
Scalable Intervention Target Estimation in Linear Models

Burak Varici
Rensselaer Polytechnic Institute
varicb@rpi.edu

Karthikeyan Shanmugam
IBM Research AI
karthikeyan.shanmugam2@ibm.com

Prasanna Sattigeri
IBM Research AI
psattig@us.ibm.com

Ali Tajer
Rensselaer Polytechnic Institute
tajer@ecse.rpi.edu

Abstract

This paper considers the problem of estimating the unknown intervention targets in a causal directed acyclic graph from observational and interventional data. The focus is on soft interventions in linear structural equation models (SEMs). Current approaches to causal structure learning either work with known intervention targets or use hypothesis testing to discover the unknown intervention targets even for linear SEMs. This severely limits their scalability and sample complexity. This paper proposes a scalable and efficient algorithm that consistently identifies all intervention targets. The pivotal idea is to estimate the intervention sites from the difference between the precision matrices associated with the observational and interventional datasets. It involves repeatedly estimating such sites in different subsets of variables. The proposed algorithm can be used to also update a given observational Markov equivalence class into the interventional Markov equivalence class. Consistency, Markov equivalency, and sample complexity are established analytically. Finally, simulation results on both real and synthetic data demonstrate the gains of the proposed approach for scalable causal structure recovery. Implementation of the algorithm and the code to reproduce the simulation results are available at <https://github.com/bvarici/intervention-estimation>.

1 Introduction

Directed acyclic graphs (DAG) are commonly used for encoding the cause-effect relationships among random variables. Extensive research has been dedicated to learning the structure of DAGs from their associated observational data. Structure learning from the observational data relies on uncovering conditional independence (CI) among the random variables. Since structurally distinct DAGs can encode the same set of CI relations, a DAG is identifiable only up to its Markov equivalence class (MEC) from the observational data. Subsequently, *interventional* data can be used to further refine the MEC obtained from the observational data and learn specific causal effects.

This paper is motivated by addressing two significant independent challenges in causal discovery. First, most of the existing approaches for learning with interventional datasets require the *intervention target set* to be known, which can be a strong assumption. For instance, gene-editing technologies are known to perform cleavage at off-target genome sites [1]. Therefore, identifying the *intervened* nodes alone is a critical problem in structure learning, and despite its significance, it remains uninvestigated. Secondly, besides learning the structures of single DAGs, there exist application domains in which the goal is learning the structural changes between two related networks and their associated DAGs, or learning the sites of interventions. For instance, structural differences between the gene regulatory

networks of different subtypes of cancers can help to identify the roles of specific genes [2]. In electroencephalography analysis, the objective is to detect different brain regions that have different interactions when the subject is performing various tasks [3]. These brain regions correspond to intervened nodes in a causal graph representation. Another application area is fault detection in large-scale Internet of things and cloud applications [4]. Faulty nodes in the system can be considered as intervened nodes, and they are localized through intervention target estimation.

In practice, fixing the target variable at a specific value or removing its causal dependencies is often difficult, while disturbing the distribution of a target variable is easier [5]. The type of interventions that do not remove the causal effects are commonly observed in the real world. For example, elements of an advertising system can be modified without removing the causal effects [6]. In another example, consider molecular biology, in which the effects of infused chemicals to the cell are not set to specific values nor are they known precisely [7,8]. Therefore, we consider a *soft intervention* setting, in which we assume the conditional distributions of the target variables are changed, but no assumption is made on the causal effects. Finally, we assume that the topological ordering remains the same after the intervention procedure.

Under the soft intervention model, we propose an algorithm for estimating the intervention targets given the data from two linear SEMs associated with the observational and interventional data. Motivated by the fact that the difference of the precision matrices associated with these two models is sparse, we focus on estimating the sparse difference between precision matrices to avoid extensive conditional independence testing. This leads to a significant improvement in the computational complexity compared to those of the alternative methods. This facilitates scaling up to high-dimensional settings. Furthermore, we show that this algorithm can be used in conjunction with an observational DAG learning procedure to refine the MEC to *interventional*-MEC (\mathcal{I} -MEC). Besides being consistent in the population setting, we provide the finite-sample guarantees for linear SEM with Gaussian noise when the soft interventional changes between the two models are sparse. Our main contributions are as follows:

- We propose an algorithm that identifies intervention targets under the intervention-faithfulness assumption. We show that our algorithm identifies \mathcal{I} -MEC given the observational MEC.
- We provide the sample complexity of our algorithm under linear SEM with Gaussian noise.
- We perform experiments on both real biological and synthetic datasets to illustrate the ability to work in the high-dimensional settings and the gains compared to the relevant methods.

2 Related work

Among the broad range of approaches to intervention recovery, there exist two methods that are closely related to the scope of this paper: (i) estimating the difference between two DAGs, and (ii) learning from a combination of observational and interventional datasets.

Direct estimation of differences: Direct estimation of differences in linear SEMs has been studied recently. The study in [9] proposes a PC-style algorithm for learning changes in the edge weights by testing invariances of regression coefficients and noise variances. Even though the differences can be sparse, individual models can be dense, and estimating these variables through regression can be inaccurate. Furthermore, the number of hypothesis tests is exponential in the number of nodes that are affected by the changes, which can be prohibitive even under sparse changes in the hub nodes. Estimating the difference of two precision matrices is a relatively easier task and has received attention recently [10–13], providing finite-sample guarantees in the high-dimensional regime when such a difference is sparse. An existing study closest to the scope of our work is [14], which proposes re-estimating precision difference to progressively eliminate nodes and estimate the difference DAG. This approach critically hinges on the assumption that the noise variance is invariant, rendering limited applicability to intervention settings. In contrast, our algorithm builds on the changes in noise variances and estimates the intervention targets efficiently. We demonstrate the effect of this difference in Appendix B.1.

Learning interventional-MEC: There is a growing number of studies on causal structure learning from both observational and interventional data. Score-based greedy interventional equivalence search (GIES) [15] and hybrid interventional greedy sparsest permutation (IGSP) [16] algorithms are proposed for settings in which there are no latent confounders. Both of these algorithms assume

that the intervention targets are known. However, knowing the intervention targets can be a strong assumption since even for controlled interventional experiments, off-target effects are common. For instance, noise variances of the off-target variables can change, resulting in the intervened nodes being unknown.

For causal structure learning without the knowledge of the intervention targets, an existing study includes the dynamic programming approach in [8]. This approach has limited scalability due to its time complexity being exponential in the model size. Structural discovery from interventions proposed in [17] is a neural network-based method that can learn from interventional data without target knowledge. Nevertheless, it requires discrete data and can have at most one intervened node in a given setting, limiting its applicability. The differential causal discovery from interventional data (DCDI) algorithm proposed in [18] extends the approach of [17] without making as strong assumptions, e.g., learning a distribution over all potential interventional families via continuous optimization. While this method is shown to converge to the real intervention targets, its runtime becomes prohibitive even in models with as few as 100 nodes. The Unknown-target IGSP (UT-IGSP) algorithm proposed in [19] learns the intervention targets simultaneously while learning the causal structure. Even though interventions refine the search space, a greedy search of the sparsest permutation is still too slow in the high-dimensional regime, especially when using non-Gaussian CI tests. A graphical characterization of soft interventions with unknown targets is proposed in [7] for causally insufficient systems. This algorithm, however, relies on CI tests and it is not scalable. We focus on causally sufficient systems in this paper.

3 Problem definition

Let $\mathcal{G} \triangleq ([p], E)$ be a DAG with the node set $[p] \triangleq \{1, \dots, p\}$ and the edge set $E \subseteq [p] \times [p]$. We denote the directed edge from $i \in [p]$ to $j \in [p]$ by $i \rightarrow j$. We associate a random variable X_i with $i \in [p]$, and accordingly, define the random vector $X \triangleq (X_1, \dots, X_p)^\top$. We consider a linear SEM, according to which

$$X = B^\top X + \epsilon, \quad (1)$$

where $B \in \mathbb{R}^{p \times p}$ is the autoregressive matrix in which $B_{ij} \neq 0$ if and only if $i \rightarrow j$ in \mathcal{G} . The random vector $\epsilon \in \mathbb{R}^{p \times 1}$ has zero-mean and covariance matrix $\Omega \triangleq \text{diag}(\sigma_1^2, \dots, \sigma_p^2)$. We denote the covariance matrix of X by Σ and its inverse (the precision matrix) by Θ , which satisfies $\Theta = (I - B)\Omega^{-1}(I - B)^\top$. For entries of Θ , we have

$$\Theta_{ij} = -\frac{B_{ij}}{\sigma_j^2} - \frac{B_{ji}}{\sigma_i^2} + \sum_{k \in \text{ch}(i) \cap \text{ch}(j)} \frac{B_{ik}B_{jk}}{\sigma_k^2}, \quad \forall i \neq j, \quad (2)$$

$$\text{and } \Theta_{ii} = \sigma_i^{-2} + \sum_{j \in \text{ch}(i)} \sigma_j^{-2} B_{ij}^2, \quad \forall i \in [p], \quad (3)$$

where $\text{ch}(i)$, $\text{pa}(i)$, $\text{de}(i)$, and $\text{an}(i)$ denote children, parents, descendants, and ancestors set of node i in DAG \mathcal{G} , respectively.

From the observational data, a DAG can be learned only up to its MEC [20]. Interventions are used to increase the identifiability of a DAG by removing all causes of the intervention target (*perfect intervention*) or modifying those relationships without removing them completely (*imperfect intervention*). We consider the following *soft intervention* setting, which does not remove causal effects (from direct parents) on intervention target nodes, and hence, is more practical.

Soft intervention model. In this model, interventions correspond to disturbing the target nodes $i \in \mathcal{I}$ by changing the variances of their noise variables, while the cause weights, i.e., the weights $B_{\text{pa}(i), i} \triangleq \{B_{j,i} : j \in \text{pa}(i)\}$, can vary freely. This intervention procedure on the initial DAG results in a second DAG with new parameters.

Let $\mathcal{G}^{(1)}$ represent a linear SEM prior to intervention with parameters $B^{(1)}, \epsilon^{(1)}$, and $\mathcal{G}^{(2)}$ be the linear SEM after the intervention, with parameters $B^{(2)}, \epsilon^{(2)}$. The intervention target set that relates these two DAGs is

$$\mathcal{I} \triangleq \{i : \sigma_i^{(1)} \neq \sigma_i^{(2)}\}. \quad (4)$$

Accordingly, denote the covariance and precision matrices of these two models by $\Sigma^{(1)}, \Sigma^{(2)}, \Theta^{(1)}$, and $\Theta^{(2)}$. Accordingly, denote the differences between the two models by $\Delta_B \triangleq B^{(1)} - B^{(2)}$ and $\Delta_\Theta \triangleq \Theta^{(1)} - \Theta^{(2)}$. For a subset of nodes $S \subseteq [p]$, denote the precision matrix of the random vector $X_S \triangleq \{X_i : i \in S\}$ by Θ_S . We also denote the set of changed nodes by $S_\Delta \triangleq \{k : (\Delta_\Theta)_{k,k} \neq 0\}$ and denote its size by $p_\Delta \triangleq |S_\Delta|$. According to (3), this set consists of all the intervened nodes and their parents.

