MESOSCOPIC CAUSAL DISCOVERY: LEARNING TARGET-CENTRIC CAUSAL SUBGRAPHS VIA DIFFERENTIABLE OPTIMIZATION

Anonymous authors
Paper under double-blind review

ABSTRACT

Causal discovery methods have traditionally focused on either recovering global causal structures or identifying local relationships around a target variable. However, many real-world applications require understanding causal mechanisms at an intermediate scale, where the goal is to uncover a structured subset of variables that are causally relevant to a target through both direct and multi-step pathways, without incurring the cost of full graph reconstruction. In this paper, we study Mesoscopic Causal Discovery in the form of target-centric k-hop causal subgraph recovery, bridging global and local causal learning. We formalize this setting as Target-Centric Causal Subgraph Discovery, which aims to recover a directed acyclic target-centric subgraph within a specified k-hop structural scope, capturing both immediate dependencies and the intermediate pathways through which causal influence propagates. This problem is challenging because extending beyond immediate neighborhoods rapidly enlarges the candidate space, while indirect causal pathways can be difficult to distinguish from spurious associations. To address these challenges, we propose Target-Conditioned Causal Subgraph Learning (TCSL), a differentiable framework that constructs a target-dependent candidate space from observational data and adaptively refines variable relevance through stochastic node gating while jointly learning causal structure and conditional distributions. TCSL combines conditional-flow modeling with a differentiable acyclicity constraint to enforce acyclicity in the learned structure and extracts the target-centric k-hop causal subgraph from the learned structure without access to ground-truth neighborhood information. Extensive experiments on both synthetic and real-world datasets demonstrate the effectiveness of TCSL.

1 INTRODUCTION

Causal discovery aims to uncover the underlying data-generating mechanisms from observational data and is central to modern machine learning and scientific analysis (Pearl, 2009; 2018). Existing methods can be broadly divided into global and local approaches. Global methods attempt to recover the full causal graph, capturing both direct and long-range dependencies, but are computationally expensive and prone to irrelevant dependencies in high-dimensional settings (Bello et al., 2022). Local methods focus on a target variable's immediate neighborhood (Wu et al., 2019), offering better scalability and interpretability (Wang et al., 2025), yet they miss multi-step causal pathways that often carry important structural information about how effects propagate through the system (Ling et al., 2025).

While direct causes can, in principle, characterize a target variable's conditional distribution, multi-step causal pathways often convey critical structural information about how effects propagate through the system (Kitson et al., 2023). Understanding these intermediate mechanisms—rather than only direct dependencies—is essential for interpretation, reasoning, and intervention. Consequently, local methods may yield an incomplete view of causal structure when pathway-level information is required (Ling et al., 2024). This limitation is particularly pronounced in real-world applications, where the goal is to identify a structured subset of variables that are causally relevant to a target through both direct and indirect connections. For instance, in scientific analysis and decision-making, practitioners often seek to trace multi-step influence chains, revealing intermediate mechanisms (Wang

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

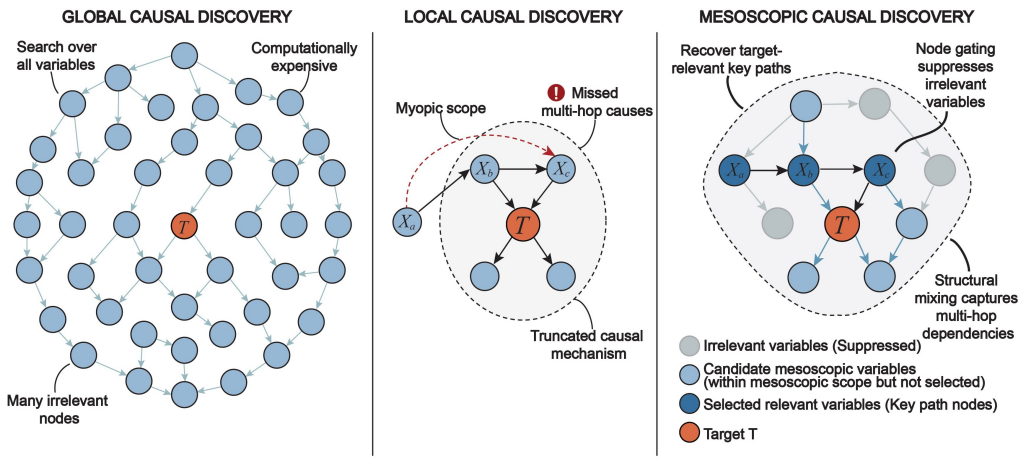

Figure 1: Illustration of global, local, and mesoscopic causal discovery. Compared with global discovery over the full graph and local discovery around direct neighbors, mesoscopic causal discovery recovers a target-centric multi-hop subgraph that captures intermediate causal pathways.

et al., 2023). Such requirements naturally call for a representation that is target-focused yet structurally expressive, as illustrated in Fig. 1, going beyond local neighborhoods while avoiding the complexity of full graph recovery.

In this work, we study Mesoscopic Causal Discovery in the form of target-centric k-hop causal subgraph recovery, bridging global and local causal discovery. Specifically, we formalize the task of Target-Centric Causal Subgraph Discovery, which aims to recover a directed acyclic subgraph centered around a target variable within a specified k-hop structural scope. This subgraph includes not only the direct causes and effects of the target, but also intermediate variables within the specified multi-hop structural scope. We provide a rigorous formulation of this problem in terms of k-hop neighborhoods and induced subgraphs, establishing a principled framework for studying causal structures at a mesoscopic scale.

Despite its practical importance, mesoscopic causal discovery presents several challenges. First, identifying relevant variables beyond the immediate neighborhood rapidly enlarges the candidate space, making naïve extensions of local methods computationally demanding (Aliferis et al., 2010). Second, indirect causality is difficult to distinguish from spurious correlations, particularly in high-dimensional settings where redundant dependencies are prevalent (Ng et al., 2024; Jin et al., 2026). Third, existing global methods, while capable of modeling long-range dependencies, operate over the full variable space rather than explicitly conditioning the learning scope on a target-specific query (Zheng et al., 2020). These challenges motivate a target-centric approach that supports multi-hop structural recovery without requiring full-graph reconstruction.

To tackle these challenges, we propose a unified framework termed Target-Conditioned Causal Subgraph Learning (TCSL). The key idea is to cast target-centric causal discovery as a differentiable optimization problem over a data-driven candidate space, jointly learning variable relevance, causal structure, and conditional distributions. Specifically, TCSL first constructs a target-dependent candidate set from observational data and then applies stochastic node gating to refine variable relevance within this space during training. Within the candidate space, we employ conditional-flow modeling to capture complex conditional dependencies and integrate it with a differentiable directed acyclic graph (DAG) constraint to enforce acyclicity in the learned structure. After thresholding, prediction-based k-hop extraction on the skeleton of the learned graph returns the final target-centric subgraph without access to ground-truth neighborhood information.

