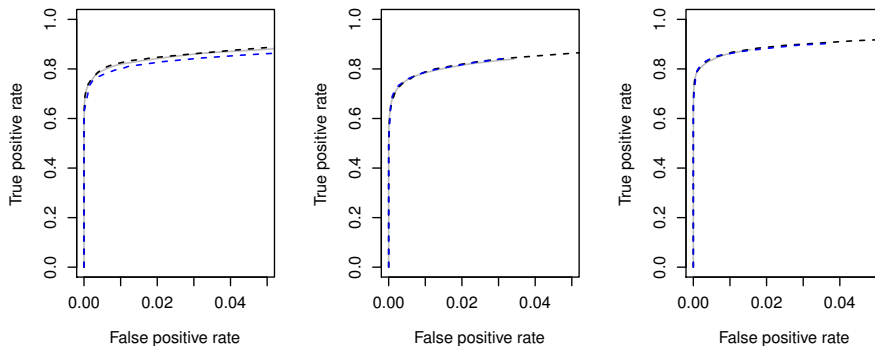


Figure 2: Average true vs. false positive rates for PC-Algorithm (grey solid line), rPC-approx (black dashed line), and rPC-full (blue dashed line) estimating Erdős-Rényi DAGs. Left: $p = 100, n = 200$; centre: $p = 200, n = 100$; right: $p = 500, n = 200$.

5.1. Pre-specified η Parameter

To facilitate the comparison with the PC-Algorithm, we generate data from Gaussian linear SEMs as in Equation 1, with the dependency structure specified by a DAG from Erdős-Rényi and power law families. We implement our algorithm with maximum separating set size $\eta = 2$ since these families are known to satisfy (η, γ) -local-separation with $\eta \leq 2$ (Anandkumar et al., 2012).

We generate a random graph with p nodes using the `igraph` library in R, assigning every edge a weight from a Uniform(0.1, 1) distribution. We then use the `rmvDAG` function from the `pcalg` library to simulate n observations from the DAG. This is repeated 20 times for each thresholding level α ; average true and false positive rates for both algorithms are reported over the grid of α values. Our grid of α values produces partial receiver operating characteristic (pROC) curves for varying sample sizes and graph structures, which are used to assess the estimation accuracy of the methods.

For both Erdős-Rényi and power law DAGs, we consider a low-dimensional setting with $p = 100$ nodes and $n = 200$ observations, and two high-dimensional settings with $(p, n) = (200, 100)$ and $(p, n) = (500, 200)$. In all settings, the DAGs are set to have an average degree of 2. The maximum degrees of the Erdős-Rényi graphs range from 5 to 7. The maximum degrees for power law graphs increase with p and are 42, 69, and 71 for the three simulation settings.

Results for Erdős-Rényi DAGs are shown in Figure 2. In this setting, all algorithms perform almost identically in both low and high-dimensional cases. This is because for the relatively small maximum degree in Erdős-Rényi graphs, the sizes of conditioning sets for PC and rPC are not very different.

Results for power law DAGs are shown in Figure 3. In this case, both versions of rPC offer improved accuracy compared to the PC-Algorithm, in both low and high-dimensional

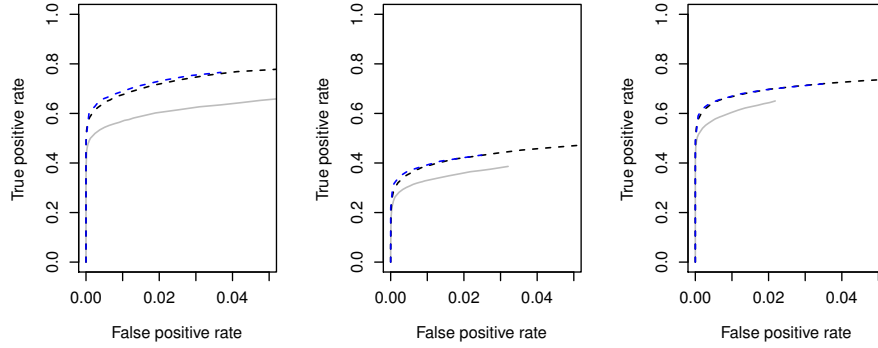


Figure 3: Average true vs. false positive rates for PC-Algorithm (grey solid line), rPC-approx (black dashed line), and rPC-full (blue dashed line) estimating power law DAGs. Left: $p = 100, n = 200$; centre: $p = 200, n = 100$; right: $p = 500, n = 200$.

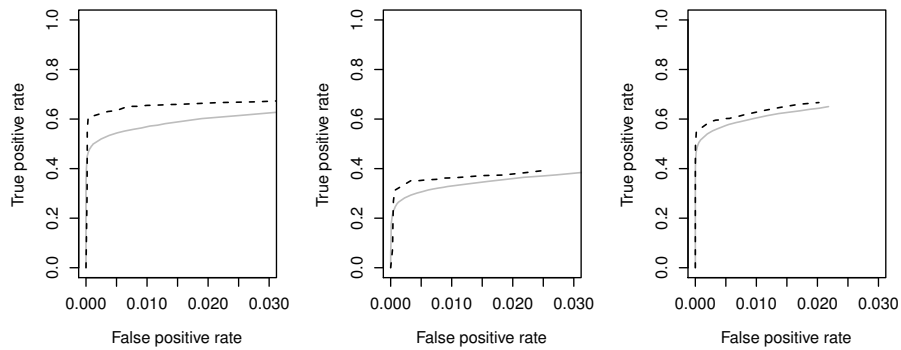


Figure 4: Average true vs. false positive rates for PC-Algorithm (grey solid line) and rPC-approx with BIC-tuned η (black dashed line) estimating power law DAGs. Left: $p = 100, n = 200$; centre: $p = 200, n = 100$; right: $p = 500, n = 200$.

settings. These results confirm our theoretical findings, and show that our algorithm performs better at estimating DAGs with hub nodes than the PC-Algorithm.

Additional simulations in Appendix B show similar results in more dense DAGs. As the underlying DAG becomes more dense, all methods perform worse; however, our algorithms maintain an advantage over the PC-Algorithm in the power law setting. All of these simulations also confirm that rPC-full and rPC-approx perform very similarly. This suggests that the situation where conditioning on non-local sets is required occurs with very low probability, indicating that rPC-approx provides suitably good estimates.

We also compare the runtimes of the algorithms for these settings. Over 100 iterations, we generate a random data set, and apply all the algorithms with a range of tuning parameters. Specifically, we set $\eta = 2$ for rPC, and $\alpha = \{10^{-6}, 10^{-5}, 10^{-4}, 10^{-3}, 10^{-2}\}$. We then take the total runtime over all parameters, and compare by considering the mean value of $100 \left(\frac{\text{time}_{rPC}}{\text{time}_{PC}} \right)$ for both rPC-approx and rPC-full. The results are shown in Table 2. We observe that rPC-approx is significantly faster than the PC-Algorithm for power law graphs. As expected, our implementation of rPC-full is slower than the PC-Algorithm in all settings, given its exhaustive search over all possible separating sets. Given the results indicating that rPC-approx performs almost identically to rPC-full, we suggest that practitioners use rPC-approx, possibly combined with BIC tuning if a suitable η bound is not known. We present relevant simulation results in the next section.