In this paper, we estimate the intervention targets \mathcal{I} given the data from two SEMs under the soft intervention model. Furthermore, we estimate the non-intervened parents of the targets, which update the given observational MEC into the \mathcal{I} -MEC. Formally, we define $\psi : \Sigma^{(1)} \times \Sigma^{(2)} \rightarrow (\hat{\mathcal{I}}, \text{pa}(\hat{\mathcal{I}}))$ as the estimator that maps the covariance matrices of the observational and interventional data to an intervention target set estimate and their parents. We aim to maximize the probabilities of ψ recovers \mathcal{I} and their parents. To this end, we define

$$P \triangleq \mathbb{P}(\mathcal{I} = \hat{\mathcal{I}}) \quad \text{and} \quad Q \triangleq \mathbb{P}(\text{pa}(\mathcal{I}) = \text{pa}(\hat{\mathcal{I}})). \quad (5)$$

We will show that the algorithm proposed in Section 4 exactly recovers \mathcal{I} and the non-intervened parents of the members of \mathcal{I} . Hence, given the observational MEC, we obtain the \mathcal{I} -MEC.

4 Algorithm and main results

In this section, we provide our proposed algorithm and present the attendant performance guarantees. Our algorithm involves repeatedly estimating the difference of the precision matrices to find the intervention target set \mathcal{I} , or equivalently, its complement \mathcal{I}^C . This algorithm consists of three key steps. In **Step 1**, instead of directly estimating \mathcal{I} , we aim to discard the nodes that are strongly deemed not to belong to \mathcal{I} . For this purpose, we start by identifying the non-intervened nodes that do not have intervened children. These nodes are not of interest and they are discarded from further consideration. We achieve this by estimating difference of precision matrices corresponding to the complete model with variables $[p]$, and we denote the changed nodes in the diagonal of this precision difference matrix with S_Δ . We continue with only the nodes contained in S_Δ for further scrutiny. A naive approach to identifying the rest of the non-intervened nodes is computing Δ_Θ exhaustively for all the 2^{p_Δ} subsets of S_Δ , which for large p_Δ is computationally prohibitive. Alternatively, we partition S_Δ into two sets: the set of *non-intervened source* nodes and its complement. We feed these two partitions for further processing to Step 2. Since the distribution of a node relies only on its ancestors, reaching a topological ordering is critical to reduce the complexity. In **Step 2**, for each source node, we find the nodes that share a common ancestor with it. Subsequently, we decompose S_Δ into equivalence classes according to these ancestral relationships with non-intervened source nodes. Decomposing S_Δ allows us to order the equivalence classes according to a topological ordering. In **Step 3**, we process these classes individually, and show that we only need to compute Δ_Θ for all the subsets of each class considered in Step 3. This results in a significant reduction in the computational complexity compared to the exhaustive search approach. Finally, we identify the non-intervened parents of \mathcal{I} from the earlier results. Since estimating the precision difference appears in our algorithm repeatedly, we describe it next and then continue with the details of the three steps.

Precision difference estimation (PDE). When the difference between two SEMs is sparse, estimating the difference of their precision matrices can be formulated as a Lasso-type problem and solved efficiently. Since it will be used repeatedly, it is important for this function to recover the support of Δ_Θ and have a feasible computational complexity. The study in [12] solves the following convex problem through the alternating direction method of multipliers (ADMM) to estimate $\Delta_\Theta = \Theta^{(1)} - \Theta^{(2)}$:

$$\hat{\Delta}_\Theta = \underset{\Delta_\Theta}{\text{argmin}} \left\{ \frac{1}{2} \text{Tr}(\Delta_\Theta^\top \hat{\Sigma}^{(1)} \Delta_\Theta \hat{\Sigma}^{(2)}) - \text{Tr}(\Delta_\Theta (\hat{\Sigma}^{(1)} - \hat{\Sigma}^{(2)})) + \lambda \|\Delta_\Theta\|_1 \right\}, \quad (6)$$

where λ is a tuning parameter. The computational complexity of this algorithm grows according to $O(p^3)$, which overcomes the limitation of the algorithm of [10] used in [14], which has complexity $O(p^4)$. Therefore, we use this ADMM-based method as our PDE procedure.

Step 1: Finding the non-intervened source nodes. By leveraging the PDE procedure discussed, we first estimate Δ_Θ over all $[p]$ nodes and obtain S_Δ . Representation of $\Theta_{i,i}$ in (3) implies that

a diagonal entry of Δ_Θ is non-zero if and only if either its corresponding node is in \mathcal{I} or it has a child in \mathcal{I} . Hence, the set $[p] \setminus S_\Delta$ contains only non-intervened nodes and can be discarded from consideration. Furthermore, some non-intervened nodes in S_Δ , namely those that do not have intervened ancestors, do not observe a change in their statistics. Therefore, it is possible to identify them from covariance matrices $\Sigma^{(1)}$ and $\Sigma^{(2)}$. Subsequently, these nodes can serve as the starting points for distinguishing the rest of the \mathcal{I}^C in S_Δ . Let us define them as the *non-intervened source nodes*, denoted by

$$J_0 \triangleq \{j : j \in S_\Delta, j \notin \mathcal{I}, \text{an}_{\mathcal{I}}(j) = \emptyset\}. \quad (7)$$

The outputs of this step, S_Δ and J_0 , are fed into the next steps of the algorithm.

Step 2: Forming equivalence classes from J_0 . We will show that for any non-intervened node $j \in S_\Delta$, there exists a minimal subset of S_Δ , which makes the corresponding diagonal entry of the precision matrix invariant, and it does not contain any descendant of j . Therefore, the non-intervened nodes that have the same ancestors are affected by the interventions similarly, and finding their ancestors is critical. We show that determining whether node $k \in S_\Delta \setminus J_0$ has a common ancestor with node $j \in J_0$ is possible by applying PDE on $\{j, k\}$. Accordingly, we define the *source ancestral set* J_0^k for each node $k \in S_\Delta \setminus J_0$ as

$$J_0^k \triangleq \{j : j \in J_0, \text{an}(j) \cap \text{an}(k) \neq \emptyset\}, \quad \forall k \in S_\Delta \setminus J_0. \quad (8)$$

Next, we decompose the set $S_\Delta \setminus J_0$ into *equivalence classes* where all the nodes in a class have the same source ancestral set. We denote these equivalence classes by $\mathcal{A}_1, \dots, \mathcal{A}_L$, and the source ancestral set corresponding to a class \mathcal{A}_ℓ by $J_0^{\mathcal{A}_\ell}$ for $\ell \in [L]$. These classes are ordered according to a topological order such that for $1 \leq \ell < \ell' \leq L$, $J_0^{\mathcal{A}_{\ell'}} \not\subseteq J_0^{\mathcal{A}_\ell}$. In other words, the class corresponding to the superset of any $J_0^{\mathcal{A}_\ell}$ should appear later than \mathcal{A}_ℓ in the sequence $\mathcal{A}_1, \dots, \mathcal{A}_L$. Source ancestral sets and equivalence classes are fed into the next step.


Step 3: Processing equivalence classes. Given Step 1 and Step 2, we can describe our exact search space of subsets for Δ_Θ estimates to declare whether a node is intervened. We process equivalence classes $\mathcal{A}_1, \dots, \mathcal{A}_L$ individually, i.e., at stage ℓ , we consider the nodes in \mathcal{A}_ℓ . Let us define $\mathcal{M}_\ell \triangleq J_0 \cup \bigcup_{1 \leq b < \ell} \mathcal{A}_b$. It suffices to estimate $\Delta_{\Theta, \mathcal{M}_\ell \cup A}$ **only for each** $A \subseteq \mathcal{A}_\ell$ to determine the intervention status of any node in \mathcal{A}_ℓ class. This key observation reduces the number of PDE steps needed. Specifically, for any non-intervened $j \in \mathcal{A}_\ell$, there exists a subset $A \subseteq \mathcal{A}_\ell$ such that the corresponding diagonal entry of $\Delta_{\Theta, \mathcal{M}_\ell \cup A}$ will be zero, and there does not exist any such set for the intervened nodes in \mathcal{A}_ℓ . Formally, the *process equivalence class* returns

$$\mathcal{I}_\ell = \{i : i \in \mathcal{A}_\ell \cap \mathcal{I}\}, \quad \text{and} \quad J_\ell = \{j : j \in \mathcal{A}_\ell \cap \mathcal{I}^C\}. \quad (9)$$

Finally, we identify the non-intervened parents of the intervened nodes without any new Δ_Θ estimates.

Computational complexity. Algorithm 1 repeatedly performs PDE in each step. The number of required instances of PDE is $(p_\Delta + 1)$ in Step 1, $O(p_\Delta^2)$ in Step 2, and $\sum_{\ell \in [L]} 2^{|\mathcal{A}_\ell|}$ in Step 3. Hence, it grows exponentially with $\max_{\ell \in [L]} |\mathcal{A}_\ell|$, which can be p_Δ in the worst case in extreme examples. Nevertheless, in almost all practical scenarios it is usually considerably smaller. To provide some insights, we provide the next example and relegate more discussions to Appendix B.4.

Example 1. Consider a DAG with nodes $\{1, 2, 3, 4, 5\}$ and the edge set $\{1 \rightarrow 3, 3 \rightarrow 4, 2 \rightarrow 4, 2 \rightarrow 5, 4 \rightarrow 5\}$ and let $\mathcal{I} = \{3, 5\}$. Hence, we have $S_\Delta = \{1, 2, 3, 4, 5\}$, $J_0 = \{1, 2\}$, $J_0^3 = \{1\}$, $J_0^4 = \{1, 2\}$, $J_0^5 = \{1, 2\}$, and accordingly, $\mathcal{A}_1 = \{3\}$, $\mathcal{A}_2 = \{4, 5\}$, $J_0^{\mathcal{A}_1} = \{1\}$, and $J_0^{\mathcal{A}_2} = \{1, 2\}$. Note that the largest \mathcal{A} class has 2 nodes whereas S_Δ has 5 nodes.



Restricted SEM. For a linear SEM \mathcal{G} with (B, ϵ) , we denote the restricted SEM (see Lemma 1 for details) that characterizes the relationship among the random variables X_S for a set S by (B_S, ϵ_S) . As defined earlier, the corresponding precision matrix is denoted by Θ_S . The entries of B_S and noise variances σ_S are characterized by the original values of B , σ , and Θ .