2 Related Work

Existing causal discovery methods can be broadly grouped into local and global approaches according to their inference scope. Local methods focus on the neighborhood of a target node to improve

scalability and interpretability. Early Markov Blanket methods, such as GSMB (Margaritis & Thrun, 1999) and IAMB (Tsamardinos et al., 2003), identify the parents, children, and spouses of a target through conditional independence tests and greedy search, but mainly recover directly related variables and provide limited edge orientation.

More recent local methods further distinguish direct causes from direct effects. PCD-by-PCD (Yin et al., 2008) recursively applies parent-child discovery with Meek rules; CMB (Gao & Ji, 2015) first recovers the target's Markov Blanket and then orients edges by analyzing changes in conditional independence relations; ELCS (Yang et al., 2021) combines local structure learning with scoring; and PSL (Ling et al., 2022a) leverages observed V-structures for orientation. While efficient, these methods primarily focus on immediate neighborhoods, Markov blankets, or other local causal structures around the target. Recent local-to-global methods further integrate locally learned structural information into larger-scale skeleton recovery (Wang et al., 2025), showing that structural information can extend beyond a single local neighborhood.

Global methods instead recover the full causal graph. Traditional constraint- or score-based approaches search over the entire structure space and become expensive as dimensionality grows. Recent continuous optimization methods, including NOTEARS (Zheng et al., 2018), GraN-DAG (Lachapelle et al., 2019), DAG-GNN (Yu et al., 2019), and GAE (Ng et al., 2019), parameterize weighted adjacency matrices with differentiable acyclicity constraints, enabling gradient-based structure learning for nonlinear dependencies. These methods optimize causal structure over the full variable space, whereas the setting considered here defines the desired inference scope through a target variable and a query hop budget.

Overall, existing methods span local, local-to-global, and global causal discovery. Accordingly, the distinction of the mesoscopic setting lies not in multi-hop structural information alone, but in the query-defined inference scope and output object: given observational data, a target variable, and a hop budget k, the goal is to recover the corresponding target-centric k-hop causal subgraph.

3 Notation and Task Formulation

A Structural Causal Model (SCM) is a mathematical framework for describing and analyzing causal relationships among a set of variables $\mathcal{V} = \{X_1, X_2, \ldots, X_d\}$. It consists of a causal graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ and a set of structural equations. The causal graph is a DAG, where $\mathcal{V}$ denotes the set of nodes corresponding to random variables and $\mathcal{E}$ denotes the set of directed edges. A directed edge $X_i \to X_j$ in $\mathcal{G}$ indicates that X_i has a direct causal effect on X_j. In this context, X_i is a parent of X_j, and X_j is a child of X_i. Let $\mathrm{Pa}(X_i)$ and $\mathrm{Ch}(X_i)$ denote the sets of parents and children of X_i in $\mathcal{G}$, respectively.

Under the Markov condition of Bayesian networks, the joint distribution $P(X)$ can be factorized as:

$$P(X) = \prod_{i=1}^{d} P(X_i \mid Pa(X_i)) \tag{1}$$

We assume causal sufficiency, mutually independent exogenous noises, and that the data-generating nonlinear location–scale SEM lies in an identifiable regime:

$$X_i = f_i(Pa(X_i), \epsilon_i), \quad i = 1, 2, \ldots, d, \tag{2}$$

Here, f_i is the causal mechanism. Under these assumptions, the recovered directions admit a causal interpretation; if the assumptions are violated, the output should instead be viewed as a model-dependent directed structure estimate.

Mesoscopic Causal Discovery. For a given target variable X_T and a query hop budget $k \in \mathbb{N}_+$, we define its reference k-hop neighborhood in the ground-truth causal graph $\mathcal{G}$ as

$$\mathcal{N}_k^{\mathcal{G}}(X_T) = \left\{ X_i \in \mathcal{V} : d_{\mathrm{skel}(\mathcal{G})}(X_i, X_T) \le k \right\}, \tag{3}$$

where $d_{\mathrm{skel}(\mathcal{G})}(\cdot, \cdot)$ denotes the shortest-path distance on the undirected skeleton of $\mathcal{G}$. The corresponding reference target-centric subgraph is the induced subgraph

$$\mathcal{G}_T^{(k)} = \mathcal{G}\left[\mathcal{N}_k^{\mathcal{G}}(X_T)\right]. \tag{4}$$

Given observational data $\mathbf{X}$, a target variable X_T, and a query hop budget k, the goal is to predict a target-centric directed subgraph $\widehat{\mathcal{G}}_T^{(k)}$. TCSL constructs a data-driven candidate set $\mathcal{S}_T^{(k)}$ as the optimization search space, which need not coincide with the reference neighborhood $\mathcal{N}_k^{\mathcal{G}}(X_T)$: it may include additional variables or omit reference-neighborhood variables. Neither the membership nor the size of the reference neighborhood is supplied to candidate construction, training, or prediction. Thus, the task is formulated as

$$(\mathbf{X}, X_T, k) \longmapsto \widehat{\mathcal{G}}_T^{(k)},$$

with $\mathcal{G}_T^{(k)}$ used only as the evaluation reference.

4 Proposed Methodology

We propose TCSL for target-centric causal subgraph recovery. Given $(\mathbf{X}, X_T, k)$, it constructs a data-driven candidate set, refines variable relevance through stochastic node gating, and jointly learns conditional distributions and a DAG-constrained weighted adjacency within the candidate space. After optimization, thresholding A yields $\widehat{A}$, from which prediction-based k-hop extraction returns $\widehat{\mathcal{G}}_T^{(k)}$.

4.1 Candidate Construction and Node Gating

To restrict the high-dimensional search space, we compute the absolute Spearman association $r_{ij} = |\rho_{\text{Spearman}}(X_i, X_j)|$. For each node, we retain the top-b associated variables satisfying $r_{ij} \geq \gamma_s$, symmetrize the screening relations, and perform at most k rounds of bounded expansion from X_T under a candidate-size cap m:

$$\begin{aligned} \mathcal{S}_T^{(k)} &= \text{CandidateSelection}(\mathbf{X}, X_T, k; b, \gamma_s, m), \\ d' &= |\mathcal{S}_T^{(k)}|. \end{aligned}$$

The resulting set contains X_T and is constructed solely from the observational data, target, and query budget, without ground-truth neighborhood membership or size. Candidate construction defines the optimization space, while stochastic gating subsequently refines variable relevance within it.