Graph family	p	n	rPC-approx	rPC-full
Erdős-Rényi	100	200	0.98	3.08
Erdős-Rényi	200	100	1.00	6.16
Erdős-Rényi	500	200	1.00	23.25
Power law	100	200	0.60	1.13
Power law	200	100	0.92	3.81
Power law	500	200	0.36	5.23

Table 2: Empirical ratio of rPC-approx and rPC-full runtimes to PC-Algorithm runtime.

5.2. BIC-tuned η Parameter

In this section, we consider simulations where the maximum size of the conditioning set for rPC-approx, η , is selected to maximize the BIC score defined in (3). To this end, we consider power law DAGs with the same low and high-dimensional settings as before. We select the value of $\eta \in \{1, 2, 3, 4\}$ which maximizes the BIC at each α value. While rPC-full could also be tuned in this way, it would be computationally expensive to consider η beyond 3 for $p > 200$. The results in Figure 4 show that our algorithm maintains an advantage over the PC-Algorithm. Interestingly, for values of α which yielded the best estimation accuracy, the optimal η selected was most frequently 1, which confirms our intuition from Section 3 that the η parameter for a graph family should be seen as an upper bound for estimating DAGs using our algorithm.

6. Application: Estimation of Gene Regulatory Networks

We apply our algorithms and the PC-Algorithm to a gene expression data set of $n = 487$ patients with prostate cancer from The Cancer Genome Atlas (Cancer Genome Atlas Research Network, 2012). We select $p = 272$ genes with known network structure from BioGRID (Stark et al., 2006), and attempt to recover this network from the data. We choose the tuning parameters for both rPC algorithms and the p-value threshold for the PC-Algorithm by searching over a grid of values and selecting those which yielded the largest BIC (3). We found that the best rPC algorithm in terms of BIC was rPC-approx with $\eta = 3$ and $\alpha = 0.09$, while the best p-value threshold for the PC-Algorithm was 0.07.

The BioGRID database provides valuable information about known gene regulatory interactions. However, this database mainly captures genetic interactions in normal cells. Thus, the information from BioGRID may not correctly capture interactions in cancerous cells, which are of interest in our application (Ideker and Krogan, 2012). Despite this limitation, highly connected hub genes in the BioGRID network, which usually correspond to transcription factors, are expected to stay highly connected in cancer cells. Therefore, to evaluate the performance of the two methods, we focus here on the identification of hub genes, which are often most clinically relevant (Horvath et al., 2006; Jeong et al., 2001; Carlson et al., 2006).

The two estimated networks and their hub genes are visualized in Figure 5. Here, we define hub genes as nodes with degree at least 8, which corresponds to the 75th percentile in the degree distribution of both estimates. rPC-approx identifies 19 of 57 true hubs, while the PC-Algorithm only identifies 6. Interestingly, several of the hub genes uniquely identified by rPC are known to be associated with prostate cancer, including ACP1, ARHGEF12, CDH1, EGFR, and PLXNB1 (Ruela-de Sousa et al., 2016; Robbins et al., 2011; Ratz et al., 2017; Sun et al., 2017; Liu et al., 2017). These results suggest that rPC may be a promising alternative for estimating biological networks, where highly-connected nodes are of clinical importance.

Examining the two networks also indicates that for nodes with small degrees, the estimated neighborhoods from rPC are very similar to those from the PC-Algorithm. To assess this observation, we consider the induced subgraph of nodes with degree at most 5 in the PC-Algorithm estimate. The F_1 score—which is a weighted average of precision and recall—between the two estimates of this sparse subnetwork is 0.86. This value indicates that the two algorithms perform very similarly over sparse nodes.

7. Discussion

Our new algorithm for learning directed acyclic graphs (DAGs) by conditioning on small sets leads to more efficient computation and estimation under a less restrictive faithfulness assumption than the PC-Algorithm. However, our weaker faithfulness condition may still not be satisfied for dense DAGs or in structural equation models with edge weights distributed over a larger parameter space. This is shown both geometrically and empirically in Uhler et al. (2013), and remains a direction for future research. Generalizing the idea of restricted conditioning to more complex probability models over DAGs, such as nonlinear SEMs (Voorman et al., 2014) would also be of interest. Finally, the idea of conditioning

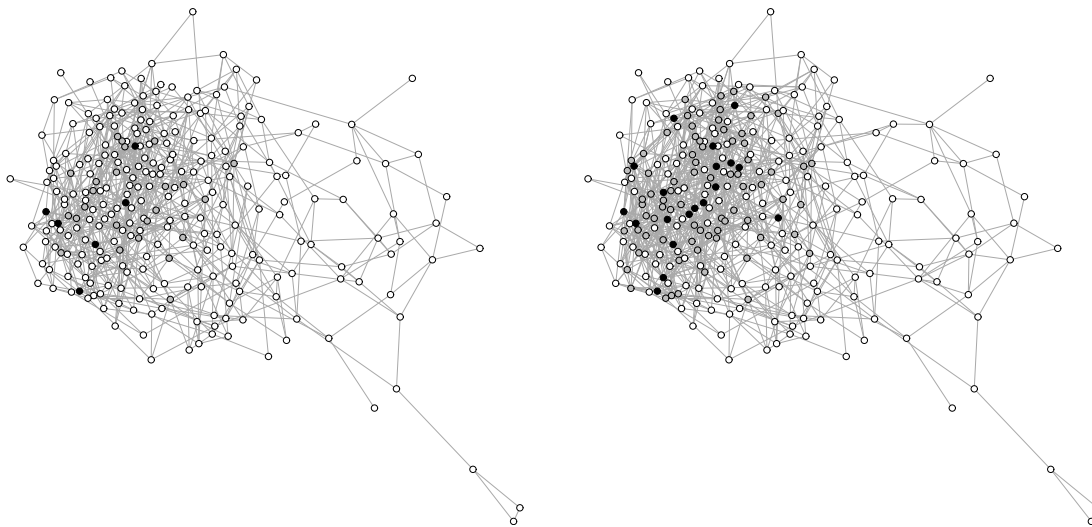


Figure 5: Estimated skeletons of gene regulatory networks in prostate cancer subjects. Black nodes are classified as hubs, having estimated degree of at least 8. Grey nodes are identified hubs that are also considered hubs in the BioGRID data. Left: PC-Algorithm; right: rPC-approx.

on small sets of variables can also be used to develop more efficient hybrid methods for learning DAGs in high dimensions.

Acknowledgments

This work was partially supported by grants from the National Science Foundation (NSF grant DMS/NIGMS-1561814 and DMS-1722246) and the National Institute of Health (NIH K01-HL124050 and R01-GM114029).

Appendix A. Proof of Theorem 1

In this section, we prove the consistency of rPC-full for estimating the skeleton of a directed acyclic graph under the assumptions stated in the main paper.