Remark 1 We remark that the invariance of the distributions for the noise term of a node and the value of the node are equivalent only for the non-intervened nodes that do not have an intervened ancestor. Therefore, only such nodes can be detected from the full linear SEM. The noise term of a

Algorithm 1 Causal Intervention Target Estimator (CITE)

- 1: **Input:** $\Sigma^{(1)}$ and $\Sigma^{(2)}$
- 2: **Output.** \mathcal{I} intervention set
- 3: Estimate $\Delta_{\Theta} \leftarrow \text{precision difference}(\Sigma^{(1)}, \Sigma^{(2)})$
- 4: Form $S_{\Delta} \triangleq \{k : k \in [p], [\Delta_{\Theta}]_{k,k} \neq 0\}$
- 5: Form J_0 using (7) and J_0^k for each $k \in S_{\Delta} \setminus J_0$ using (8)
- 6: Form equivalence classes $\mathcal{A}_1, \dots, \mathcal{A}_L$.
- 7: **for** $\ell \in [L]$ **do**
- 8: Take \mathcal{A}_{ℓ} set and the corresponding $J_0^{\mathcal{A}_{\ell}}$ set
- 9: $\mathcal{B}_{\ell} \leftarrow \{b : J_0^{\mathcal{A}_b} \subset J_0^{\mathcal{A}_{\ell}}, 1 \leq b < \ell\}$
- 10: $\mathcal{M}_{\ell} = J_0^{\mathcal{A}_{\ell}} \cup \bigcup_{b \in \mathcal{B}_{\ell}} \mathcal{A}_b$
- 11: $J_{\ell}, \mathcal{I}_{\ell} \leftarrow \text{process equivalence class}(\mathcal{M}_{\ell}, \mathcal{A}_{\ell}, \Sigma^{(1)}, \Sigma^{(2)})$
- 12: **end for**
- 13: $\mathcal{I} = \bigcup_{\ell \in [L]} \mathcal{I}_{\ell}$
- 14: $\text{pa}(\mathcal{I}) \leftarrow \text{parent finder}(\mathcal{I}, \mathcal{J}, \mathcal{M}_1, \mathcal{A}_1, \dots, \mathcal{M}_L, \mathcal{A}_L, \Sigma^{(1)}, \Sigma^{(2)})$
- 15: **Return** \mathcal{I} and $\text{pa}(\mathcal{I})$

Algorithm 2 Functions for the main algorithm

Precision Difference Estimation (PDE) $(\Sigma^{(1)}, \Sigma^{(2)})$

- 1: Using ADMM based algorithm of [12] method estimate $\Delta_{\Theta} = (\Sigma^{(1)})^{-1} - (\Sigma^{(2)})^{-1}$
- 2: Symmetrize Δ_{Θ} : set $\Delta_{\Theta} = (\Delta_{\Theta} + \Delta_{\Theta}^{\top})/2$
- 3: Threshold Δ_{Θ} : set $[\Delta_{\Theta}]_{i,j} = 0$ if $|[\Delta_{\Theta}]_{i,j}| < \varepsilon$.
- 4: **Return** Δ_{Θ}

Process Equivalence Class $(\mathcal{M}, \mathcal{A}, \Sigma^{(1)}, \Sigma^{(2)})$

- 1: For each subset $A \subseteq \mathcal{A}$, estimate $\Delta_{\Theta_{\mathcal{M} \cup A}} \leftarrow \text{precision difference}((\Sigma^{(1)})_{A,A}, (\Sigma^{(2)})_{A,A})$
- 2: **for** $k \in \mathcal{A}$ **do**
- 3: **if** $\exists A \subseteq \mathcal{A}$, where $k \in A$, and $[\Delta_{\Theta_{\mathcal{M} \cup A}}]_{k,k} = 0$ **then**
- 4: $J \leftarrow J \cup k$
- 5: **else**
- 6: $\mathcal{I} \leftarrow \mathcal{I} \cup k$
- 7: **end if**
- 8: **end for**
- 9: **Return** J, \mathcal{I}

Parent Finder $(\mathcal{I}, \mathcal{J}, \mathcal{M}_1, \mathcal{A}_1, \dots, \mathcal{M}_L, \mathcal{A}_L, \Sigma^{(1)}, \Sigma^{(2)})$

- 1: **for** $i \in \mathcal{I}$ **do**
- 2: $c_i \leftarrow c_i \triangleq \ell : i \in \mathcal{A}_{\ell}$
- 3: **for** $j \in \mathcal{M}_{c_i}$ **do**
- 4: **if** $\nexists A \subseteq \mathcal{A}_{c_i}$ such that $[\Delta_{\Theta_{\mathcal{M}_{c_i} \cup A}}]_{j,i} = 0$ **then**
- 5: Add j to $\text{pa}(i)$
- 6: **end if**
- 7: **end for**
- 8: **end for**
- 9: **Return** $\text{pa}(i)$ for $i \in \mathcal{I}$

non-intervened node maintains its invariance in a restricted SEM in which we keep its intervened ancestors and their parents. On the other hand, the noise term of an intervened node is always variant for any choice of restricted SEM.

In the following subsections, we will provide different analytical guarantees of Algorithm 1. Specifically, we will comment on the consistency of \mathcal{I} recovery, the refinement of MEC to \mathcal{I} -MEC, and the sample complexity.

4.1 Consistency of \mathcal{I} recovery

We provide the consistency of Algorithm 1 for estimating \mathcal{I} in this subsection. First, we need the following assumption to ensure that interventions are successful.

Assumption 1 (\mathcal{I} -faithfulness) For any choice of $i, j \in S \subseteq [p]$, we have the following properties:

1. If $\sigma_i^{(1)} \neq \sigma_i^{(2)}$, then $\sigma_{S,i}^{(1)} \neq \sigma_{S,i}^{(2)}$.
2. If $\sigma_{S,i}^{(1)} \neq \sigma_{S,i}^{(2)}$, then $[\Theta_S^{(1)}]_{i,i} \neq [\Theta_S^{(2)}]_{i,i}$.
3. If $\sigma_{S,i}^{(1)} \neq \sigma_{S,i}^{(2)}$ and $[B_S]_{j,i} \neq 0$ in either model, then $[\Theta_S^{(1)}]_{i,j} \neq [\Theta_S^{(2)}]_{i,j}$.

Next, we characterize the parameters of a restricted SEM, and then formalize the observations stated in Remark 1 in the subsequent proposition.

Lemma 1 ([14]) Corresponding to a subset $S \subseteq [p]$, denote the removed set of nodes by $U \triangleq [p] \setminus S$ and define $U_j \triangleq U \cap \text{an}(j)$, for $j \in S$. We have

$$\sigma_{S,j}^2 = \sigma_j^4 \left(\sigma_j^2 - B_{U_j,j}^\top [\Theta_{\text{an}(j)}]_{U_j,U_j}^{-1} B_{U_j,j} \right)^{-1}, \quad (10)$$

$$[B_S]_{k,j} = \frac{\sigma_{S,j}^2}{\sigma_j^2} \left(B_{k,j} - B_{U_j,j}^\top [\Theta_{\text{an}(j)}]_{U_j,U_j}^{-1} [\Theta_{\text{an}(j)}]_{U_j,k} \right). \quad (11)$$

Proposition 1 Denote the ancestors of $j \notin \mathcal{I}$ in \mathcal{I} by $\text{an}_{\mathcal{I}}(j)$. If a set S contains $\text{an}_{\mathcal{I}}(j)$ and their parents $\text{pa}(\text{an}_{\mathcal{I}}(j))$, then $\sigma_{S,j}^{(1)} = \sigma_{S,j}^{(2)}$. Furthermore, for $i \in \mathcal{I}$ and any set S we have $[\Delta_{\Theta_S}]_{i,i} \neq 0$. Additionally, if $[B_S]_{j,i} \neq 0$ in either model, then we have $[\Delta_{\Theta_S}]_{j,i} \neq 0$.

Remark 2 We repeatedly use the restricted SEM characterization in Lemma 1 with various strategic choices of subsets S in Algorithm 1 to eliminate the non-intervened nodes from S_Δ using the criterion of Proposition 1. In Step 1, to identify J_0 in (7), we set $S = \{j\}$ for each $j \in S_\Delta$. In Step 2, to identify J_0^k in (8) for each $k \in S_\Delta \setminus J_0$, we set $S = \{j, k\}$ for each $j \in J_0$. In Step 3, to process the nodes in \mathcal{A}_ℓ , we use subsets of the form $\mathcal{M}_\ell \cup A$ for subsets A in \mathcal{A}_ℓ .

Theorem 1 (Consistency) If Assumption 1 holds and covariance estimates are perfect, then Algorithm 1 consistently estimates soft interventions with $\mathbb{P} = 1$.

4.2 \mathcal{I} -Markov equivalence

Interventions in a DAG change the conditional distributions of the intervened variables, and hence, they reveal orientations of some edges that were previously undirected in observational CPDAG, resulting in the interventional CPDAG (\mathcal{I} -CPDAG). The DAGs that have the same \mathcal{I} -CPDAG under soft intervention \mathcal{I} form the \mathcal{I} -Markov equivalence class (\mathcal{I} -MEC). This is shown and discussed next.

For a DAG \mathcal{G} and an intervention set \mathcal{I} , an additional \mathcal{I} -vertex ζ and corresponding \mathcal{I} -edges $\{\zeta \rightarrow i\}_{i \in \mathcal{I}}$ are added to form the interventional DAG (\mathcal{I} -DAG). Note that vertex ζ creates a new v-structure $\zeta - i - j$ for any non-intervened $j \in \text{pa}(i)$. However, if j is also in \mathcal{I} , then \mathcal{I} -DAG also contains the $\zeta \rightarrow j$ edge, and there is no new v-structure that can orient the edge $i - j$.

We call the edges in \mathcal{G} that are not directed in the original CPDAG but are directed in \mathcal{I} -CPDAG as \mathcal{I} -directed edges. In the *parent finder* step of Algorithm 1, we find the edge set $\{j \rightarrow i\}_{j \notin \mathcal{I}, i \in \mathcal{I}}$, and subsequently, obtain the \mathcal{I} -MEC. Therefore, we can use Algorithm 1 in conjunction with an observational algorithm to perform causal structure learning, and establish the following theorem.

Theorem 2 (\mathcal{I} -MEC) If Assumption 1 holds and covariance estimates are perfect, then Algorithm 1 consistently recovers non-intervened parents of an $i \in \mathcal{I}$ with $\mathbb{Q} = 1$. This result modifies the original MEC, which can be obtained via any observational structure learning algorithm, into the \mathcal{I} -MEC.




Figure 1: Average precision of estimating intervention targets. Algorithm 1 reaches high precision with increasing number of samples even in large models for all settings.

4.3 Sample complexity

In this subsection, we provide the finite-sample counterparts of Theorem 1 and Theorem 2. Our choice of the PDE algorithm, the ADMM-based method of [12], enjoys finite-sample results when noise ϵ has a Gaussian distribution. The following theorem establishes the sample complexity of Algorithm 1 for estimating \mathcal{I} and the non-intervened parents of the nodes in \mathcal{I} .