Within $\mathcal{S}_T^{(k)}$, each non-target variable X_i is assigned a learnable gating logit η_i. During training, we sample $u_i \sim \text{Uniform}(0, 1)$ and instantiate the effective gate stochastically (Yamada et al., 2020; Louizos et al., 2017):

$$g_i = \sigma\left(\frac{\eta_i + \text{logit}(u_i)}{\tau_g}\right), \quad \text{logit}(u_i) = \ln u_i - \ln(1 - u_i), \tag{5}$$

Here, $\tau_g > 0$ controls relaxation smoothness. At inference, the perturbation is removed and $g_i = \sigma(\eta_i/\tau_g)$ for $i \neq T$; the target gate is fixed at $g_T = 1$ during both training and inference.

Let $W \in \mathbb{R}^{d' \times d'}$ be the learnable weighted adjacency, where W_{ij} represents the edge $X_i \to X_j$. The mask $M = \mathbf{1}_{d' \times d'} - I_{d'}$ removes only self-loops, giving $W^{(m)} = W \odot M$; all cross-variable directed edges remain learnable, and M performs neither candidate selection nor ground-truth encoding. Combining the masked adjacency with the node gates gives

$$A = \text{diag}(g) W^{(m)} \text{diag}(g), \quad \text{i.e., } A_{ij} = g_i W_{ij}^{(m)} g_j. \tag{6}$$

Because $A_{ij} = g_i W_{ij}^{(m)} g_j$, an edge is attenuated whenever either endpoint has a small gate. The same effective adjacency A is used for structural mixing, sparsity regularization, and differentiable acyclicity enforcement.

4.2 Target-Centric Causal Subgraph Learning

We impose the NOTEARS acyclicity characterization on the candidate-space effective adjacency:

$$h(A) = \text{tr}\left(\exp(A \odot A)\right) - d' = 0, \tag{7}$$

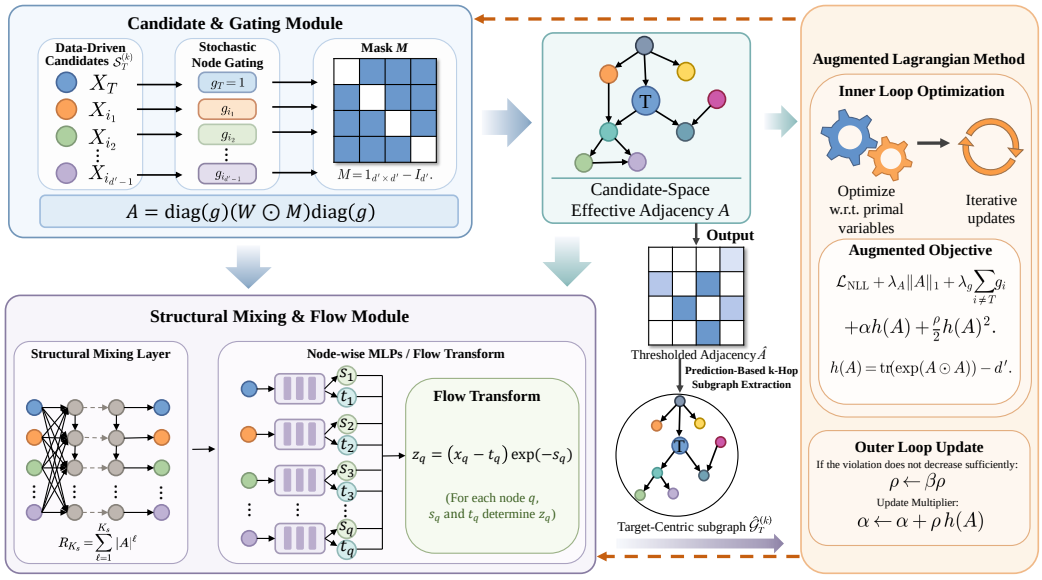

Figure 2: Flowchart of the TCSL framework.

Here, $\odot$ is the Hadamard product and $\exp(\cdot)$ is the matrix exponential. The smooth equality $h(A) = 0$ exactly characterizes acyclicity and permits gradient-based optimization.

Within the candidate subspace, we model each structural mechanism with a conditional affine flow:

$$X_j = \mu_{\psi_j}(Pa(X_j)) + \sigma_{\psi_j}(Pa(X_j)) \cdot Z_j, \tag{8}$$

Here, $Z_j \sim \mathcal{N}(0,1)$, while μ_{ψ_j} and σ_{ψ_j} model the conditional location and scale, providing a likelihood signal for nonlinear dependencies and edge directions.

Following Section 4.1, we perform structure learning on the effective adjacency

$$A = \text{diag}(g)(W \odot M)\text{diag}(g), \tag{9}$$

We combine this flow model with sparsity and DAG constraints on A. A Linear Aggregation–Nonlinear Transformation architecture first mixes candidate-parent information and then uses node-specific networks to produce the flow parameters.

The Structural Mixing Layer aggregates information through the effective adjacency:

$$R_{K_s} = \sum_{\ell=1}^{K_s} |A|^\ell, \tag{10}$$

Here, K_s is the internal mixing order, distinct from the query budget k used for candidate expansion and final extraction, and $|A|$ denotes element-wise absolute value. The resulting features are

$$\widetilde{\mathbf{X}} = \mathbf{X} R_{K_s}, \quad \text{i.e., } \tilde{x}_j = \sum_{i=1}^{d'} (R_{K_s})_{ij} X_i. \tag{11}$$

A node-specific MLP maps each mixed feature $\widetilde{x}_j$ to $[s_j, t_j] = \text{MLP}_j(\widetilde{x}_j; \psi_j)$, where ψ_j denotes all parameters of the node-specific network and $\Psi = \{\psi_j\}_{j=1}^{d'}$. Thus, $t_j = \mu_{\psi_j}(\widetilde{x}_j)$ and $\exp(s_j) = \sigma_{\psi_j}(\widetilde{x}_j)$. Following normalizing flows (Javaloy et al., 2023), the affine flow computes

$$z_j = (x_j - t_j)\exp(-s_j). \tag{12}$$

The exponential parameterization ensures a positive scale and an invertible transformation (Dinh et al., 2016).

At an acyclic feasible solution, a topological ordering makes the Jacobian triangular because t_j and s_j depend only on predecessors of X_j, giving $\log|\det(\partial Z/\partial X)| = -\sum_{j=1}^{d'} s_j$. With independent standard-normal latent variables and omitted constants, the negative log-likelihood on the feasible set is

$$\mathcal{L}_{NLL}(W,\eta,\Psi) = \frac{1}{n}\sum_{i=1}^{n}\sum_{j=1}^{d'}\left[\frac{1}{2}z_{ij}^2 + s_{ij}\right], \tag{13}$$

During ALM optimization, this expression is used as an off-feasible surrogate objective, while the augmented-Lagrangian terms involving $h(A)$ penalize acyclicity violations. We then add edge sparsity, non-target gate regularization, and the acyclicity constraint:

$$\begin{aligned} \min_{W,\eta,\Psi} \quad & \mathcal{F} = \mathcal{L}_{\mathrm{NLL}}(W,\eta,\Psi) + \lambda_A\|A\|_1 + \lambda_g\sum_{i\neq T} g_i \\ \text{s.t.} \quad & h(A) = 0, \end{aligned} \tag{14}$$

where λ_A controls edge sparsity and λ_g regularizes non-target gates; the fixed target gate $g_T = 1$ is excluded from the gate penalty.