We begin by establishing that correlations decay over long paths in the graph, and use this to show that the partial correlation between two non-adjacent nodes, conditional on a suitable set S with small cardinality, is bounded above. We establish this through two possible sufficient conditions: Lemma 1 is based on Assumption 4, which directly assumes that the total weight over long paths is sufficiently small; on the other hand, Lemma 2 uses Assumption 4', which assumes the underlying model is directed walk-summable. Combining this result with Assumption 1, which says that the relevant partial correlations between two adjacent nodes is bounded below, we have oracle consistency of rPC-full. In Lemma 3, we invoke a concentration inequality for sample partial correlations to bound their deviations from population quantities. Using these results, we then prove that there exists a threshold level that consistently recovers the true skeleton in the finite sample setting.

In our first two lemmas, we make use of the local separation property in Assumption 3. While Assumption 3 does not directly concern *local separating sets*, since a d -separating set is a subset of a general separating set, Assumption 3 asymptotically guarantees the existence of a γ -local d -separator of size at most η for any two non-neighbouring nodes.

Definition 10 *Given a graph G , a γ -local d -separator $S_\gamma(i, j) \subset V \setminus \{i, j\}$ between non-neighbours i and j minimally d -separates i and j over paths of length at most γ .*

Lemma 11 *Under Assumptions 1-4, the partial correlation between non-neighbours i and j satisfies*

$$\min_{S \in S_{\eta, \gamma}} |\rho(i, j \mid S)| = O(\beta^\gamma).$$

where $S_{\eta, \gamma}$ is the set of γ -local d -separators of size at most η .

Proof

Recall the form of the linear structural equation model:

$$X_k = \sum_{j \in pa(k)} \rho_{jk} X_j + \epsilon_k,$$

where ϵ_k are independent and $Var(\epsilon_k) = \sigma_k^2 < \infty$ for all k .

Let A_G denote the lower-triangular weighted adjacency matrix for the graph G , obtained by ordering the nodes according to a causal order (Shojaie and Michailidis, 2010), so that $j \in pa(k)$ implies $j < k$. Then, as shown by Shojaie and Michailidis (2010) for the Gaussian case and by Loh and Bühlmann (2014) in general linear structural equation models,

$$\Sigma_G = (I - A_G)^{-1} D (I - A_G)^{-T},$$

where $D = \text{diag}(\sigma_1^2, \dots, \sigma_p^2)$.

First, suppose that $\sigma_i^2 = 1$ for all i . We consider the conditional covariance $\Sigma_G(i, j \mid S)$ where i and j are non-neighbours, and S is the γ -local d -separator, as defined above. Let

π denote a d -connecting path between i and j of length $l(\pi)$, and ρ_1, \dots, ρ_l denote the edge weights along the path. By conditioning on S , we have that covariance is only induced through d -connecting paths of length greater than γ , as referenced in the main paper. Then,

$$\Sigma_G(i, j | S) = \sum_{\substack{\pi: i \leftrightarrow j \\ \pi \cap S = \emptyset}} \sum_{\pi: l(\pi)=l} \prod_{k=1}^l \rho_k = \sum_{\substack{\pi: i \leftrightarrow j \\ l(\pi)=\gamma+1}}^{p-1} \sum_{\pi: l(\pi)=l} \prod_{k=1}^l \rho_k.$$

Therefore:

$$\begin{aligned} |\Sigma_G(i, j | S)| &= \left| \sum_{\substack{\pi: i \leftrightarrow j \\ l(\pi)=\gamma+1}}^{p-1} \sum_{\pi: l(\pi)=l} \prod_{k=1}^l \rho_k \right| \\ &\leq \sum_{\substack{\pi: i \leftrightarrow j \\ l(\pi)=\gamma+1}}^{p-1} \sum_{\pi: l(\pi)=l} \prod_{k=1}^l |\rho_k| && \text{(by triangle inequality)} \\ &= O(\beta^\gamma). && \text{(by Assumption 4)} \end{aligned}$$

Now, suppose that not all $\sigma_i^2 = 1$. Then, let $\sigma_{max}^2 = \max_i \sigma_i^2$. We have:

$$\begin{aligned} |\Sigma_G(i, j | S)| &\leq \sum_{\substack{\pi: i \leftrightarrow j \\ l(\pi)=\gamma+1}}^{p-1} \sum_{\pi: l(\pi)=l} \sigma_{max}^2 \prod_{k=1}^l |\rho_k| \\ &= \sigma_{max}^2 \sum_{\substack{\pi: i \leftrightarrow j \\ l(\pi)=\gamma+1}}^{p-1} \sum_{\pi: l(\pi)=l} \prod_{k=1}^l |\rho_k| \\ &= \sigma_{max}^2 O(\beta^\gamma) \\ &= O(\beta^\gamma). && \text{(by Assumption 2)} \end{aligned}$$

Finally, we have $|\rho(i, j | S)| = \frac{|\Sigma_G(i, j | S)|}{\sqrt{\Sigma_G(i, i | S)\Sigma_G(j, j | S)}} = O(\beta^\gamma)$ by Assumption 2, since the conditional variances are functions of the marginal variances, which are bounded. \blacksquare

Next, we show the same result by assuming directed walk-summability of the model.

Lemma 12 *Under Assumptions 1-3 and Assumption 4', the partial correlation between non-neighbours i and j satisfies*

$$\min_{S \in S_{\eta, \gamma}} |\rho(i, j | S)| = O(\beta^\gamma).$$

where $S_{\eta, \gamma}$ is the set of γ -local d -separators of size at most η .

Proof

Recall from the proof of Lemma 1,

$$\Sigma_G = (I - A_G)^{-1} D (I - A_G)^{-T},$$

where $D = \text{diag}(\sigma_1^2, \dots, \sigma_p^2)$. First, suppose that $\sigma_i^2 = 1$ for all i . Then, we can write:

$$\begin{aligned} \Sigma_G &= \left(\sum_{r=0}^{\infty} A_G^r \right) \left(\sum_{r=0}^{\infty} A_G^r \right)^T \\ &= \left(\sum_{r=0}^{\gamma} A_G^r + \sum_{r=\gamma+1}^{\infty} A_G^r \right) \left(\sum_{r=0}^{\gamma} A_G^r + \sum_{r=\gamma+1}^{\infty} A_G^r \right)^T. \end{aligned}$$

Now, let Σ_H denote the covariance matrix induced by only considering d-connecting paths of length at most γ . For convenience, let $\Lambda_H := \sum_{r=0}^{\gamma} A_G^r$ and $R_\gamma := \sum_{r=\gamma+1}^{\infty} A_G^r$. Considering their spectral norms, denoted by $\|\cdot\|$, we have by walk-summability that $\|\Lambda_H\| \leq \frac{1 - \beta^{\gamma+1}}{1 - \beta}$ and $\|R_\gamma\| \leq \frac{\beta^{\gamma+1}}{1 - \beta}$. Then,

$$\begin{aligned} \Sigma_G &= (\Lambda_H + R_\gamma)(\Lambda_H + R_\gamma)^T \\ &= \Lambda_H \Lambda_H^T + \Lambda_H R_\gamma^T + R_\gamma \Lambda_H^T + R_\gamma R_\gamma^T \\ &= \Sigma_H + \Lambda_H R_\gamma^T + R_\gamma \Lambda_H^T + R_\gamma R_\gamma^T \end{aligned}$$