Theorem 3 (Sample complexity) *Let d denote the maximum degree of an intervened node and set $\Gamma \triangleq \Sigma^{(2)} \otimes \Sigma^{(1)}$ and $\alpha \triangleq 1 - \max_{e \notin \text{supp}} |\Gamma_{e, \text{supp}} \Gamma_{\text{supp}, \text{supp}}^{-1}|_1$. Accordingly, define $M \triangleq \max\{\|\Sigma^{(1)}\|_\infty, \|\Sigma^{(2)}\|_\infty\}$, $M_\Sigma \triangleq \max\{\|\Sigma^{(1)}\|_{1, \infty}, \|\Sigma^{(2)}\|_{1, \infty}\}$, $M_{\Gamma, \Gamma^\tau} \triangleq \max\{\|\Gamma_{S, S}\|_{1, \infty}, \|\Gamma_{S, S}^T\|_{1, \infty}\}$, where S is the support of $(\Sigma^{(2)})^{-1} - (\Sigma^{(1)})^{-1}$. When $\alpha > 0$ and $M_\Sigma M_{\Gamma, \Gamma^\tau} < +\infty$, with $n = O\left(\frac{d^4 \log p}{\epsilon^2 \delta}\right)$ samples, Algorithm 1*

1. identifies \mathcal{I} with a probability at least $\geq 1 - \delta$;
2. identifies the non-intervened parents $\{j \rightarrow i\}_{j \notin \mathcal{I}, i \in \mathcal{I}}$ with a probability at least $\geq 1 - \delta$.

Note that we have assumed that the product $M_\Sigma M_{\Gamma, \Gamma^\tau}$ is bounded. This is necessary to avoid a linear scaling of the sample complexity in p . More discussion on the necessity and implications of this assumption is provided in the proof of Theorem 3 in Appendix A.

5 Empirical results

5.1 Synthetic data - intervention recovery

We start by testing our algorithm for estimating intervention targets, i.e., the set \mathcal{I} . We generate 100 realizations of Erdős-Rényi [21] DAGs with expected neighborhood size $c = 1.5$, and $|\mathcal{I}| = 5$. We sample the entries of B , i.e., the edge weights, independently at random according to the uniform distribution on $[-1, -0.25] \cup [0.25, 1]$. The additive Gaussian noise terms have distribution $\mathcal{N}(0, I_p)$. We select the intervention set \mathcal{I} by randomly selecting 5 nodes from $[p]$. We consider three different models to intervene on the nodes in \mathcal{I} : (i) *shift intervention* model in which mean of the noise ϵ_i is shifted from 0 to 1, (ii) *variance increase* model in which the variance of the noise ϵ_i is increased from 1 to 2, and (iii) *randomized intervention* model, in which $(B^{(2)})_{\text{pa}(i), i} = 0$ and the noise variance varies from 1 to 1.5. All the simulations are run on a MacBook Pro with 2.7 GHz Dual-Core i5 core and 8 GB RAM.

We first run our algorithm by varying the graph size p and the number of samples. Figure 1 illustrates that our algorithm is able to recover the intervention targets with high precision under all three intervention models. Having high precision is especially important in high dimensions, since a large false positive rate severely affects any downstream task such as structure learning. Recall rates are close to 1 and they are omitted from the graph.

Next, we compare our results with that of the UT-IGSP algorithm [19] for the shift intervention model. We note that UT-IGSP performs a greedy search to identify the sparsest permutation through CI tests, and it returns intervention targets as a by-product along with the learned causal structure. The computation time of UT-IGSP, hence, grows quickly with the size of the graph, reaching an

average of 61.2 seconds for $p = 100$. Therefore, the complexity of structure learning and intervention target discovery in the high-dimensional regime is prohibitive. In contrast, our algorithm has comparable performance to UT-IGSP when $p = 100$, while requiring less than a second of runtime. Our algorithm’s runtime scales gracefully when the dimension is in the hundreds.

Table 1: \mathcal{I} estimation in the shift intervention model - 50 repetitions with 5000 samples - density 1.5

p	UT-IGSP ([19])			Algorithm 1		
	Precision	Recall	Time(s)	Precision	Recall	Time(s)
40	0.99 (0.04)	0.99 (0.04)	0.8	0.96 (0.09)	0.94 (0.09)	0.1
60	0.95 (0.07)	0.99 (0.05)	5.2	0.97 (0.07)	0.95 (0.10)	0.2
80	0.96 (0.08)	0.99 (0.04)	17.8	0.96 (0.08)	0.96 (0.10)	0.3
100	0.93 (0.11)	1 (0)	61.2	0.94 (0.09)	0.98 (0.07)	0.3

5.2 Synthetic data - causal structure learning

In Section 4.2, we have shown that Algorithm 1 recovers the intervention targets. It can be further used to refine the observational MEC to an \mathcal{I} -MEC. Accordingly, we take the correct CPDAG of $\mathcal{G}^{(1)}$, and then apply our algorithm’s findings to obtain \mathcal{I} -CPDAG. We report the accuracy of additional edge orientations and in particular recovery of parents (if possible) of intervention targets in Appendix B.2.

5.3 Application to real data

We apply our algorithm to two real datasets with observational and interventional data to learn their causal structures. When there exist multiple interventional environments, we apply our algorithm to them individually to estimate the intervened nodes and their parents. Subsequently, we combine the results from all environments in order to form the final estimated structure. There is a large number of interventional settings in which finding the targets and their non-intervened parents by our algorithm yields a good estimate of the associated DAG, which we use for our evaluation.

We compare our results with those of the algorithms UT-IGSP, and UT-IGSP* in [19]¹, where the former works with partially known intervention targets and the latter does not require any target input. We use both parametric (Gaussian) and non-parametric (Hilbert-Schmidt independence criterion) CI tests for UT-IGSP methods. We note that the non-parametric tests have significant runtimes. We note that our algorithm uses PDE at several stages, which calls for different λ regularization parameters. Namely, let us denote the parameters used for Step 1 and Step 3 by λ_1 , Step 2 by λ_2 and the *parent finder* of Step 3 by λ_3 . Similarly, UT-IGSP needs a cut-off value α for CI tests. We run the algorithms with different values of these parameters to obtain the receiver operating characteristic (ROC) curves.

Protein signaling data. We first consider the dataset in [22] for discovering the protein signaling network of 11 nodes. It consists of measurements of proteins and phospholipids under different interventional environments. In each environment, signaling nodes are inhibited or activated. Hence, these sites form intervention targets. The conventionally accepted ground truth has been updated over the years, and we compare with the recent version in [23], which consists of 16 edges. We follow the process of [19] and work with 1755 observational and 4091 interventional samples aggregated from 5 different interventional environments. In Fig. 2a, we report the results of running Algorithm 1 and UT-IGSP with various parameters.

Perturb-seq gene expression data. We analyze the performance of our algorithm on the perturb-seq dataset by in [24]. The dataset consists of observational data and the interventional data from bone marrow-derived dendritic cells (BMDCs). A single gene has been targeted for deletion in each interventional environment. Similarly to [24], we have focused on 24 genes that are known to have regulatory effects, and we have followed [16] to select interventional data from 8 gene deletions along with observational samples. We take the Fig. 4D in [24] as the ground truth, which has 34

¹The code and preprocessed real data are taken from <https://github.com/csquires/utigsp> for fair comparison, and CausalDAG package which is distributed under 3-Clause BSD licence is used.

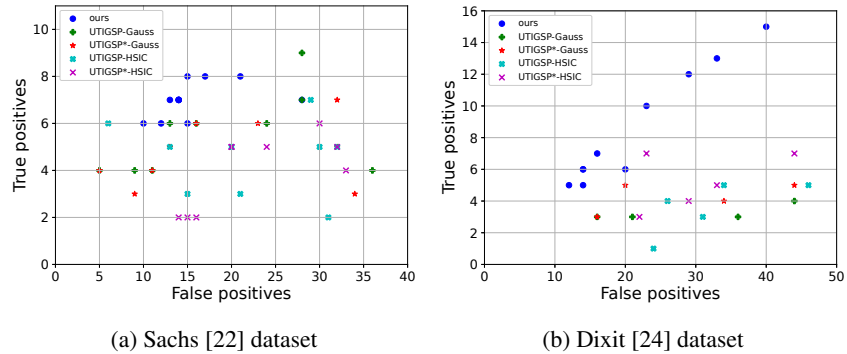


Figure 2: ROC curves for directed edge recovery. Algorithm 1: (a) is more robust in smaller graphs; (b) handles larger graphs and multiple interventional settings successfully, and conforms to real datasets better than CI testers.

edges among 24 nodes. We use 23 interventional settings for the targeted 8 gene deletions. In Fig. 2b, we plot the results of running Algorithm 1 and UT-IGSP with various parameters.

In both real datasets, our algorithm achieves higher accuracy in recovering directed edges. The comparison with UT-IGSP is more striking in Fig. 2b, and shows that our ability to work with many interventional environments on a relatively larger graph. Furthermore, Fig. 2a shows that Algorithm 1 handles smaller graphs more robustly. These results illustrate that even though our algorithm is designed for linear models, on real datasets, it performs better than the current state-of-the-art methods that rely on CI tests.

6 Conclusion

In this paper, we have considered the problem of estimating intervention targets in linear structural equation models (SEMs) under soft interventions. We have proposed an algorithm that consistently identifies intervened nodes that can scale to larger graphs and have sample complexity guarantees in Gaussian linear SEMs. The algorithm can be used also to infer interventional Markov equivalence class (MEC) from the observational MEC. We have demonstrated comparable or better performance compared to the existing methods in a number of settings.

The limitation of our method is that it only applies to linear SEMs. The dataset in an application should be evaluated carefully to confirm whether the assumptions are satisfied. This avoids any adverse effects arising from wrong interpretations of cause-effect relationships. Extending the similar ideas for scalable and efficient intervention target estimation to the non-linear DAGs is an open question that we aim to address in future work. Finally, the condition number of the optimization problem is assumed to be bounded in the sample complexity results. We note that our algorithm is independent of the specific precision difference estimation (PDE) algorithms and can be used in a modular way. In this regard, it can benefit from any potential relaxation on this limitation of PDE algorithms.

Acknowledgments and Disclosure of Funding

This work was supported by the Rensselaer-IBM AI Research Collaboration (<http://airc.rpi.edu>), part of the IBM AI Horizons Network (<http://ibm.biz/AIHorizons>).