We solve the constrained problem using the augmented Lagrangian method:

$$\mathcal{L}_{\mathrm{total}} = \mathcal{F} + \alpha h(A) + \frac{\rho}{2}[h(A)]^2, \tag{15}$$

where α and ρ are the Lagrange multiplier and penalty coefficient. With α and ρ fixed, Adam (Kingma & Ba, 2015) updates W, η, Ψ; the outer loop, indexed by r, then updates ρ and α based on the acyclicity violation:

$$W_{r+1}, \eta_{r+1}, \Psi_{r+1} \leftarrow \arg\min_{W,\eta,\Psi} \mathcal{L}_{\mathrm{total}}(W,\eta,\Psi,\alpha_r,\rho_r), \tag{16}$$

$$\rho_{r+1} = \begin{cases} \beta\rho_r, & \text{if } h(A_{r+1}) > \gamma h(A_r), \\ \rho_r, & \text{otherwise}, \end{cases} \tag{17}$$

$$\alpha_{r+1} = \alpha_r + \rho_{r+1}h(A_{r+1}). \tag{18}$$

Equation (16) is solved approximately with automatic differentiation and Adam. Equations (17)–(18) update the penalty coefficient and Lagrange multiplier to progressively enforce acyclicity.

After the fixed optimization budget, we form A using deterministic gates and threshold it to obtain the binary DAG adjacency $\widehat{A}$ (Ng et al., 2022):

$$\widehat{A}_{ij} = \mathbb{I}(|A_{ij}| > \tau), \tag{19}$$

where $\widehat{A}_{ij} = 1$ denotes an inferred edge $X_i \to X_j$.

Prediction-Based k-Hop Subgraph Extraction. Let $\widehat{\mathcal{G}}$ be the predicted directed graph represented by $\widehat{A}$. We define

$$\widehat{\mathcal{N}}_k(X_T) = \left\{X_i :\ d_{\mathrm{skel}(\widehat{\mathcal{G}})}(X_i, X_T) \le k\right\},$$
$$\widehat{\mathcal{G}}_T^{(k)} = \widehat{\mathcal{G}}\big[\widehat{\mathcal{N}}_k(X_T)\big].$$

The output is the directed induced subgraph on the predicted k-hop node set. Extraction uses only $\widehat{\mathcal{G}}$, X_T, and the supplied k, without ground-truth neighborhood membership or oracle projection. Breadth-first search simply computes unweighted shortest-path distances on the predicted skeleton.

4.3 Algorithm Analysis

Algorithm 1 summarizes TCSL. After constructing $\mathcal{S}_T^{(k)}$, the inner loop jointly updates W, η, Ψ, while scheduled ALM updates adjust ρ and α to enforce acyclicity. At inference, deterministic gates define A, which is thresholded to obtain $\widehat{A}$; prediction-based k-hop extraction then returns $\widehat{\mathcal{G}}_T^{(k)}$.

Algorithm 1: TCSL

Require: $\mathbf{X}, X_T, k$; candidate hyperparameters (b, γ_s, m); model and optimization hyperparameters, including budget T_{opt}

Ensure: $\widehat{\mathcal{G}}_T^{(k)}$: predicted target-centric directed subgraph

$\mathcal{S}_T^{(k)} \leftarrow \text{CandidateSelection}(\mathbf{X}, X_T, k; b, \gamma_s, m)$

$\mathbf{X} \leftarrow \mathbf{X}_{:,\mathcal{S}_T^{(k)}}$, $d' \leftarrow |\mathcal{S}_T^{(k)}|$

Initialize W, η, Ψ, α, and ρ; set $M \leftarrow \mathbf{1}_{d' \times d'} - I_{d'}$

for $t = 1, \ldots, T_{\text{opt}}$ **do**

  $g \leftarrow \text{ComputeGates}(\eta, \tau_g)$; set $g_T \leftarrow 1$

  $A \leftarrow \text{diag}(g)(W \odot M)\text{diag}(g)$

  Take an Adam step on the augmented-Lagrangian objective in Eq. (15)

  if *an ALM update is scheduled* **then**

    Recompute A and update ρ, α using Eqs. (17)–(18)

  end

end

$g_i \leftarrow \sigma(\eta_i/\tau_g)$ for $i \neq T$; set $g_T \leftarrow 1$

$A \leftarrow \text{diag}(g)(W \odot M)\text{diag}(g)$

$\widehat{A}_{ij} \leftarrow \mathbb{I}(|A_{ij}| > \tau)$; form the predicted graph $\widehat{\mathcal{G}}$

Extract $\widehat{\mathcal{N}}_k(X_T)$ using unweighted distances on $\text{skel}(\widehat{\mathcal{G}})$

return $\widehat{\mathcal{G}}_T^{(k)} = \widehat{\mathcal{G}}[\widehat{\mathcal{N}}_k(X_T)]$

5 EXPERIMENTS

We evaluate TCSL on synthetic and real-world datasets. For each synthetic dataset, the five nodes with the largest parent–child sets in the ground-truth DAG are selected as targets, and their induced k-hop subgraphs for $k \in \{1, 2, 3\}$ serve as references. Distance, Precision, and Recall are reported as mean $\pm$ standard deviation across these targets; the standard deviation reflects target-wise rather than graph-wise variation.

We compare TCSL with the local methods ELCS, PCD-by-PCD, MB-by-MB Wang et al. (2014), and PSL, and the global methods NOTEARS-MLP, GraN-DAG, and DAG-GNN. TCSL and the local methods are run separately for each target, whereas each global method is trained once and its target neighborhoods are extracted for evaluation (Ling et al., 2022b).

Precision and Recall evaluate directed edges, and $\text{Distance} = \sqrt{(1 - \text{Precision})^2 + (1 - \text{Recall})^2}$, with lower values being better. Runtime is reported in seconds and averaged across target runs for TCSL and local baselines; for global baselines, it denotes a single full-graph run. If a run reaches the 24-hour limit, its structural metrics and runtime are marked as "–".

5.1 SYNTHETIC DATASETS

Synthetic DAGs follow the Erdős–Rényi model (Erdős & Rényi, 1959) with $d \in \{50, 100, 200, 300, 500\}$ nodes and $2d$ edges. For each DAG, we generate $n \in \{1000, 5000\}$ samples under nonlinear ANM-GP and Additive-GP mechanisms. The former models causal relations with Gaussian processes and additive noise, while the latter sums independent Gaussian-process functions for each input (Hoyer et al., 2008); the noise terms satisfy $N_i \sim \mathcal{N}(0, 1)$ (Xie et al., 2022).