Now, defining $E_\gamma := \Sigma_G - \Sigma_H$ and taking spectral norms, we get:

$$\begin{aligned} \|E_\gamma\| &= \|\Sigma_G - \Sigma_H\| = \|\Lambda_H R_\gamma^T + R_\gamma \Lambda_H^T + R_\gamma R_\gamma^T\| \\ &\leq \|\Lambda_H R_\gamma^T\| + \|R_\gamma \Lambda_H^T\| + \|R_\gamma R_\gamma^T\| \\ &\leq 2\|\Lambda_H\| \|R_\gamma\| + \|R_\gamma\|^2 \\ &\leq 2 \left(\frac{1 - \beta^{\gamma+1}}{1 - \beta} \right) \left(\frac{\beta^{\gamma+1}}{1 - \beta} \right) + \left(\frac{\beta^{\gamma+1}}{1 - \beta} \right)^2 \\ &= \frac{2\beta^{\gamma+1} - 2\beta^{2\gamma+2}}{(1 - \beta)^2} + \frac{\beta^{2\gamma+2}}{(1 - \beta)^2} \\ &= \frac{2\beta^{\gamma+1} - \beta^{2\gamma+2}}{(1 - \beta)^2} \\ &= \frac{\beta^{\gamma+1}(2 - \beta^{\gamma+1})}{(1 - \beta)^2} \\ &= O(\beta^\gamma). \end{aligned} \tag{S1}$$

Now, suppose that not all $\sigma_i^2 = 1$. Then, following the same expansion of Σ_G as above, we have:

$$\|E_\gamma\| = \|D\| \frac{\beta^{\gamma+1}(2 - \beta^{\gamma+1})}{(1 - \beta)^2} = \sigma_{\max}^2 \frac{\beta^{\gamma+1}(2 - \beta^{\gamma+1})}{(1 - \beta)^2} = O(\beta^\gamma), \tag{S2}$$

by Assumption 2, where $\sigma_{max}^2 = \max_i \sigma_i^2$.

We now show that $|\rho(i, j \mid S)| = O(\|E_\gamma\|) = O(\beta^\gamma)$ where S is a γ -local d-separator between i and j . Let $A = \{i, j\} \cup S$ and $B = V \setminus A$. Consider the marginal precision matrix, $P := \{\Sigma_G(A, A)\}^{-1}$. Then, using the Schur complement, we can write this as

$$P = \Sigma_G^{-1}(A, A) - \Sigma_G^{-1}(A, B)\{\Sigma_G^{-1}(B, B)\}^{-1}\Sigma_G^{-1}(B, A).$$

Specifically, the partial correlation of X_i and X_j conditional on S is given by $\frac{P_{1,2}}{(P_{1,1}P_{2,2})^{1/2}} = O(P_{1,2})$, by Assumption 2.

Recall from (S1) that $\Sigma_G = \Sigma_H + E_\gamma$. Let F_γ be the matrix such that $\Sigma_G^{-1} = \Sigma_H^{-1} + F_\gamma$. Because Σ_H only considers covariance induced by paths of length at most γ , we have that $\Sigma_H^{-1}(A, A)_{1,2} = 0$.

Thus,

$$\begin{aligned} |\{\Sigma_G(A, A)\}_{1,2}^{-1}| &= |\Sigma_G^{-1}(A, A)_{1,2} - \Sigma_G^{-1}(A, B)\{\Sigma_G^{-1}(B, B)\}^{-1}\Sigma_G^{-1}(B, A)_{1,2}| \\ &= |\Sigma_H^{-1}(A, A)_{1,2} + F_\gamma(A, A)_{1,2} - \Sigma_G^{-1}(A, B)\{\Sigma_G^{-1}(B, B)\}^{-1}\Sigma_G^{-1}(B, A)_{1,2}| \\ &= |F_\gamma(A, A)_{1,2} - \Sigma_G^{-1}(A, B)\{\Sigma_G^{-1}(B, B)\}^{-1}\Sigma_G^{-1}(B, A)_{1,2}| \\ &\leq \|F_\gamma(A, A) - \Sigma_G^{-1}(A, B)\{\Sigma_G^{-1}(B, B)\}^{-1}\Sigma_G^{-1}(B, A)\|_\infty \\ &\leq \|F_\gamma(A, A) - \Sigma_G^{-1}(A, B)\{\Sigma_G^{-1}(B, B)\}^{-1}\Sigma_G^{-1}(B, A)\|. \end{aligned}$$

However, since $\Sigma_G^{-1}(A, B)\{\Sigma_G^{-1}(B, B)\}^{-1}\Sigma_G^{-1}(B, A)$ is positive semi-definite, $|\{\Sigma_G(A, A)\}_{1,2}^{-1}| \leq \|F_\gamma(A, A)\|$. We next show that $\|F_\gamma\| = O(\|E_\gamma\|) = O(\beta^\gamma)$. First, note that:

$$\begin{aligned} F_\gamma &= \Sigma_G^{-1} - \Sigma_H^{-1} \\ &= (\Sigma_H + E_\gamma)^{-1} - \Sigma_H^{-1} \\ &= \Sigma_H^{-1} - \Sigma_H^{-1}(E_\gamma^{-1} + \Sigma_H^{-1})^{-1}\Sigma_H^{-1} - \Sigma_H^{-1} && \text{(by Woodbury)} \\ &= -\Sigma_H^{-1}(E_\gamma^{-1} + \Sigma_H^{-1})^{-1}\Sigma_H^{-1} \end{aligned}$$

Then, taking spectral norms, and noting that $\Sigma_G = \Sigma_H + E_\gamma$:

$$\begin{aligned} \|F_\gamma\| &\leq \|\Sigma_H^{-1}\| \|(E_\gamma^{-1} + \Sigma_H^{-1})^{-1}\| \|\Sigma_H^{-1}\| && \text{(by sub-multiplicity)} \\ &= \|\Sigma_H^{-1}\|^2 \|E_\gamma - E_\gamma(\Sigma_H + E_\gamma)^{-1}E_\gamma\| && \text{(by Woodbury)} \\ &= \|\Sigma_H^{-1}\|^2 \|E_\gamma(I - \Sigma_G^{-1}E_\gamma)\| \\ &\leq \|\Sigma_H^{-1}\|^2 \|E_\gamma\| \|I - \Sigma_G^{-1}E_\gamma\| && \text{(by sub-multiplicity)} \\ &\leq \|\Sigma_H^{-1}\|^2 \|E_\gamma\| (1 + \|\Sigma_G^{-1}E_\gamma\|) && \text{(by triangle inequality)} \\ &\leq \|\Sigma_H^{-1}\|^2 \|E_\gamma\| (1 + \|\Sigma_G^{-1}\| \|E_\gamma\|) && \text{(by sub-multiplicity)} \\ &\leq J \|E_\gamma\|^2, && \text{(by boundedness of } \|\Sigma_G^{-1}\| \geq \|\Sigma_H^{-1}\|) \end{aligned}$$

for some constant J . Then, by walk-summability, $\|F_\gamma\| = O(\|E_\gamma\|)$. Hence, $|\rho(i, j \mid S)| = O(\beta^\gamma)$ by (S1).