References

- [1] Y. Fu, J. A. Foden, C. Khayter, M. L. Maeder, D. Reyon, J. K. Joung, and J. D. Sander, “High-frequency off-target mutagenesis induced by CRISPR-Cas nucleases in human cells,” *Nature Biotechnology*, vol. 31, no. 9, pp. 822–826, 2013.
- [2] R. W. Tothill, A. V. Tinker, J. George, R. Brown, S. B. Fox, S. Lade, D. S. Johnson, M. K. Trivett, D. Etemadmoghadam, B. Locandro, N. Traficante, S. Fereday, J. A. Hung, Y.-E. Chiew,

- I. Haviv, Australian Ovarian Cancer Study Group, D. Gertig, A. deFazio, and D. D. Bowtell, “Novel molecular subtypes of serous and endometrioid ovarian cancer linked to clinical outcome,” *Clinical Cancer Research*, vol. 14, no. 16, pp. 5198–5208, 2008.
- [3] S. Saney and J. A. Chambers, *EEG Signal Processing*. John Wiley & Sons, 2013.
- [4] J. Bogatinovski, S. Nedelkoski, A. Acker, F. Schmidt, T. Wittkopp, S. Becker, J. Cardoso, and O. Kao, “Artificial Intelligence for IT Operations workshop white paper,” *arXiv:2101.06054*, 2021.
- [5] F. Eberhardt and R. Scheines, “Interventions and causal inference,” *Philosophy of Science*, vol. 74, no. 5, pp. 981–995, 2007.
- [6] L. Bottou, J. Peters, J. Quiñero-Candela, D. X. Charles, D. M. Chickering, E. Portugaly, D. Ray, P. Simard, and E. Snelson, “Counterfactual reasoning and learning systems: The example of computational advertising,” *Journal of Machine Learning Research*, vol. 14, no. 65, pp. 3207–3260, 2013.
- [7] A. Jaber, M. Kocaoglu, K. Shanmugam, and E. Bareinboim, “Causal discovery from soft interventions with unknown targets: Characterization and learning,” in *Proc. Advances in Neural Information Processing Systems*, Dec. 2020, pp. 9551–9561.
- [8] D. Eaton and K. Murphy, “Exact Bayesian structure learning from uncertain interventions,” in *Proc. International Conference on Artificial Intelligence and Statistics*, San Juan, Puerto Rico, Mar. 2007, pp. 107–114.
- [9] Y. Wang, C. Squires, A. Belyaeva, and C. Uhler, “Direct estimation of differences in causal graphs,” in *Proc. Advances in Neural Information Processing Systems*, Montreal, Canada, Dec. 2018, pp. 3770–3781.
- [10] S. D. Zhao, T. T. Cai, and H. Li, “Direct estimation of differential networks,” *Biometrika*, vol. 101, no. 2, pp. 253–268, 2014.
- [11] H. Yuan, R. Xi, C. Chen, and M. Deng, “Differential network analysis via lasso penalized d -trace loss,” *Biometrika*, vol. 104, no. 4, pp. 755–770, 2017.
- [12] B. Jiang, X. Wang, and C. Leng, “A direct approach for sparse quadratic discriminant analysis,” *The Journal of Machine Learning Research*, vol. 19, no. 1, pp. 1098–1134, 2018.
- [13] Z. Tang, Z. Yu, and C. Wang, “A fast iterative algorithm for high-dimensional differential network,” *Computational Statistics*, vol. 35, no. 1, pp. 95–109, 2020.
- [14] A. Ghoshal and J. Honorio, “Direct estimation of difference between structural equation models in high dimensions,” *arXiv:1906.12024*, 2019.
- [15] A. Hauser and P. Bühlmann, “Characterization and greedy learning of interventional Markov equivalence classes of directed acyclic graphs,” *Journal of Machine Learning Research*, vol. 13, no. 1, pp. 2409–2464, 2012.
- [16] Y. Wang, L. Solus, K. Yang, and C. Uhler, “Permutation-based causal inference algorithms with interventions,” in *Proc. Advances in Neural Information Processing Systems*, Long Beach, CA, Dec. 2017, pp. 5822–5831.
- [17] N. R. Ke, O. Bilaniuk, A. Goyal, S. Bauer, H. Larochelle, B. Schölkopf, M. C. Mozer, C. Pal, and Y. Bengio, “Learning neural causal models from unknown interventions,” *arXiv:1910.01075*, 2019.
- [18] P. Brouillard, S. Lachapelle, A. Lacoste, S. Lacoste-Julien, and A. Drouin, “Differentiable causal discovery from interventional data,” in *Proc. Advances in Neural Information Processing Systems*, Dec. 2020, pp. 21 865–21 877.
- [19] C. Squires, Y. Wang, and C. Uhler, “Permutation-based causal structure learning with unknown intervention targets,” in *Proc. Conference on Uncertainty in Artificial Intelligence*, Aug. 2020, pp. 1039–1048.
- [20] T. Verma and J. Pearl, “An algorithm for deciding if a set of observed independencies has a causal explanation,” in *Proc. Conference on Uncertainty in Artificial Intelligence*, Stanford, CA, Jul. 1992, pp. 323–330.
- [21] P. Erdős and A. Rényi, “On Random Graphs I,” *Publicationes Mathematicae Debrecen*, vol. 6, pp. 290–297, 1959.

- [22] K. Sachs, O. Perez, D. Pe'er, D. A. Lauffenburger, and G. P. Nolan, "Causal protein-signaling networks derived from multiparameter single-cell data," *Science*, vol. 308, no. 5721, pp. 523–529, 2005.
- [23] R. O. Ness, K. Sachs, P. Mallick, and O. Vitek, "A Bayesian active learning experimental design for inferring signaling networks," in *Proc. Research in Computational Molecular Biology*, Hong Kong, May 2017, pp. 134–156.
- [24] A. Dixit, O. Parnas, B. Li, J. Chen, C. P. Fulco, L. Jerby-Arnon, N. D. Marjanovic, D. Dionne, T. Burks, R. Raychowdhury, B. Adamson, T. M. Norman, E. S. Lander, J. S. Weissman, N. Friedman, and A. Regev, "Perturb-seq: Dissecting molecular circuits with scalable single-cell RNA profiling of pooled genetic screens," *Cell*, vol. 167, no. 7, pp. 1853–1866.e17, 2016.
- [25] A. Ghoshal and J. Honorio, "Learning linear structural equation models in polynomial time and sample complexity," in *Proc. International Conference on Artificial Intelligence and Statistics*, Playa Blanca, Lanzarote, Canary Islands, Apr. 2018, pp. 1466–1475.
- [26] M. Pourahmadi, "Covariance estimation: The glm and regularization perspectives," *Statistical Science*, vol. 26, no. 3, pp. 369–387, 2011.
- [27] P. Ravikumar, M. J. Wainwright, G. Raskutti, and B. Yu, "High-dimensional covariance estimation by minimizing ℓ_1 -penalized log-determinant divergence," *Electronic Journal of Statistics*, vol. 5, pp. 935–980, 2011.

"Scalable Intervention Target Estimation in Linear Models" Supplementary Material

A Theoretical analysis

Proof of Lemma 1. This lemma is proved in [14]. We provide an alternative and simplified proof for completeness, and with an approach that fits our subsequent analysis. The first observation is that the noise variance of a terminal node j is the inverse of the corresponding diagonal entry of the precision matrix obtained by removing all descendants of j . Thus, if j has no descendants in a set S , then $\sigma_{S,j}^{-2} = [\Theta_S]_{j,j}$. The second observation is that the noise variance of a node in a restricted SEM is affected by only its ancestors. Therefore, the noise variance of a node j in a restricted SEM over S is equal to that of over a set $S \cap \text{an}(j)$, i.e., $\sigma_{S,j} = \sigma_{S \cap \text{an}(j),j}$. Due to the second observation, we will consider only the restricted SEMs over the sets of the form $S = \text{an}(j) \setminus U_j$, where U_j denotes the ancestors of j that lie out of this restricted SEM. Let us denote the precision matrix of the restricted SEM over $\text{an}(j)$ by $\Phi \triangleq \Theta_{\text{an}(j)}$. We obtain the variance of the noise term for a node $j \in S$ as follows:

$$\Phi_S = \Phi_{S,S} - \Phi_{S,U_j}(\Phi_{U_j,U_j})^{-1}\Phi_{U_j,S}, \quad (12)$$

$$\frac{1}{\sigma_{S,j}^2} = \Phi_{j,j} = \frac{1}{\sigma_j^2} - \Phi_{j,U_j}(\Phi_{U_j,U_j})^{-1}\Phi_{U_j,j} \quad (13)$$

$$= \frac{1}{\sigma_j^2} - \frac{B_{U_j,j}^\top [\Theta_{\text{an}(j)}]_{U_j,U_j}^{-1} B_{U_j,j}}{\sigma_j^4}, \quad (14)$$

where the first line is due to Schur's complement; the second line is due to the observations mentioned above since j is a terminal node in both sets S and $\text{an}(j)$; and the last line is due to (2) and Proposition 4 of [25].

We use the Markov property to characterize edge weights in a restricted SEM. Conditioned on all of its parents, node j is independent of the remaining nodes. Hence, $[B_{\text{an}(j)}]_{k,j} = B_{k,j}$. We consider the same set $S = \text{an}(j) \setminus U_j$, $\Phi = \Theta_{\text{an}(j)}$ and derive the edge weights as follows:

$$[\Phi_S]_{j,k} = \Phi_{j,k} - \Phi_{j,U_j}(\Phi_{U_j,U_j})^{-1}\Phi_{U_j,k} \quad (15)$$

$$= -\frac{[B_{\text{an}(j)}]_{k,j}}{\sigma_{\text{an}(j),j}^2} + \frac{[B_{\text{an}(j)}]_{U_j,j}^\top}{\sigma_{\text{an}(j),j}^2} [\Theta_{\text{an}(j)}]_{U_j,U_j}^{-1} [\Theta_{\text{an}(j)}]_{U_j,k} \quad (16)$$

$$= -\frac{B_{k,j}}{\sigma_j^2} + \frac{B_{U_j,j}^\top}{\sigma_j^2} [\Theta_{\text{an}(j)}]_{U_j,U_j}^{-1} [\Theta_{\text{an}(j)}]_{U_j,k}, \quad (17)$$

$$B_{k,j}^S = \frac{\sigma_{S,j}^2}{\sigma_j^2} (B_{k,j} - B_{U_j,j}^\top [\Theta_{\text{an}(j)}]_{U_j,U_j}^{-1} [\Theta_{\text{an}(j)}]_{U_j,k}), \quad (18)$$

where the last line follows from $[\Phi_S]_{j,k} = -[B_S]_{k,j}/\sigma_{S,j}^2$. Note that this last equality is correct since S contains only the ancestors of j . Similarly, we can write $[\Theta_S]_{j,j} = 1/\sigma_{S,j}^2$ if S contains only the ancestors of j . ■

Proof of Proposition 1. Let us consider the restricted SEM over set S and let $U_j = \text{an}(j) \setminus S$ denote the ancestors of j that are not included in the restricted SEM. Note that the restricted SEM over $\text{an}(j)$ has edge weights $B_{\text{an}(j)} = [B]_{\text{an}(j),\text{an}(j)}$ and noise covariance $\Omega_{\text{an}(j)} = [\Omega]_{\text{an}(j),\text{an}(j)}$.