Table 1 shows that TCSL achieves the lowest Distance on the selected ANM-GP graph sizes while maintaining competitive Precision and Recall. Figure 3 shows a similar Distance pattern on Additive-GP datasets. Local methods are computationally efficient but exhibit less stable recovery as the target scope and graph size increase, whereas the global neural baselines become substantially more expensive and may fail on the largest graphs.

Table 1: Representative results on nonlinear ANM-GP datasets with $n = 5000$ and query hop budget $k = 3$ across selected graph sizes. Results across all evaluated graph sizes are reported in Appendix Table 3. Bold indicates the best values. Distance, Precision, and Recall are reported as mean ± std across the five selected target variables. Time is reported in seconds. Dist/Prec/Rec/Time denote Distance/Precision/Recall/Runtime.

Nodes	Algorithm	Dist↓	Prec↑	Rec↑	Time (s)↓
100	Notears-MLP	0.80±0.09	0.74±0.12	0.26±0.13	1032.63
	DAG-GNN	1.06±0.05	0.45±0.05	0.10±0.06	1592.47
	GraN-DAG	1.05±0.20	**0.80**±0.45	0.03±0.03	4592.33
	ELCS	0.82±0.11	0.49±0.03	0.37±0.12	1.80
	PCD-by-PCD	0.90±0.10	0.40±0.04	0.33±0.12	**0.72**
	MB-by-MB	1.01±0.05	0.16±0.01	**0.45**±0.09	4.98
	PSL	0.79±0.08	0.44±0.02	0.44±0.09	0.95
	TCSL	**0.66**±0.04	0.67±0.03	0.43±0.07	49.26
300	Notears-MLP	0.89±0.06	**0.79**±0.06	0.14±0.06	13124.16
	DAG-GNN	1.20±0.15	0.32±0.23	0.03±0.02	4892.13
	GraN-DAG	1.24±0.28	0.40±0.55	0.01±0.01	49323.28
	ELCS	0.94±0.04	0.38±0.04	0.29±0.05	9.48
	PCD-by-PCD	0.93±0.04	0.42±0.04	0.28±0.07	**3.07**
	MB-by-MB	1.09±0.06	0.12±0.04	0.36±0.05	76.86
	PSL	0.88±0.05	0.40±0.06	**0.37**±0.03	6.85
	TCSL	**0.72**±0.12	0.66±0.03	**0.37**±0.14	292.33
500	Notears-MLP	1.01±0.03	**0.61**±0.04	0.08±0.02	63334.11
	DAG-GNN	1.29±0.12	0.19±0.19	0.01±0.01	13022.77
	GraN-DAG	–	–	–	–
	ELCS	1.02±0.09	0.34±0.05	0.22±0.08	27.67
	PCD-by-PCD	0.98±0.02	0.40±0.03	0.22±0.01	**8.84**
	MB-by-MB	1.13±0.03	0.07±0.01	**0.36**±0.03	235.17
	PSL	0.95±0.02	0.37±0.04	0.29±0.05	17.88
	TCSL	**0.89**±0.05	0.59±0.01	0.21±0.06	724.65

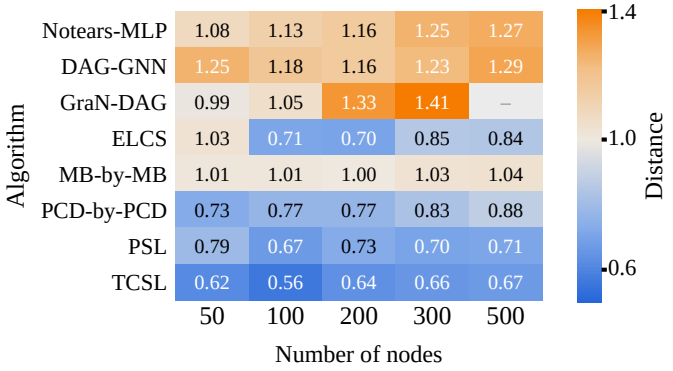

Figure 3: Distance on Additive-GP datasets with $n = 5000$ under query hop budget $k = 3$.

These results indicate a favorable overall structural-recovery trade-off rather than uniform dominance on individual metrics. By restricting optimization to a target-relevant candidate space, TCSL avoids the main bottleneck of global structure learning, although its conditional-flow architecture remains slower than simple discrete local methods. Complete results across sample sizes, graph sizes, query budgets, and mechanisms are reported in Appendix A (Tables 3–14).

5.2 Real-world Datasets

We further evaluate TCSL on the real-world biological Sachs dataset (Sachs et al., 2005), which records expression levels of proteins and phospholipids in human cells. Its consensus network contains 11 nodes and 17 edges, making it a standard benchmark for causal structure learning. Since the Sachs dataset is relatively small and deeper neighborhood structures are difficult to distinguish, we evaluate only $k = 1$, corresponding to the target's direct structural neighborhood, to ensure comparability with existing methods. Each of the 11 variables is used as a target, and the structural metrics are macro-averaged across all target queries. Table 2 reports the results. TCSL achieves the best Distance and Precision on the Sachs dataset, indicating a strong structural recovery trade-off with more accurate edge orientations. Although MB-by-MB obtains slightly higher Recall and PCD-by-PCD is faster, TCSL recovers target-centric causal subgraphs closer to the ground-truth network overall. These results suggest that TCSL remains effective on real-world biological data with complex conditional dependencies.

Table 2: Structural metrics on the Sachs benchmark, macro-averaged across all 11 target variables. Time is reported in seconds.

Algorithm	Dist	Prec	Rec	Time (s)↓
Notears-MLP	1.26	0.07	0.16	40.93
DAG-GNN	1.30	0.09	0.07	1599.30
GraN-DAG	1.35	0.06	0.04	187.25
ELCS	0.70	0.66	0.43	0.55
PCD-by-PCD	0.77	0.60	0.38	**0.16**
MB-by-MB	0.71	0.52	**0.50**	0.30
PSL	0.67	0.65	0.47	0.22
TCSL	**0.64**	**0.68**	0.46	11.94

5.3 Ablation Study

We conduct ablations at query hop budget $k = 1$, using the same data-driven candidate-construction configuration and the corresponding $k = 1$ reference subgraphs. Removing stochastic node gating substantially increases Distance, especially on larger graphs, with only minor runtime changes (Fig. 4). This supports its role in suppressing irrelevant nodes and stabilizing target-centric subgraph learning. Additional results on the structural mixing order K_s are provided in Appendix A.3.