■

By combining the result from either Lemma 1 or 2 with the λ -path-faithfulness assumption, we achieve oracle consistency for our algorithm given a threshold level α such that $\alpha = O(\lambda)$, $\alpha = \Omega(\beta^\gamma)$.

Next, we consider the finite sample setting, and establish a concentration inequality for sample partial correlations, under sub-Gaussian distributions, using a result from Ravikumar et al. (2011).

Lemma 13 *Assume $X = (X_1, \dots, X_p)$ is a zero-mean random vector with covariance matrix Σ such that each $X_i/\Sigma_{ii}^{1/2}$ is sub-Gaussian with parameter σ . Assume $\|\Sigma\|_\infty$ and σ are bounded. Then, the empirical partial correlation obtained from n samples satisfies, for some bounded $M > 0$:*

$$P\left(\max_{i \neq j, |S| \leq \eta} |\hat{\rho}(i, j | S) - \rho(i, j | S)| > \epsilon\right) \leq 4\left(3 + \frac{3}{2}\eta + \frac{1}{2}\eta^2\right)p^{\eta+2} \exp\left(-\frac{n\epsilon^2}{M}\right)$$

for all $\epsilon \leq \max_i(\Sigma_{ii})8(1 + 4\sigma^2)$.

Proof

Using the recursive formula for partial correlation, for any $k \in S$

$$\rho(i, j | S) = \frac{\rho(i, j | S \setminus k) - \rho(i, k | S \setminus k)\rho(k, j | S \setminus k)}{(1 - \rho^2(i, k | S \setminus k))^{1/2}(1 - \rho^2(k, j | S \setminus k))^{1/2}}$$

For example, with $S = \{k\}$, we can simplify this to:

$$\rho(i, j | S) = \frac{\rho(i, j) - \rho(i, k)\rho(k, j)}{(1 - \rho^2(i, k))^{1/2}(1 - \rho^2(k, j))^{1/2}},$$

where $\rho(i, j) = \Sigma_{ij}/(\Sigma_{ii}\Sigma_{jj})^{1/2}$.

Rewriting in terms of elements of Σ , we then decompose the empirical partial correlation deviance from the true partial correlation into the deviances of covariance terms. Here, the event of the empirical partial correlation being within ϵ distance of the true partial correlation is contained in the union of the empirical covariance terms being within $C\epsilon$ distance of the true covariance terms for a sufficiently large $C > 0$:

$$\begin{aligned} \left[|\hat{\rho}(i, j | S) - \rho(i, j | S)| > \epsilon\right] &\subset \left[|\hat{\Sigma}_{ij} - \Sigma_{ij}| > C\epsilon\right] \cup \left[|\hat{\Sigma}_{ii} - \Sigma_{ii}| > C\epsilon\right] \cup \left[|\hat{\Sigma}_{jj} - \Sigma_{jj}| > C\epsilon\right] \\ &\cup \bigcup_{k \in S} \left[|\hat{\Sigma}_{ik} - \Sigma_{ik}| > C\epsilon\right] \\ &\cup \bigcup_{k \in S} \left[|\hat{\Sigma}_{jk} - \Sigma_{jk}| > C\epsilon\right] \\ &\cup \bigcup_{k \leq k' \in S} \left[|\hat{\Sigma}_{kk'} - \Sigma_{kk'}| > C\epsilon\right], \end{aligned}$$

The number of events on the right-hand side is $3 + |S| + |S| + |S|^2 - \binom{|S|}{2}$. For $|S| \leq \eta$, the number of events is then bounded by $3 + \frac{3}{2}\eta + \frac{1}{2}\eta^2$. Then, by applying Lemma 1 in Ravikumar et al. (2011), we have that, for any i, j :

$$Pr\left(|\hat{\rho}(i, j | S) - \rho(i, j | S)| > \epsilon\right) \leq 4\left(3 + \frac{3}{2}\eta + \frac{1}{2}\eta^2\right) \exp\left(-\frac{n\epsilon^2}{K}\right),$$

for some $K > 0$, bounded when $\|\Sigma\|_\infty$ and σ are bounded. From here, the result follows. \blacksquare

Combining the results established in Lemmas 1-3, we now prove the consistency of our algorithm in the finite sample setting.

Let α denote the threshold where if $\hat{\rho}_G(i, j | S) < \alpha$, the edge (i, j) is deleted. Let G_S denote the true undirected skeleton of G , and let $S_{\eta, \gamma}$ denote the set of γ -local d-separators or size at most η .

For any $(i, j) \notin G_S$, define the false positive event as

$$F_1(i, j) = \left[\min_{S \in S_{\eta, \gamma}} |\hat{\rho}_G(i, j | S)| > \alpha \right].$$

Define

$$\theta_{max} = \max_{(i, j) \notin G_S} \min_{S \in S_{\eta, \gamma}} |\rho_G(i, j | S)|$$

and

$$\hat{\theta}_{max} = \max_{(i, j) \notin G_S} \min_{S \in S_{\eta, \gamma}} |\hat{\rho}_G(i, j | S)|.$$

Consider

$$\begin{aligned} Pr\left\{ \bigcup_{(i, j) \notin G_S} F_1(i, j) \right\} &= Pr(\hat{\theta}_{max} > \alpha) \\ &= Pr(|\hat{\theta}_{max} - \theta_{max}| > |\alpha - \theta_{max}|) \\ &= O\left[p^{\eta+2} \exp\left\{-\frac{n(\alpha - \theta_{max})^2}{M}\right\}\right] \quad (\text{by Lemma 3}) \end{aligned}$$

where $\theta_{max} = O(\beta^\gamma)$ by Lemma 1 and 2.

For any true edge $(i, j) \in G_S$, define the false negative event as

$$F_2(i, j) = \left[\min_{S \subset V \setminus \{i, j\}, |S| \leq \eta} |\hat{\rho}_G(i, j | S)| < \alpha \right].$$

Define

$$\theta_{min} = \min_{(i, j) \in G_S} \min_{S \subset V \setminus \{i, j\}, |S| \leq \eta} |\rho_G(i, j | S)|$$

and

$$\hat{\theta}_{min} = \min_{(i, j) \in G_S} \min_{S \subset V \setminus \{i, j\}, |S| \leq \eta} |\hat{\rho}_G(i, j | S)|.$$

Graph family	Expected degree	$Pr(\text{RSF})$	$Pr(\text{PF})$
Erdős-Rényi	2	0.94	0.98
Erdős-Rényi	5	0.10	0.61
Power law	2	0.95	0.97
Power law	6	0.08	0.60

Table 3: Empirical probabilities of random DAGs of size $p = 10$ satisfying faithfulness conditions; RSF refers to restricted strong faithfulness of the PC-Algorithm, and PF refers to path faithfulness of reduced PC (rPC).