Therefore, for nodes $u, v \in U_j$, we can use Lemma 1 to obtain,

$$[\Theta_{\text{an}(j)}]_{u,v} = -\frac{[B_{\text{an}(j)}]_{u,v}}{\sigma_{\text{an}(j),v}^2} - \frac{[B_{\text{an}(j)}]_{v,u}}{\sigma_{\text{an}(j),u}^2} + \sum_{l \in \text{an}(j)} \frac{[B_{\text{an}(j)}]_{u,l}[B_{\text{an}(j)}]_{v,l}}{\sigma_{\text{an}(j),l}^2} \quad (19)$$

$$= -\frac{B_{u,v}}{\sigma_v^2} - \frac{B_{v,u}}{\sigma_u^2} + \sum_{l \in \text{an}(j)} \frac{B_{u,l}B_{v,l}}{\sigma_l^2}, \quad (20)$$

$$[\Theta_{\text{an}(j)}]_{u,u} = \frac{1}{\sigma_{\text{an}(j),u}^2} + \sum_{l \in \text{an}(j)} \frac{[B_{\text{an}(j)}]_{u,l}^2}{\sigma_{\text{an}(j),l}^2} \quad (21)$$

$$= \frac{1}{\sigma_u^2} + \sum_{l \in \text{an}(j)} \frac{B_{u,l}^2}{\sigma_l^2}. \quad (22)$$

Now we will prove the first statement. If S contains $\text{an}_{\mathcal{I}}(j)$ and their parents $\text{pa}(\text{an}_{\mathcal{I}}(j))$, we know that neither u, v nor their children belong to \mathcal{I} . Therefore, $[\Theta_{\text{an}(j)}]_{u,v}$ and $[\Theta_{\text{an}(j)}]_{u,u}$ are invariant due to (20) and (22), respectively. Subsequently, we have $[\Delta_{\Theta_{\text{an}(j)}}]_{U_j, U_j} = 0$. Furthermore, since $j \notin \mathcal{I}$, we have that $[\Delta_B]_{k,j} = 0$ for $k \in [p]$. Using the Lemma 1 again, we obtain

$$\sigma_{S,j}^2 = \sigma_j^2 \left(\sigma_j^4 - B_{U_j,j}^\top [\Theta_{\text{an}(j)}]_{U_j, U_j}^{-1} B_{U_j,j} \right)^{-1}, \quad (23)$$

where we note that σ_j , $B_{U_j,j}$, and $[\Theta_{\text{an}(j)}]_{U_j, U_j}$ are all invariant and, subsequently, $\sigma_{S,j}^{(1)} = \sigma_{S,j}^{(2)}$ is invariant. This proves the first statement regarding the invariance of the noise term for a non-intervened node under certain restricted SEMs.

For the last part, Assumption 1 ensures that $\sigma_{S,i}^{(1)} \neq \sigma_{S,i}^{(2)}$ for $i \in \mathcal{I}$, and we have $[\Delta_{\Theta_S}]_{i,i} \neq 0$. Similarly, Assumption 1 states that if $[B_S]_{j,i} \neq 0$ for either model, then $[\Delta_{\Theta_S}]_{j,i} \neq 0$. ■

Proof of Theorem 1. We will follow the steps of the Algorithm 1 to obtain the consistency results. We assume that the covariance estimates are perfect, i.e., they are equal to the population-level statistics. Therefore, we can compute Δ_{Θ_S} for any $S \subseteq [p]$ correctly. Instead of estimating \mathcal{I} directly, we, equivalently, aim to identify its complement \mathcal{I}^C .

Forming S_Δ . In Step 1, we first estimate Δ_Θ over $[p]$ to obtain the nodes that are affected by the interventions. Note that $\sigma_i^{(1)} \neq \sigma_i^{(2)}$ for intervened nodes $i \in \mathcal{I}$ and $[\Delta_B]_{k,j} = 0$ for non-intervened nodes $j \notin \mathcal{I}$ and $k \in [p]$. According to (2) and (3), $[\Delta_B]_{k,k} \neq 0$ if and only if either $k \in \mathcal{I}$ or there exists $k \rightarrow i$ for which $i \in \mathcal{I}$. In other words, by forming the set $S_\Delta = \{k : k \in [p], [\Delta_\Theta]_{k,k} \neq 0\} = \mathcal{I} \cup \bigcup_{i \in \mathcal{I}} \text{pa}(i)$, we can discard the nodes in $[p] \setminus S_\Delta$. The discarded nodes consist of the non-intervened nodes that do not have children in \mathcal{I} . Next, we will show computationally, some of the non-intervened nodes in S_Δ can be identified easier than the others.

Forming non-intervened source nodes J_0 . Note that if a node j has an intervened ancestor, the distribution of X_j changes and, subsequently, $\Sigma_{j,j}^{(1)} \neq \Sigma_{j,j}^{(2)}$. If a node i is intervened, the distribution of X_i changes too, and it results in $\Sigma_{i,i}^{(1)} \neq \Sigma_{i,i}^{(2)}$. Therefore, we are able to find non-intervened source nodes directly from $\Sigma^{(1)}$ and $\Sigma^{(2)}$. Since we have already narrowed down our focus to set S_Δ , we define *non-intervened source nodes* as

$$J_0 \triangleq \{j : j \in S_\Delta, j \notin \mathcal{I}, \text{an}_{\mathcal{I}}(j) = \emptyset\} \quad (24)$$

$$= \{j : j \in S_\Delta, \Sigma_{j,j}^{(1)} = \Sigma_{j,j}^{(2)}\}. \quad (25)$$

Sets S_Δ and J_0 are subsequently fed into the next steps of the algorithm.

Forming source ancestral sets J_0^k . In Proposition 1 we have shown that for any non-intervened node j , there exists sets S that makes $\sigma_{S,j}$ invariant, and the condition is closely related to ancestors of j that are affected by the intervention being included in S . On the other hand, such a restricted SEM does not exist for any intervened node. Therefore, we can identify all the non-intervened nodes in $S_\Delta \setminus J_0$ by finding a proper restricted SEM over a subset of S_Δ . Hence, finding the ancestors of non-intervened nodes is critical. Now consider pair $\{j, k\}$ such that $j \in J_0, k \in S_\Delta \setminus J_0$. $\Sigma_{j,j}$ is invariant and $\Sigma_{k,k}$ is changing. If j and k have a common ancestor, which can be j itself, then $\Sigma_{j,k}$ is

nonzero and $[\Delta_{\Theta_{\{j,k\}}}]_{j,k} \neq 0$. Otherwise, $\Sigma_{j,k} = 0$ and $[\Delta_{\Theta_{\{j,k\}}}]_{j,k} = 0$. Subsequently, we define the *source ancestral set* J_0^k for each node $k \in S_\Delta \setminus J_0$, that consists of the nodes in J_0 that have a common ancestor with k , i.e.,

$$\begin{aligned} J_0^k &\triangleq \{j : j \in J_0, [\Delta_{\Theta_{\{j,k\}}}]_{j,k} \neq 0\}, \quad \forall k \in S_\Delta \setminus J_0 \\ &= \{j : j \in J_0, \text{an}(j) \cap \text{an}(k) \neq \emptyset\}. \end{aligned} \quad (26)$$

Next, we will use these source ancestral sets to group the nodes that have similar ancestors together.

Forming equivalence classes from J_0 . We note that some of the nodes in $S_\Delta \setminus J_0$ will have identical source ancestral sets. Therefore, we can decompose the set $S_\Delta \setminus J_0$ into *equivalence classes* such that all the nodes in a class have the same source ancestral sets. We denote these equivalence classes by $\mathcal{A}_1, \dots, \mathcal{A}_L$, and the source ancestral set corresponding to the class \mathcal{A}_ℓ by $J_0^{\mathcal{A}_\ell}$ for $\ell \in [L]$. Formally,

$$S_\Delta \setminus J_0 = \bigcup_{\ell \in [L]} \mathcal{A}_\ell, \quad (27)$$

$$\mathcal{A}_{\ell_1} \cap \mathcal{A}_{\ell_2} = \emptyset, \quad \text{for } \ell_1 \neq \ell_2, \quad (28)$$

$$J_0^{\mathcal{A}_\ell} \triangleq J_0^{k_1} = J_0^{k_2}, \quad \forall k_1, k_2 \in \mathcal{A}_\ell, \quad \text{for } \ell \in [L]. \quad (29)$$

We note that we order these classes according to a topological order such that for $1 \leq \ell < \ell' \leq L$, $J_0^{\mathcal{A}_{\ell'}} \not\subseteq J_0^{\mathcal{A}_\ell}$. In other words, the class corresponding to the superset of any $J_0^{\mathcal{A}_\ell}$ should appear after \mathcal{A}_ℓ in the sequence $\mathcal{A}_1, \dots, \mathcal{A}_L$. This ordering is important since we do not need descendants of a non-intervened node in a restricted SEM to conclude its invariance. In the next step, we will show how searching for such restricted SEMs for non-intervened nodes is simplified with this decomposition to equivalence classes.

d-separation property for invariance. We establish the connection between d-separation in interventional graphs and the precision differences. Consider the augmented graph characterization of interventions presented in [7]. A new node, F , is introduced to the graph to represent the interventional distribution. There are edges from F to i for any intervened node $i \in \mathcal{I}$ in the augmented graph. As there is no edge between F and non-intervened node j , there exists a set S that d-separates F and j in the augmented graph. This implies that the probability distribution of the node j is invariant given $S \setminus \{j\}$, which in turn implies that both conditional mean and variance of the node j does not change. Subsequently, $\sigma_{S,j}$ is invariant for this set S . Applying the results of [26] and [9], $[\Theta_S]_{j,j} = \sigma_{S,j}^{-2}$ is also invariant. Therefore, the set S that d-separates F and non-intervened j results in $[\Delta_{\Theta_S}]_{j,j} = 0$.

Processing equivalence classes. We process equivalence classes $\mathcal{A}_1, \dots, \mathcal{A}_L$ individually, i.e., at stage ℓ , we consider the nodes in \mathcal{A}_ℓ . Let us define $\mathcal{M}_\ell = J_0 \cup \bigcup_{1 \leq b < \ell} \mathcal{A}_b$. We will prove that for a non-intervened node $j \in \mathcal{A}_\ell$, we can determine its invariance via $2^{|\mathcal{A}_\ell|}$ PDE. Due to our ordering of the equivalence classes, any ancestor of j in S_Δ will lie in either \mathcal{M}_ℓ or \mathcal{A}_ℓ . Consider the set $S = \mathcal{M}_\ell \cup \text{an}_{\mathcal{I}}(j) \cup \text{pa}(\text{an}_{\mathcal{I}}(j))$ which is also of the form $\mathcal{M}_\ell \cup A$ for some $A \subseteq \mathcal{A}_\ell$. Note that S does not contain any descendant of j .