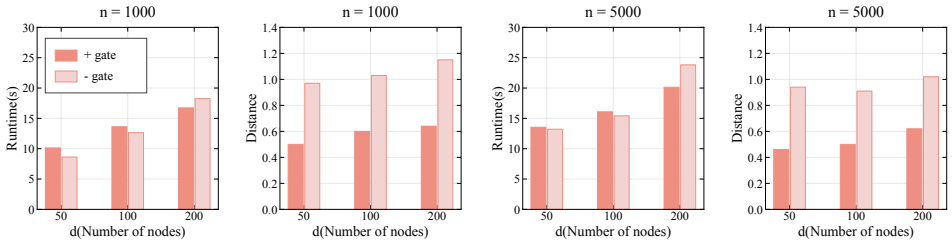

Figure 4: Runtime and Distance of TCSL with and without stochastic node gating. For each sample size $n = 1000$ and $n = 5000$, the Runtime panel is shown first, followed by the Distance panel.

6 Conclusion

We study mesoscopic causal discovery as target-centric causal subgraph recovery and propose TCSL, a continuous optimization framework for recovering direct causal relations and intermediate variables within multi-hop structural scopes. Experiments on synthetic and real-world datasets show favorable structural-recovery and scalability trade-offs. Future work will explore adaptive candidate selection and gating, as well as extensions to time-series and dynamic settings.

Limitations. TCSL assumes causal sufficiency and depends on candidate quality; latent confounding or omitted variables may degrade recovery. Reported synthetic variation is target-wise, not across independent graphs or seeds, and selecting five targets with the largest parent–child sets may limit representativeness. TCSL is costlier than simple local methods, and real-world evaluation is limited to Sachs at $k = 1$.

AI USE STATEMENT

Generative AI tools were used for language editing, manuscript polishing, and LaTeX formatting assistance. The research ideas, methodology, experiments, analyses, and conclusions were developed and verified by the authors. All AI-assisted changes were reviewed by the authors, who take full responsibility for the final content of this work.

REFERENCES

Constantin F Aliferis, Alexander Statnikov, Ioannis Tsamardinos, Subramani Mani, and Xenofon D Koutsoukos. Local causal and markov blanket induction for causal discovery and feature selection for classification part i: algorithms and empirical evaluation. *Journal of Machine Learning Research*, 11(1), 2010.

Kevin Bello, Bryon Aragam, and Pradeep Ravikumar. Dagma: Learning dags via m-matrices and a log-determinant acyclicity characterization. *Advances in Neural Information Processing Systems*, 35:8226–8239, 2022.

Laurent Dinh, Jascha Sohl-Dickstein, and Samy Bengio. Density estimation using real nvp. *arXiv preprint arXiv:1605.08803*, 2016.

Paul Erdős and Alfréd Rényi. On random graphs i. *Publicationes Mathematicae Debrecen*, 6: 290–297, 1959.

Tian Gao and Qiang Ji. Local causal discovery of direct causes and effects. *Advances in Neural Information Processing Systems*, 28, 2015.

Patrik Hoyer, Dominik Janzing, Joris M Mooij, Jonas Peters, and Bernhard Schölkopf. Nonlinear causal discovery with additive noise models. *Advances in neural information processing systems*, 21, 2008.

Adrián Javaloy, Pablo Sánchez-Martín, and Isabel Valera. Causal normalizing flows: from theory to practice. *Advances in Neural Information Processing Systems*, 36:58833–58864, 2023.

Kaifeng Jin, Ignavier Ng, Kun Zhang, and Biwei Huang. Revisiting differentiable structure learning: Inconsistency of l1 penalty and beyond. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 40, pp. 22399–22407, 2026.

Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representations*, 2015.

Neville Kenneth Kitson, Anthony C Constantinou, Zhigao Guo, Yang Liu, and Kiattikun Chobtham. A survey of bayesian network structure learning. *Artificial Intelligence Review*, 56(8):8721–8814, 2023.

Sébastien Lachapelle, Philippe Brouillard, Tristan Deleu, and Simon Lacoste-Julien. Gradient-based neural dag learning. *arXiv preprint arXiv:1906.02226*, 2019.

Zhaolong Ling, Kui Yu, Lin Liu, Jiuyong Li, Yiwen Zhang, and Xindong Wu. Psl: an algorithm for partial bayesian network structure learning. *ACM Transactions on Knowledge Discovery from Data (TKDD)*, 16(5):1–25, 2022a.

Zhaolong Ling, Kui Yu, Yiwen Zhang, Lin Liu, and Jiuyong Li. Causal learner: A toolbox for causal structure and markov blanket learning. *Pattern Recognition Letters*, 163:92–95, 2022b.

Zhaolong Ling, Honghui Peng, Yiwen Zhang, Debo Cheng, Xingyu Wu, Peng Zhou, and Kui Yu. Hybrid local causal discovery. *arXiv preprint arXiv:2412.19507*, 2024.

Zhaolong Ling, Jiale Yu, Yiwen Zhang, Debo Cheng, Peng Zhou, Xingyu Wu, Bingbing Jiang, and Kui Yu. Local causal discovery without causal sufficiency. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pp. 18737–18745, 2025.

Christos Louizos, Karen Ullrich, and Max Welling. Bayesian compression for deep learning. *Advances in neural information processing systems*, 30, 2017.

Dimitris Margaritis and Sebastian Thrun. Bayesian network induction via local neighborhoods. *Advances in neural information processing systems*, 12, 1999.

Ignavier Ng, Shengyu Zhu, Zhitang Chen, and Zhuangyan Fang. A graph autoencoder approach to causal structure learning. *arXiv preprint arXiv:1911.07420*, 2019.

Ignavier Ng, Shengyu Zhu, Zhuangyan Fang, Haoyang Li, Zhitang Chen, and Jun Wang. Masked gradient-based causal structure learning. In *Proceedings of the 2022 SIAM International Conference on Data Mining (SDM)*, pp. 424–432. SIAM, 2022.

Ignavier Ng, Biwei Huang, and Kun Zhang. Structure learning with continuous optimization: A sober look and beyond. In *Causal Learning and Reasoning*, pp. 71–105. PMLR, 2024.

Judea Pearl. *Causality: Models, Reasoning, and Inference*. Cambridge University Press, Cambridge, UK, 2nd edition, 2009.

Judea Pearl. Theoretical impediments to machine learning with seven sparks from the causal revolution. *arXiv preprint arXiv:1801.04016*, 2018.

Karen Sachs, Omar Perez, Dana Pe'er, Douglas A Lauffenburger, and Garry P Nolan. Causal protein-signaling networks derived from multiparameter single-cell data. *Science*, 308(5721):523–529, 2005.

Ioannis Tsamardinos, Constantin F Aliferis, Alexander R Statnikov, and Er Statnikov. Algorithms for large scale markov blanket discovery. In *FLAIRS*, volume 2, pp. 376–81, 2003.