Consider

$$\begin{aligned}
Pr\left\{\bigcup_{(i,j)\in G_S} F_2(i,j)\right\} &= Pr(\hat{\theta}_{min} < \alpha) \\
&= Pr(|\theta_{min} - \hat{\theta}_{min}| > |\theta_{min} - \alpha|) \\
&= O\left[p^{\eta+2} \exp\left\{-\frac{n(\alpha - \theta_{min})^2}{K}\right\}\right] \quad (\text{by Lemma 3})
\end{aligned}$$

where $\theta_{min} = \Omega(\lambda)$ by restricted path-faithfulness assumption.

Under our assumptions, we have that $n = \Omega\{(\log p)^{1/1-2c}\}$, and $\lambda = \Omega(n^{-c})$ with $c \in (0, 1/2)$. Rewriting in terms of λ , we have $n = \Omega(\frac{\log p}{\lambda^2})$. Then, by selecting α such that $\alpha = O(\lambda)$, $\alpha = \Omega(\beta^\gamma)$, we have $Pr\{\bigcup_{(i,j)\notin G_S} F_1(i,j)\} = o(1)$ and $Pr\{\bigcup_{(i,j)\in G_S} F_2(i,j)\} = o(1)$. This completes the proof of Theorem 1.

Appendix B. Additional Simulation Results

In this section, we present some simulation results comparing our algorithms to the PC-Algorithm in estimating dense graphs. The simulation setup is otherwise identical to that described in the main paper. We consider a low-dimensional setting, with $p = 100$ and $n = 200$, as well as a high-dimensional setting, with $p = 200$ and $n = 100$. For Erdős-Rényi graphs, we use a constant edge probability of 0.05, corresponding to an expected degree of 5 for the low-dimensional graph and 10 for the high-dimensional graph. For power law graphs, we use an expected degree of 6 in both graphs.

While estimation quality is worse in the dense setting for all methods, we observe in Figure 6 and Figure 7 similar trends as in the main paper. Our methods perform as well as the PC-Algorithm in estimating Erdős-Rényi graphs, and shows some improvement for power law graphs. Both versions of rPC give similar results.

We also include more simulation results comparing path faithfulness to the restricted strong faithfulness assumption. As in Section 4.2, 1000 random DAGs were generated from Erdős-Rényi and power law families, with edge weights drawn independently from a Uniform(-1, 1) distribution. Here, we consider graphs of size $p = 10$ and $p = 20$, with faithfulness condition parameters $\lambda = 0.001$ and $\eta = 2$. The results are shown in Table 3 and Table 4.

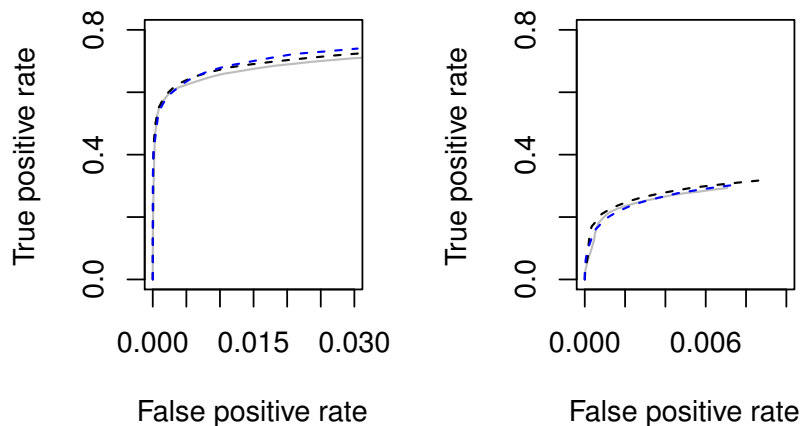


Figure 6: Average true vs. false positive rates for PC-Algorithm (grey solid line), rPC-approx (black dashed line), and rPC-full (blue dashed line) estimating Erdős-Rényi graphs. Left: $p = 100, n = 200$, average degree 5; right: $p = 200, n = 100$, average degree 10.

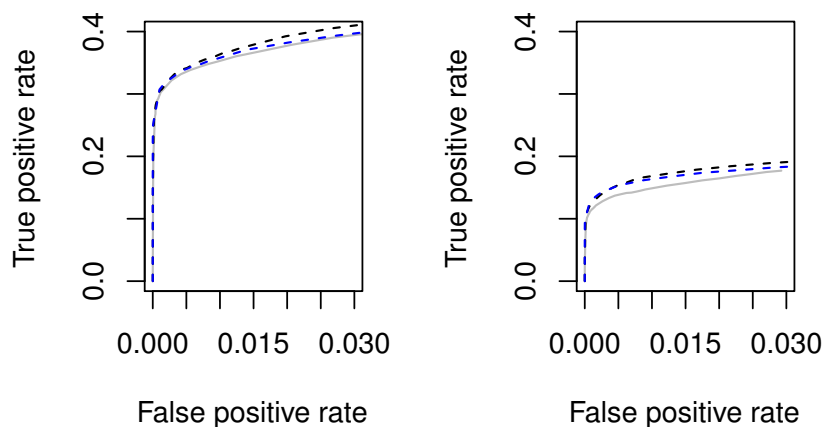


Figure 7: Average true vs. false positive rates for PC-Algorithm (grey solid line), rPC-approx (black dashed line), and rPC-full (blue dashed line) estimating power law graphs with average degree 6. Left: $p = 100, n = 200$; right: $p = 200, n = 100$.

Although it is harder for both conditions to be met as graph size and density increase, we see that path faithfulness is still more likely to be satisfied than restricted strong faithfulness, particularly with power law graphs.

Graph family	Expected degree	$Pr(\text{RSF})$	$Pr(\text{PF})$
Erdős-Rényi	2	0.66	0.86
Erdős-Rényi	5	0	0.01
Power law	2	0.04	0.51
Power law	6	0	0.01

Table 4: Empirical probabilities of random DAGs of size $p = 30$ satisfying faithfulness conditions; RSF refers to restricted strong faithfulness of the PC-Algorithm, and PF refers to path faithfulness of reduced PC (rPC).