We will use *d-separation property for invariance* to show that this set S yields $[\Delta_{\Theta_S}]_{j,j} = 0$. Specifically, we will show that there does not exist a d-connecting path between the augmented node F and j . Suppose the contrary and let $\pi : \langle F \rightarrow i \dots j \rangle$ be a d-connecting path where $i \in \mathcal{I}$. If j has a tail end on π , there is a collider node k on the path that is a descendant of j . Since S does not contain any descendant of j , neither node k nor its descendants are in S , and it blocks the path. Therefore, the path should be of the form $\langle F \rightarrow i \dots \rightarrow j \rangle$. If i is a collider and not in S , it means it is not an ancestor of S . Therefore, its descendants are also not in S , and i blocks the path. If i is a collider and in S , it is either in \mathcal{M}_ℓ or in $\text{an}_{\mathcal{I}}(j)$. In either case, the parent of i on the path is also in S and it blocks the path. If i is not a collider, the path will be $\langle F \rightarrow i \rightarrow \dots \rightarrow j \rangle$. If i is in S , it blocks the path. If i is not in S , it is not an ancestor of j . Then, there is a collider k on the path that is a descendant of i . Since i is not in S , none of its descendants are neither in S . Therefore, k blocks the path. We have ruled out all possible active paths and shown that there does not exist a d-connecting path between F and j for $S = \mathcal{M}_\ell \cup \text{an}_{\mathcal{I}}(j) \cup \text{pa}(\text{an}_{\mathcal{I}}(j))$. Subsequently, $[\Delta_{\Theta_S}]_{j,j} = 0$ due to d-separation for invariance property. As we have noted before, set S can be written as $S = \mathcal{M}_\ell \cup A$ for some $A \subseteq \mathcal{A}_\ell$, and we can check the existence of such A , i.e., whether j is non-intervened by

using PDE only $2^{|\mathcal{A}_\ell|}$ times. Formally, the *process equivalence class* returns

$$\mathcal{I}_\ell = \{i : i \in \mathcal{A}_\ell \cap \mathcal{I}\}, \quad \text{and} \quad J_\ell = \{j : j \in \mathcal{A}_\ell \cap \mathcal{I}^C\}. \quad (30)$$

This concludes the proof that Algorithm 1 consistently estimates \mathcal{I} set. \blacksquare

Remark 3 After forming $\mathcal{A}_1, \dots, \mathcal{A}_L$ classes with corresponding sets $J_0^{\mathcal{A}_1}, \dots, J_0^{\mathcal{A}_L}$, consider a pair $\mathcal{A}_\ell, \mathcal{A}_{\ell'}$ where $1 \leq \ell < \ell' \leq L$. Note that for any node pair (u, v) where $u \in \mathcal{A}_\ell$ and $v \in \mathcal{A}_{\ell'}$, u is not a descendant of v . Additionally, if $J_0^{\mathcal{A}_\ell} \not\subset J_0^{\mathcal{A}_{\ell'}}$, u is not an ancestor of v . Hence, while considering \mathcal{A}_ℓ step of Algorithm 1, taking $\mathcal{M}_\ell = J_0^{\mathcal{A}_\ell} \cup \bigcup_{b \in \mathcal{B}_\ell} \mathcal{A}_b$ where $\mathcal{B}_\ell \triangleq \{b : J_0^{\mathcal{A}_b} \subset J_0^{\mathcal{A}_\ell}, 1 \leq b < \ell\}$ is equivalent to taking $\mathcal{M}_\ell = J_0 \cup \bigcup_{1 \leq b < \ell} \mathcal{A}_b$. We use the former simplified approach to reduce the computational burden by having fewer nodes for subsequent Δ_Θ estimates.

Proof of Theorem 2. While processing a class \mathcal{A}_ℓ in Algorithm 1, we declare a node j non-intervened if there exist a set $A \subset \mathcal{A}_\ell$ such that $[\Delta_{\Theta_{\mathcal{M}_\ell \cup A}}]_{j,j} = 0$. Note that there may exist more than one such $\mathcal{M}_\ell \cup A$, in which case we denote the smallest one by \mathcal{N}_j .

Now, define $c_j \triangleq \ell$ for all $j \in \mathcal{A}_\ell$, where ℓ is the index of the equivalence class that contains node j . We have shown in Section 4.2 that finding $\{j \rightarrow i\}_{j \notin \mathcal{I}, i \in \mathcal{I}}$ is sufficient to update MEC into \mathcal{I} -MEC. Therefore, our goal for a non-intervened node is to find all of its intervened children. Consider $j \in \mathcal{J}_{c_j}$ and $i \in \mathcal{I}_{c_i}$ such that $c_j \leq c_i$. If $i \in \mathcal{N}_j$, it immediately implies that j is not a parent of i . Suppose that $i \notin \mathcal{N}_j$.

Consider $S = \mathcal{M}_{c_i} \cup \text{pa}(i) \cup \{i\}$ that is also of the form $\mathcal{M}_{c_i} \cup A$ for some $A \subseteq \mathcal{A}_{c_i}$. Therefore, we compute PDE for this S in c_i -th stage of *process equivalence class*. If $j \notin \text{pa}(i)$, all the paths $j \cdots \rightarrow i$ are blocked with a parent of i that is given in S . On the other hand, if the path ends with $\leftarrow i$, the path contains a collider node k that is a descendant of i . Since i is the youngest node in S , that collider k blocks the path. Therefore, $[\Theta_S]_{j,i} = 0$ and $[\Delta_{\Theta_S}]_{j,i} = 0$ if $j \notin \text{pa}(i)$. From Assumption 1, if $j \in \text{pa}(i)$, $[\Delta_{\Theta_S}]_{j,i} \neq 0$. Therefore, we identify all the non-intervened parents of intervened node i .

Orienting more edges. In addition to finding $\{j \rightarrow i\}_{j \notin \mathcal{I}, i \in \mathcal{I}}$, which is the main objective of Theorem 2, we can also recover the edges $\{k \rightarrow i\}_{\{k,i\} \in \mathcal{I}, c_k \neq c_i}$. Consider nodes $k \in \mathcal{I}_{c_k}$ and $i \in \mathcal{I}_{c_i}$ such that $c_k < c_i$. In other words, k and i are both intervened but they belong to different equivalence classes. Similar to the previous case, by considering set $S = \mathcal{M}_{c_i} \cup \text{pa}(i) \cup i$, we obtain $[\Delta_{\Theta_S}]_{k,i} \neq 0$ if $k \in \text{pa}(i)$ and $[\Delta_{\Theta_S}]_{k,i} = 0$ otherwise. Therefore, $k \notin \text{pa}(i)$, and we can orient all $k \rightarrow i$ edges if both nodes are intervened and belong to different equivalence classes.

Proof of Theorem 3. We use the ADMM-based approach of [12] as our PDE function to estimate $\Delta = \Theta^{(1)} - \Theta^{(2)}$. Theorem 1 of [12] gives the sample complexity of this estimation as $O(M_\Sigma M_{\Gamma, \Gamma^T} d^4 \log p)$. In Theorem 3, we further assume that the product $M_\Sigma M_{\Gamma, \Gamma^T}$ is bounded. Accordingly, with $n = O\left(\frac{d^4 \log p}{\varepsilon^2 \delta}\right)$ samples, PDE's output $\hat{\Delta}$ satisfies $\|\hat{\Delta} - \Delta\|_\infty < \varepsilon$ with a probability at least $1 - \delta$. We note that the conditions in Theorem 3 are given for the linear SEM over $[p]$ and the associated covariance matrices. If these conditions hold, they also hold for the restricted SEM over any $S \subset [p]$. Therefore, if we have $\|\hat{\Delta}_\Theta - \Delta_\Theta\|_\infty < \varepsilon$, we also have $\|\hat{\Delta}_{\Theta_S} - \Delta_{\Theta_S}\|_\infty < \varepsilon$ for any set S . Subsequently, we can threshold PDE outputs $\hat{\Delta}_{\Theta_S}$ by ε to exactly recover the support of Δ_{Θ_S} for any set S .

Note that Algorithm 1 requires only the support of Δ_{Θ_S} for a number of sets S . Therefore, with $n = O\left(\frac{d^4 \log p}{\varepsilon^2 \delta}\right)$ samples, Algorithm 1 identifies \mathcal{I} with a probability at least $1 - \delta$. We have shown in the proof of Theorem 2 that finding $\{j \rightarrow i\}_{j \notin \mathcal{I}, i \in \mathcal{I}}$ does not require any additional Δ_Θ estimates. Therefore, with $n = O\left(\frac{d^4 \log p}{\varepsilon^2 \delta}\right)$ samples, Algorithm 1 also identifies the non-intervened parents of the intervened nodes $\{j \rightarrow i\}_{j \notin \mathcal{I}, i \in \mathcal{I}}$ with a probability at least $1 - \delta$. \blacksquare

We finally note that Corollary 1 of [12] explicitly assumes that both M_Σ and M_{Γ, Γ^T} are bounded to remove $M_\Sigma M_{\Gamma, \Gamma^T}$ from the sample complexity. However, it can be readily relaxed to $M_\Sigma M_{\Gamma, \Gamma^T} < +\infty$ since both terms always appear within the same product. We note that this relaxation brings about a significant level of flexibility in choosing covariance matrices. Indeed, this

product is closely related to the condition number of the estimation problem. Two terms correspond to the norm of the inverse of Hessian of the optimization problem and the norm of the covariance, respectively. Product of these terms, the condition number, appears in similar matrix inference problems such as graphical lasso [27].

B Additional experiments

B.1 Intervention recovery

We have compared the results of our algorithm and those of UT-IGSP for estimating intervention targets under shift intervention model in Section 5.1. We expand the simulations to various settings in this subsection. Specifically, we report the results for shift intervention model with higher density $c = 2.5$ in Table 2, increased variance setting with $c = 2.5$ in Table 3, and randomized intervention setting with $c = 2.5$ in Table 4.

Our algorithm works well in all settings. Especially, increasing the dimension does not adversely affect accuracy and time complexity.

Table 2: \mathcal{I} estimation in the shift intervention model - 50 repetitions with 5000 samples - density 2.5

p	UT-IGSP ([19])				Algorithm 1			
	Precision	Recall	F1	Time(s)	Precision	Recall	F1	Time(s)
20	0.95	0.99	0.97	0.2	0.90	0.86	0.88	0.2
40	0.89	0.99	0.94	0.6	0.87	0.91	0.89	0.3
60	0.88	1	0.94	2.0	0.86	0.96	0.91	0.4
80	0.80	1	0.89	7.0	0.86	0.94	0.90	0.5
100	0.77	1	0.87	17.7	0.87	0.98	0.92	0.5

Table 3: \mathcal{I} estimation in the increased variance model - 50 repetitions with 5000 samples - density 2.5

p	UT-IGSP ([19])				Algorithm 1			
	Precision	Recall	F1	Time(s)	Precision	Recall	F1	Time(s)
20	0.90	0.99	0.95	0.2	0.89	0.86	0.87	0.2
40	0.85	1	0.92	0.6	0.87	0.93	0.90	0.3
60	0.88	1	0.93	2.4	0.89	0.97	0.92	0.3
80	0.80	1	0.89	5.8	0.86	0.97	0.91	0.4

Table 4: \mathcal{I} estimation in the randomized intervention - 50 repetitions with 5000 samples - density 2.5

p	UT-IGSP ([19])				Algorithm 1			
	Precision	Recall	F1	Time(s)	Precision	Recall	F1	Time(s)
20	0.92	1	0.96	0.2	0.86	0.91	0.88	0.2
40	0.82	1	0.90	0.7	0.88	0.94	0.91	0.3
60	0.81	1	0.90	2.8	0.84	0.96	0.90	0.5
80	0.74	1	0.85	8.4	0.86	0.92	0.89	0.6

Comparison with Ghoshal’s algorithm [14]. Ghoshal’s algorithm in [14] is designed to estimate Δ_B , and its performance critically hinges on the noise variances to be invariant. Even though it is not designed to return intervention targets, we can define the estimated intervention set of Ghoshal’s algorithm as $\hat{\mathcal{I}} \triangleq \{i : i, \exists j, (\Delta_B)_{j,i} \neq 0\}$. We run our algorithm and Ghoshal’s algorithm on the randomized intervention setting described in Section 5.1 and report the results in Table 5. Expectedly, Ghoshal’s algorithm does not perform well due to violation of the invariant noise variance assumption.