Changzhang Wang, You Zhou, Qiang Zhao, and Zhi Geng. Discovering and orienting the edges connected to a target variable in a dag via a sequential local learning approach. *Computational statistics & data analysis*, 77:252–266, 2014.

Niantai Wang, Haoran Liu, Yunhe Zhao, Lin Li, Ye Gao, and Fan Zhang. An efficient local-to-global bayesian network skeleton learning algorithm. *Expert Systems with Applications*, 283:127593, 2025.

Xiaoxiao Wang, Minda Zhao, Fanyu Meng, Xin Liu, Zhaodan Kong, and Xin Chen. Quantifying causal path-specific importance in structural causal model. *Computation*, 11(7):133, 2023.

Xingyu Wu, Bingbing Jiang, Kui Yu, Huanhuan Chen, et al. Accurate markov boundary discovery for causal feature selection. *IEEE transactions on cybernetics*, 50(12):4983–4996, 2019.

Feng Xie, Biwei Huang, Zhengming Chen, Yangbo He, Zhi Geng, and Kun Zhang. Identification of linear non-gaussian latent hierarchical structure. In *International Conference on Machine Learning*, pp. 24370–24387. PMLR, 2022.

Yutaro Yamada, Ofir Lindenbaum, Sahand Negahban, and Yuval Kluger. Feature selection using stochastic gates. In *International conference on machine learning*, pp. 10648–10659. PMLR, 2020.

Shuai Yang, Hao Wang, Kui Yu, Fuyuan Cao, and Xindong Wu. Towards efficient local causal structure learning. *IEEE Transactions on Big Data*, 8(6):1592–1609, 2021.

Jianxin Yin, You Zhou, Changzhang Wang, Ping He, Cheng Zheng, and Zhi Geng. Partial orientation and local structural learning of causal networks for prediction. In *Causation and prediction challenge*, pp. 93–105. PMLR, 2008.

Yue Yu, Jie Chen, Tian Gao, and Mo Yu. Dag-gnn: Dag structure learning with graph neural networks. In *International conference on machine learning*, pp. 7154–7163. PMLR, 2019.

Xun Zheng, Bryon Aragam, Pradeep K Ravikumar, and Eric P Xing. Dags with no tears: Continuous optimization for structure learning. *Advances in neural information processing systems*, 31, 2018.

Xun Zheng, Chen Dan, Bryon Aragam, Pradeep Ravikumar, and Eric Xing. Learning sparse nonparametric dags. In *International conference on artificial intelligence and statistics*, pp. 3414–3425. Pmlr, 2020.

A Additional Methodological and Experimental Details

In this appendix, k denotes the query hop budget. It controls the depth of the bounded data-driven candidate expansion and indexes the corresponding target-centered reference and output scope. In contrast, K_s denotes the internal structural-mixing order in $R_{K_s} = \sum_{\ell=1}^{K_s} |A|^\ell$. The two quantities have different roles, although their numerical values may coincide. We evaluate all methods for $k \in \{1, 2, 3\}$, sample sizes $n \in \{1000, 5000\}$, graph sizes $d \in \{50, 100, 200, 300, 500\}$, and two nonlinear data-generating mechanisms: nonlinear ANM-GP and Additive-GP.

A.1 Evaluation and Aggregation Details

For each target query (X_T, k), let P_T and R_T denote the directed-edge Precision and Recall, respectively. Distance is defined as

$$D_T = \sqrt{(1 - P_T)^2 + (1 - R_T)^2}, \tag{20}$$

where a lower value indicates a better Precision–Recall trade-off.

For each synthetic dataset, let $\mathcal{T} = \{T_1, \ldots, T_5\}$ denote the five selected target variables. For any structural metric $m_T \in \{D_T, P_T, R_T\}$, we report

$$\overline{m} = \frac{1}{|\mathcal{T}|} \sum_{T \in \mathcal{T}} m_T, \qquad s_m = \sqrt{\frac{1}{|\mathcal{T}| - 1} \sum_{T \in \mathcal{T}} (m_T - \overline{m})^2}. \tag{21}$$

Thus, the reported standard deviation reflects variation across target queries rather than variation across independently generated graphs.

For the Sachs benchmark, every variable is used as a target, and the reported structural metrics are macro-averaged as

$$\overline{m}_{\text{Sachs}} = \frac{1}{11} \sum_{T \in \mathcal{V}} m_T. \tag{22}$$

For TCSL and the local baselines, given an evaluation target set $\mathcal{Q}$, runtime is averaged across the target-specific runs:

$$\bar{t}_{\text{query}} = \frac{1}{|\mathcal{Q}|} \sum_{T \in \mathcal{Q}} t_T. \tag{23}$$

For each global baseline, runtime denotes a single full-graph training run used to evaluate all target queries.

A.2 Path-Based Interpretation of Structural Mixing

Recall that the structural-mixing matrix is

$$R_{K_s} = \sum_{\ell=1}^{K_s} |A|^\ell, \tag{24}$$

where $|A|$ denotes the element-wise absolute value of A. Let $[d'] = \{1, \ldots, d'\}$ index the variables in the candidate space. For any $\ell \geq 1$, the (i, j)-th entry of $|A|^\ell$ can be expanded as

$$\left(|A|^\ell\right)_{ij} = \sum_{\substack{(v_0, \ldots, v_\ell) \in [d']^{\ell+1} \\ v_0 = i,\ v_\ell = j}} \prod_{r=0}^{\ell-1} |A_{v_r v_{r+1}}|. \tag{25}$$

Each nonzero product represents the weighted contribution of a directed walk of length ℓ from X_i to X_j. When A satisfies the DAG constraint, such walks cannot repeat a node and therefore correspond to directed paths.

Consequently,

$$(R_{K_s})_{ij} = \sum_{\ell=1}^{K_s} \left(|A|^\ell\right)_{ij} \tag{26}$$

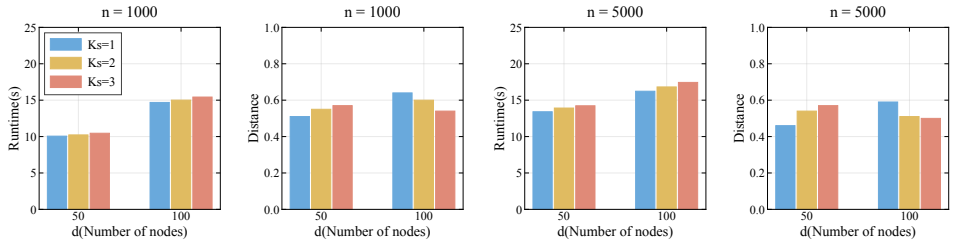

Figure 5: Runtime and Distance of TCSL with $K_s \in \{1, 2, 3\}$ for $n = 1000$ and $n = 5000$; for each n, the Runtime panel is shown first, followed by Distance.

aggregates the weighted contributions of directed walks of lengths up to K_s, which become directed paths under the DAG constraint. The mixed feature associated with node j is

$$\widetilde{\mathbf{X}}_{:j} = \sum_{i=1}^{d'} (R_{K_s})_{ij} \mathbf{X}_{:i}. \tag{27}$$

Thus, the structural-mixing layer combines information reaching node j through higher-order candidate-space connections.