Finally, we consider the plausibility of Assumption 4 when the SEM coefficients are not bounded by 1 in absolute value. Suppose that the data matrix X has been normalized by column-wise standard deviations. We show that this transformation preserves the original network structure, and leads to most edge weights being bounded by 1 in absolute value. Consider an edge $j \rightarrow k$ and let $W_j := \{X_i : i \in pa(k) \setminus \{j\}\}$. Then, taking the conditional covariance:

$$\begin{aligned} Cov(X_j, X_k | W_j) &= Cov \left(X_j, \rho_{jk} X_j + \sum_{i \in pa(k) \setminus j} \rho_{ik} X_i + \epsilon_k \middle| W_j \right) \\ &= \rho_{jk} Var(X_j | W_j). \end{aligned}$$

We can therefore write:

$$\rho_{jk} = \frac{Cov(X_j, X_k | W_j)}{Var(X_j | W_j)}.$$

Now, let $\tilde{X}_k = X_k / sd(X_k)$ for all k . Consider the edge weights $\tilde{\rho}_{jk}$ corresponding to the SEM for this transformed data. We have:

$$\tilde{\rho}_{jk} = \frac{sd(X_j) Cov(X_j, X_k | W_j)}{sd(X_k) Var(X_j | W_j)} = \frac{sd(X_j)}{sd(X_k)} \rho_{jk}$$

Clearly, $\tilde{\rho}_{jk} = 0$ if and only if $\rho_{jk} = 0$. Therefore, we recover the same network by applying our algorithm to the transformed data. Furthermore,

$$\begin{aligned} |\tilde{\rho}_{jk}| &= \sqrt{\frac{Var(X_j)}{Var(X_k)}} \rho_{jk}^2 \\ &= \sqrt{\frac{Var(X_j) \rho_{jk}^2}{Var(X_j) \rho_{jk}^2 + \sigma_k^2 + \sum_{i \in pa(k) \setminus \{j\}} \rho_{ik}^2 Var(X_i) + 2 \sum_{i, i' \in pa(k)} \rho_{ik} \rho_{i'k} Cov(X_i, X_{i'})}} \end{aligned}$$

Thus, $|\tilde{\rho}_{jk}| > 1$ only if

$$\sigma_k^2 + \sum_{i \in pa(k) \setminus \{j\}} \rho_{ik}^2 Var(X_i) + 2 \sum_{i, i' \in pa(k)} \rho_{ik} \rho_{i'k} Cov(X_i, X_{i'}) < 0$$

Graph family	p	Distribution	$Pr(\rho > 1)$
Erdős-Rényi	200	Uniform(-10, 10)	0.0036
	200	Normal(0, 3 ²)	0.0017
	500	Uniform(-10, 10)	0.0036
	500	Normal(0, 3 ²)	0.0012
Power law	200	Uniform(-10, 10)	0.0045
	200	Normal(0, 3 ²)	0.0045
	500	Uniform(-10, 10)	0.0034
	500	Normal(0, 3 ²)	0.0037

Table 5: Empirical average probabilities of standardized coefficients exceeding 1 in absolute value.

Rewriting the variance and covariance terms as trek sums, we obtain:

$$\sigma_k^2 + \sum_{i \in pa(k) \setminus \{j\}} \rho_{ik}^2 \sum_{\substack{\pi: u \leftrightarrow i \\ u \in V \setminus \{i\}}} \sigma_w^2 \prod_{\rho_e \in \pi} \rho_e^2 + 2 \sum_{i, i' \in pa(k)} \rho_{ik} \rho_{i'k} \sum_{\pi: i \leftrightarrow i'} \sigma_w \prod_{\rho_e \in \pi} \rho_e < 0$$

where w denotes the source or common node in the trek.

Since the first two terms in this expression will always be positive, $|\tilde{\rho}_{jk}| > 1$ only if the covariance sum is both negative and greater in magnitude than the sum of the first two terms. A term in the covariance sum will be negative if there is an odd number of negative terms in the product, and the entire sum will be negative only if there are a sufficient number of these to make the whole sum negative. Even then, the sum over squared trek products would need to be smaller than the covariance term sum in absolute value. Moreover, even if some $|\tilde{\rho}_{jk}| > 1$, the product over the entire trek may still be less than 1 if most of the other coefficients along a trek are small. Thus, we may still expect the product over the trek to be small.

Next, we provide simulation results showing that standardized SEM coefficients in graph configurations we consider rarely exceed 1 in absolute value. Specifically, we consider Erdős-Rényi and power law graphs, with $p \in \{200, 500\}$ having average degree 2. We consider edge weights from both a $\text{Unif}(-10, 10)$ and $N(0, 3^2)$ distribution. For each graph, we estimate the standardized coefficients using the `fitDag` function in the `ggm` library, and compute the proportion of edge weights that are greater than 1 in absolute value. We repeat this process over 1,000 iterations. The results are shown in Table 5. This experiment shows that, with high probability, the covariance induced over long treks between any two nodes is indeed small.

References

- Silvia Acid and Luis M de Campos. Approximations of causal networks by polytrees: an empirical study. In *Advances in Intelligent Computing*, pages 149–158. Springer, 1994.
- Animashree Anandkumar, Vincent YF Tan, Furong Huang, and Alan S Willsky. High-dimensional Gaussian graphical model selection: walk summability and local separation criterion. *Journal of Machine Learning Research*, 13(1):2293–2337, 2012.
- Cancer Genome Atlas Research Network. Comprehensive genomic characterization of squamous cell lung cancers. *Nature*, 489(7417):519–525, 2012.
- Marc RJ Carlson, Bin Zhang, Zixing Fang, Paul S Mischel, Steve Horvath, and Stanley F Nelson. Gene connectivity, function, and sequence conservation: predictions from modular yeast co-expression networks. *BMC Genomics*, 7(1):40, 2006.
- Hao Chen and Burt M Sharp. Content-rich biological network constructed by mining pubmed abstracts. *BMC Bioinformatics*, 5(1):147, 2004.
- David Maxwell Chickering. Learning Bayesian networks is NP-complete. In *Learning From Data*, pages 121–130. Springer, 1996.
- Richard Durrett. *Random Graph Dynamics*, volume 200. Cambridge University Press, 2007.
- Daniel Eaton and Kevin Murphy. Bayesian structure learning using dynamic programming and MCMC. *arXiv preprint arXiv:1206.5247*, 2012.
- Rina Foygel and Mathias Drton. Extended Bayesian information criteria for Gaussian graphical models. In *Neural Information Processing Systems (NeurIPS)*, pages 604–612, 2010.
- Nir Friedman, Michal Linial, and Iftach Nachman. Using Bayesian networks to analyze expression data. *Journal of Computational Biology*, 7:601–620, 2000.
- Min Jin Ha, Wei Sun, and Jichun Xie. PenPC: a two-step approach to estimate the skeletons of high-dimensional directed acyclic graphs. *Biometrics*, 72(1):146–155, 2016.
- S Horvath, B Zhang, M Carlson, KV Lu, S Zhu, RM Felciano, MF Laurance, W Zhao, S Qi, Z Chen, et al. Analysis of oncogenic signaling networks in glioblastoma identifies aspm as a molecular target. *Proceedings of the National Academy of Sciences*, 103(46):17402–17407, 2006.
- Juan F. Huete and Luis M. Campos. Learning causal polytrees. In *Symbolic and Quantitative Approaches to Reasoning and Uncertainty*, pages 180–185. Springer, 1993. doi:10.1007/BFb0028199.
- Trey Ideker and Nevan J Krogan. Differential network biology. *Molecular Systems Biology*, 8(1):565, 2012.