Table 5: \mathcal{I} estimation in the randomized intervention model - 100 repetitions with 10000 samples - density 2.5

p	Ghoshal [14]				Algorithm 1			
	Precision	Recall	F1	Time(s)	Precision	Recall	F1	Time(s)
20	0.74	0.62	0.67	<0.1	0.92	0.92	0.92	0.6
40	0.73	0.68	0.70	0.1	0.91	0.94	0.93	0.6
60	0.70	0.69	0.69	0.2	0.91	0.96	0.94	0.6
80	0.69	0.66	0.67	0.3	0.91	0.96	0.93	0.6
100	0.66	0.63	0.64	0.4	0.91	0.95	0.93	0.7

Increased number of samples. Theorem 1 states that our algorithm is consistent. Figure 1 shows that the performance of the algorithm increases significantly with the increased number of samples in all of the considered settings. We provide additional evidence of this fact. We generate 50 random graphs with density $c = 2.5$ for each of the shift intervention, increased variance, and randomized intervention settings. We report the F1 scores for each setting with 5000, 10000, and 20000 samples in Table 6.

Table 6: \mathcal{I} estimation with increased number of samples - 50 repetitions - density 2.5

p	Shift Intervention			Increased Variance			Randomized Intervention		
	5000	10000	20000	5000	10000	20000	5000	10000	20000
40	0.87	0.90	0.91	0.95	0.96	0.96	0.90	0.92	0.94
60	0.90	0.92	0.93	0.93	0.96	0.97	0.91	0.93	0.95
80	0.90	0.91	0.94	0.93	0.97	0.98	0.91	0.94	0.95
100	0.93	0.94	0.96	0.94	0.97	0.97	0.89	0.93	0.92

B.2 Causal structure learning

In Section 4.2, we have shown that our method recovers the new information that can be gained through interventions. Hence, Algorithm 1 refines the given MEC into the \mathcal{I} -MEC. Accordingly, we test our algorithm for the causal structure recovery task in this subsection.

First, we take the correct CPDAG of $\mathcal{G}^{(1)}$ and apply our algorithm’s findings to obtain \mathcal{I} -CPDAG. We run 100 realizations of Erdős-Rényi graphs with $c = 2$ and 10000 samples. For different values of graph size p , we consider fixed target set size $|\mathcal{I}| = 5$ or growing target set size $|\mathcal{I}| = p/10$. We report the results for recovery of \mathcal{I} -directed edges in Table 7.

Table 7: Recovery of \mathcal{I} -directed edges in the increased variance model

p	$ \mathcal{I} = 5$				$ \mathcal{I} = p/10$			
	Precision	Recall	F1	Time(s)	Precision	Recall	F1	Time(s)
40	0.69	0.93	0.80	0.15	0.73	0.94	0.82	0.11
60	0.73	0.93	0.82	0.24	0.73	0.93	0.82	0.25
80	0.75	0.93	0.83	0.28	0.73	0.96	0.83	0.45
100	0.82	0.97	0.89	0.42	0.72	0.93	0.81	0.86

Next, we consider recovering the non-intervened parents of the intervened nodes, i.e., $\{j \rightarrow i\}_{j \notin \mathcal{I}, i \in \mathcal{I}}$. We note that we do not use any given MEC information in this setting. Therefore, a comparison with UT-IGSP algorithm becomes feasible. We report the results for $|\mathcal{I}| = 5$ in Table 8. Similar to the intervention recovery task, our algorithm’s runtime does not suffer from increasing the dimension while the runtime of UT-IGSP grows very quickly.

Table 8: Recovery of non-intervened parents of intervened nodes

p	UT-IGSP ([19])				Algorithm 1			
	Precision	Recall	F1	Time(s)	Precision	Recall	F1	Time(s)
20	0.76	0.98	0.86	0.32	0.81	0.81	0.81	0.15
40	0.82	0.98	0.89	2.30	0.85	0.79	0.82	0.22
60	0.84	0.98	0.91	10.11	0.88	0.85	0.86	0.26
80	0.89	0.99	0.93	32.97	0.92	0.78	0.85	0.28

B.3 Application to real data

We have investigated directed edge recovery results for two real biological datasets in Section 5.3. In this subsection, we give the skeleton recovery results for the same datasets. Figure 3 illustrates that our observations from the directed edge recovery hold for the skeleton recovery as well. Comparison of figures 2 and 3 reveals that our algorithm orients fewer number of edges incorrectly with respect to UT-IGSP algorithm.

Hyperparameters. We have defined the regularization parameters λ_1, λ_2 , and λ_3 for our algorithm and cut-off value α for UT-IGSP in Section 5. Specifically, we have used $\lambda_1 \in [0.1, 0.3]$, $\lambda_2 = 0.2$, and $\lambda_3 \in [0.05, 0.2]$ for Algorithm 1, and $\alpha \in [0.0001, 0.5]$ for UT-IGSP while creating figures 2a and 3a. Similarly, we have used $\lambda_1 = 0.1$, $\lambda_2 = 0.05$, and $\lambda_3 \in [0.005, 0.1]$ for Algorithm 1, and $\alpha \in [0.005, 0.1]$ for UT-IGSP while creating figures 2b and 3b.

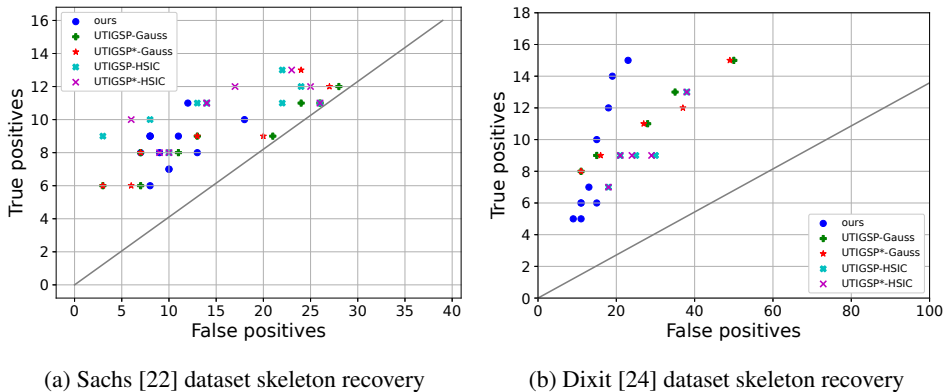


Figure 3: ROC curves for skeleton recovery. The solid grey line corresponds to random guessing.

B.4 Computational complexity

We have stated in Section 4 that the computational complexity of our algorithm is exponential in the size of the largest equivalence class, $\max |\mathcal{A}_\ell|$. This can be as large as p_Δ in some extreme examples. One possible scenario for this case is if the parents of intervention targets are also intervened. In this case, J_0 will be the empty set and all nodes in S_Δ will belong to the same group. However, this requires the interventions to concentrate in one neighborhood such that parents of the intervened nodes will also be intervened. In reality, such scenarios happen rarely, and interventions are generally distributed.

We generate 1000 instances of random graphs with $p = 100$, various densities, and target set sizes to demonstrate the much smaller size of \mathcal{A}_ℓ groups with respect to S_Δ . Figure 4 illustrates that $\max |\mathcal{A}_\ell|$ is much smaller than p_Δ . Indeed, Fig. 4 also shows the limitations of some of the related work that has computational complexity exponential in p_Δ strictly. For instance, for $p = 100$, $|\mathcal{I}| = 5$, and $c = 5$ in Fig. 4a, the 90%-th percentile of p_Δ is 25, whereas $\max_\ell |\mathcal{A}_\ell|$ is only 4. Gains of our algorithm become more dramatic when the target set is larger. For instance, for $p = 100$, $|\mathcal{I}| = 10$, and $c = 5$ in Fig. 4b, the 50%-th percentile of p_Δ is 34, whereas the 90%-th percentile of $\max_\ell |\mathcal{A}_\ell|$ is only 10. Therefore, our algorithm can scale up to higher dimensions.

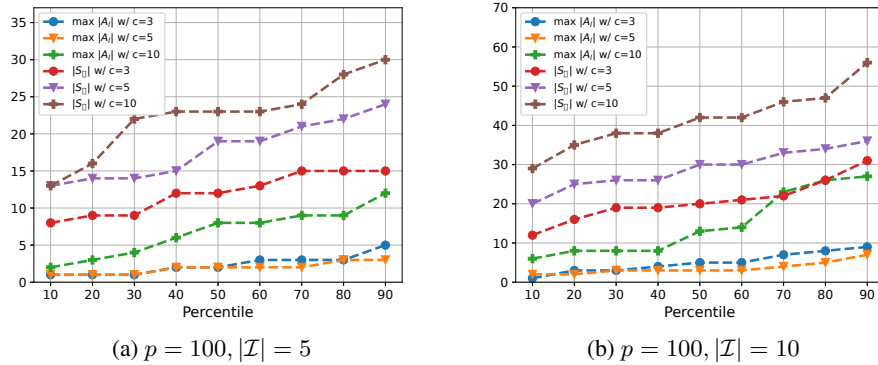


Figure 4: Exponential factor in the computational complexity of our algorithm, $\max_{\ell} |\mathcal{A}_{\ell}|$, is much smaller than the size of the affected nodes $p_{\Delta} = |S_{\Delta}|$. x-axis shows the percentile values over 1000 different random DAG instances. Largest class size $\max_{\ell} |\mathcal{A}_{\ell}|$ and p_{Δ} are plotted for three different density values.

We finally comment on the computational complexity of the PDE routine. The ADMM-based PDE algorithm of [12] has $O(p^3)$ complexity. We note that we run PDE with all $[p]$ nodes only once during the S_{Δ} estimation in Step 1. Hence, the estimation with $O(p^3)$ complexity will only be performed once. The rest of the PDE instances require much smaller number of nodes as stated in Remark 3. We note that a related study in [14] uses another PDE algorithm that has complexity $O(p^4)$. Reducing it to $O(p^3)$ is a significant gain, which allows us to process hundreds of nodes.

Checklist

1. For all authors...
 - (a) Do the main claims made in the abstract and introduction accurately reflect the paper’s contributions and scope? [\[Yes\]](#)
 - (b) Did you describe the limitations of your work? [\[Yes\]](#) See the last paragraph in Section 6.
 - (c) Did you discuss any potential negative societal impacts of your work? [\[Yes\]](#) See the last paragraph in Section 6.
 - (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [\[Yes\]](#)
2. If you are including theoretical results...
 - (a) Did you state the full set of assumptions of all theoretical results? [\[Yes\]](#) See Section 3 for exact problem definition and Assumption 1 in Section 4 for assumptions.
 - (b) Did you include complete proofs of all theoretical results? [\[Yes\]](#) See Appendix A.
3. If you ran experiments...
 - (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [\[Yes\]](#) See the link at the end of the abstract.
 - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [\[Yes\]](#) See Section 5 and the code in supplemental material.
 - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [\[Yes\]](#) Standard deviation of the precision and recall rates of intervention recovery are given in Table 1 in Section 5.1.
 - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [\[Yes\]](#) See first paragraph in Section 5.
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
 - (a) If your work uses existing assets, did you cite the creators? [\[Yes\]](#) See Section 5.

- (b) Did you mention the license of the assets? [Yes] See the footnote in Section 5.3.
 - (c) Did you include any new assets either in the supplemental material or as a URL? [Yes]
The code for our algorithm and simulations are provided in the supplemental material.
 - (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
 - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
- (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
 - (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
 - (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]