- Ronald Jansen, Haiyuan Yu, Dov Greenbaum, Yuval Kluger, Nevan J. Krogan, Sambath Chung, Andrew Emili, Michael Snyder, Jack F. Greenblatt, and Mark Gerstein. A Bayesian networks approach for predicting protein-protein interactions from genomic data. *Science*, 302(5644):449–453, 2003. ISSN 0036-8075. doi: 10.1126/science.1087361.
- Hawoong Jeong, Bálint Tombor, Réka Albert, Zoltan N Oltvai, and A-L Barabási. The large-scale organization of metabolic networks. *Nature*, 407(6804):651–654, 2000.
- Hawoong Jeong, Sean P Mason, A-L Barabási, and Zoltan N Oltvai. Lethality and centrality in protein networks. *Nature*, 411(6833):41–42, 2001.
- Markus Kalisch and Peter Bühlmann. Estimating high-dimensional directed acyclic graphs with the PC-Algorithm. *Journal of Machine Learning Research*, 8:613–636, May 2007. ISSN 1532-4435.
- Jon Kleinberg. The small-world phenomenon: an algorithmic perspective. In *ACM Symposium on Theory of Computing*, pages 163–170. ACM, 2000.
- Bide Liu, Xiao Gu, Tianbao Huang, Yang Luan, and Xuefei Ding. Identification of tmprss2-erg mechanisms in prostate cancer invasiveness: involvement of mmp-9 and plexin b1. *Oncology Reports*, 37(1):201–208, 2017.
- Po-Ling Loh and Peter Bühlmann. High-dimensional learning of linear causal networks via inverse covariance estimation. *Journal of Machine Learning Research*, 15(1):3065–3105, 2014.
- Dmitry M Malioutov, Jason K Johnson, and Alan S Willsky. Walk-sums and belief propagation in Gaussian graphical models. *Journal of Machine Learning Research*, 7(Oct):2031–2064, 2006.
- Michael Molloy and Bruce Reed. A critical point for random graphs with a given degree sequence. *Random Structures & Algorithms*, 6(2-3):161–180, 1995.
- Preetam Nandy, Alain Hauser, and Marloes H Maathuis. High-dimensional consistency in score-based and hybrid structure learning. *arXiv preprint arXiv:1507.02608*, 2015.
- Jonas Peters, Joris M Mooij, Dominik Janzing, and Bernhard Schölkopf. Causal discovery with continuous additive noise models. *Journal of Machine Learning Research*, 15:2009–2053, 2014.
- Leonie Ratz, Mark Laible, Lukasz A Kacprzyk, Stephanie M Wittig-Blaich, Yanis Tolstov, Stefan Duensing, Peter Altevogt, Sabine M Klauck, and Holger Sültmann. Tmprss2: Erg gene fusion variants induce tgf- β signaling and epithelial to mesenchymal transition in human prostate cancer cells. *Oncotarget*, 8(15):25115, 2017.
- Pradeep Ravikumar, Martin J Wainwright, Garvesh Raskutti, and Bin Yu. High-dimensional covariance estimation by minimizing l1-penalized log-determinant divergence. *Electronic Journal of Statistics*, 5:935–980, 2011.

- Christiane M Robbins, Waibov A Tembe, Angela Baker, Shripad Sinari, Tracy Y Moses, Stephen Beckstrom-Sternberg, James Beckstrom-Sternberg, Michael Barrett, James Long, Arul Chinnaiyan, et al. Copy number and targeted mutational analysis reveals novel somatic events in metastatic prostate tumors. *Genome Research*, 21(1):47–55, 2011.
- Roberta R Ruela-de Sousa, Elmer Hoekstra, A Marije Hoogland, Karla C Souza Queiroz, Maikel P Peppelenbosch, Andrew P Stubbs, Karin Pelizzaro-Rocha, Geert JLH van Leenders, Guido Jenster, Hiroshi Aoyama, et al. Low-molecular-weight protein tyrosine phosphatase predicts prostate cancer outcome by increasing the metastatic potential. *European Urology*, 69(4):710–719, 2016.
- Eric E Schadt, John Lamb, Xia Yang, Jun Zhu, Steve Edwards, Debraj GuhaThakurta, Solveig K Sieberts, Stephanie Monks, Marc Reitman, Chunsheng Zhang, et al. An integrative genomics approach to infer causal associations between gene expression and disease. *Nature Genetics*, 37(7):710–717, 2005.
- Mark Schmidt, Alexandru Niculescu-Mizil, and Kevin Murphy. Learning graphical model structure using l1-regularization paths. In *Association for the Advancement of Artificial Intelligence (AAAI)*, volume 7, pages 1278–1283, 2007.
- Marco Scutari. Learning Bayesian networks with the bnlearn r package. *Journal of Statistical Software*, 35(1):1–22, 2010. ISSN 1548-7660. doi: 10.18637/jss.v035.i03.
- A. Shojaie, A. Jauhiainen, M. Kallitsis, and G. Michailidis. Inferring Regulatory Networks by Combining Perturbation Screens and Steady State Gene Expression Profiles. *PLoS One*, 9(2):e82393, 2014.
- Ali Shojaie and George Michailidis. Penalized likelihood methods for estimation of sparse high-dimensional directed acyclic graphs. *Biometrika*, 97(3):519–538, 2010.
- P. Spirtes, C. Glymour, and R. Scheines. *Causation, Prediction, and Search*. MIT press, 2nd edition, 2000.
- Chris Stark, Bobby-Joe Breitkreutz, Teresa Regulý, Lorrie Boucher, Ashton Breitkreutz, and Mike Tyers. Biogrid: a general repository for interaction datasets. *Nucleic Acids Research*, 34(suppl 1):D535–D539, 2006.
- Seth Sullivant, Kelli Talaska, and Jan Draisma. Trek separation for Gaussian graphical models. *The Annals of Statistics*, pages 1665–1685, 2010.
- Liang Sun, Jiaju Lü, Sentai Ding, Dongbin Bi, Kejia Ding, Zhihong Niu, and Ping Liu. Hcrp1 regulates proliferation, invasion, and drug resistance via egfr signaling in prostate cancer. *Biomedicine & Pharmacotherapy*, 91:202–207, 2017.
- 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*, 41(2):436–463, 2013.

Arend Voorman, Ali Shojaie, and Daniela Witten. Graph estimation with joint additive models. *Biometrika*, 101(1):85, 2014.

Changwon Yoo, Vesteinn Thorsson, and Gregory F Cooper. Discovery of causal relationships in a gene-regulation pathway from a mixture of experimental and observational DNA microarray data. In *Pacific Symposium on Biocomputing*, volume 7, pages 498–509, 2002.

Jiji Zhang and Peter Spirtes. Strong faithfulness and uniform consistency in causal inference. In *Uncertainty in Artificial Intelligence (UAI)*, pages 632–639. Morgan Kaufmann Publishers Inc., 2002.