The structural-mixing order K_s controls this internal aggregation and is distinct from the query hop budget k, which controls candidate expansion and final prediction-based subgraph extraction.

A.3 Effect of Structural Mixing Order

We conduct the structural mixing order ablation at query hop budget $k = 1$, using the same data-driven candidate-construction configuration and the corresponding $k = 1$ reference subgraphs. Increasing K_s has limited impact on small graphs but improves Distance on larger graphs (Fig. 5), suggesting that higher-order structural mixing provides useful multi-hop aggregation even when evaluation is restricted to first-order neighborhoods.

A.4 Extended Results on Nonlinear ANM with GP

Tables 4–6 summarize the results on nonlinear ANM-GP datasets with $n = 1000$ for query hop budgets $k \in \{1, 2, 3\}$. For $n = 5000$, Tables 7–8 report the results for $k \in \{1, 2\}$, while the complete $k = 3$ results are reported in Table 3.

Across these settings, local methods such as ELCS, PCD-by-PCD, and PSL show strong computational efficiency, with particularly low runtimes on small graphs. However, their structural recovery tends to degrade as the target-hop radius increases, especially for larger query radii such as $k = 3$ and for larger graphs with $d \geq 300$. In contrast, gradient-based global methods such as NOTEARS-MLP and GraN-DAG often achieve competitive structural accuracy, but their runtime increases substantially on larger graphs. TCSL generally achieves a favorable Distance–Recall trade-off across varying neighborhood bounds. This suggests that stochastic node gating combined with continuous DAG optimization can help isolate target-relevant subgraphs and reduce the influence of irrelevant variables as the candidate search space expands.

A.5 Extended Results on Additive Model with GP

To further evaluate the robustness of TCSL under different nonlinear causal mechanisms, we report results on the Additive-GP model. Tables 9–11 report the Additive-GP results with $n = 1000$, and Tables 12–14 report the corresponding results with $n = 5000$.

Table 3: Results on nonlinear ANM-GP datasets for query hop budget $k = 3$ with $n = 5000$ samples across all evaluated graph sizes. This table extends the $k = 3$ results in Table 1. Bold indicates the best values. Distance, Precision, and Recall are reported as mean $\pm$ std across the five selected target variables. Time is reported in seconds. Dist/Prec/Rec/Time denote Distance/Precision/Recall/Runtime.

Nodes	Algorithm	Distance	Precision	Recall	Time (s)↓
50	Notears-MLP	**0.77±0.06**	**0.80±0.11**	0.26±0.06	234.46
	DAG-GNN	1.15±0.25	0.60±0.55	0.02±0.02	1336.21
	GraN-DAG	0.99±0.24	0.80±0.45	0.10±0.06	1401.21
	ELCS	0.94±0.06	0.39±0.12	0.30±0.13	0.83
	PCD-by-PCD	0.87±0.08	0.39±0.05	0.37±0.07	**0.35**
	MB-by-MB	1.02±0.04	0.15±0.02	**0.43±0.05**	1.55
	PSL	0.90±0.05	0.34±0.03	0.38±0.04	0.38
	TCSL	0.78±0.04	0.52±0.03	0.39±0.03	28.21
100	Notears-MLP	0.80±0.09	0.74±0.12	0.26±0.13	1032.63
	DAG-GNN	1.06±0.05	0.45±0.05	0.10±0.06	1592.47
	GraN-DAG	1.05±0.20	**0.80±0.45**	0.03±0.03	4592.33
	ELCS	0.82±0.11	0.49±0.03	0.37±0.12	1.80
	PCD-by-PCD	0.90±0.10	0.40±0.04	0.33±0.12	**0.72**
	MB-by-MB	1.01±0.05	0.16±0.01	**0.45±0.09**	4.98
	PSL	0.79±0.08	0.44±0.02	0.44±0.09	0.95
	TCSL	**0.66±0.04**	0.67±0.03	0.43±0.07	49.26
200	Notears-MLP	0.83±0.11	**0.73±0.14**	0.22±0.08	4783.10
	DAG-GNN	1.20±0.16	0.38±0.37	0.02±0.01	2097.40
	GraN-DAG	1.24±0.25	0.40±0.55	0.01±0.02	18277.16
	ELCS	0.96±0.11	0.42±0.09	0.23±0.08	2.17
	PCD-by-PCD	0.84±0.09	0.49±0.07	0.33±0.09	**0.91**
	MB-by-MB	0.99±0.02	0.19±0.02	**0.43±0.01**	20.87
	PSL	**0.78±0.06**	0.46±0.02	**0.43±0.07**	2.83
	TCSL	0.79±0.11	0.61±0.06	0.32±0.10	47.41
300	Notears-MLP	0.89±0.06	**0.79±0.06**	0.14±0.06	13124.16
	DAG-GNN	1.20±0.15	0.32±0.23	0.03±0.02	4892.13
	GraN-DAG	1.24±0.28	0.40±0.55	0.01±0.01	49323.28
	ELCS	0.94±0.04	0.38±0.04	0.29±0.05	9.48
	PCD-by-PCD	0.93±0.04	0.42±0.04	0.28±0.07	**3.07**
	MB-by-MB	1.09±0.06	0.12±0.04	0.36±0.05	76.86
	PSL	0.88±0.05	0.40±0.06	**0.37±0.03**	6.85
	TCSL	**0.72±0.12**	0.66±0.03	**0.37±0.14**	292.33
500	Notears-MLP	1.01±0.03	**0.61±0.04**	0.08±0.02	63334.11
	DAG-GNN	1.29±0.12	0.19±0.19	0.01±0.01	13022.77
	GraN-DAG	–	–	–	–
	ELCS	1.02±0.09	0.34±0.05	0.22±0.08	27.67
	PCD-by-PCD	0.98±0.02	0.40±0.03	0.22±0.01	**8.84**
	MB-by-MB	1.13±0.03	0.07±0.01	**0.36±0.03**	235.17
	PSL	0.95±0.02	0.37±0.04	0.29±0.05	17.88
	TCSL	**0.89±0.05**	0.59±0.01	0.21±0.06	724.65

Table 4: Experimental results for query hop budget $k = 1$ with $n = 1000$ samples (Nonlinear ANM with GP).

Nodes	Algorithm	Distance	Precision	Recall	Time (s)↓
50	Notears-MLP	0.56±0.29	**0.90±0.22**	0.47±0.25	178.61
	DAG-GNN	0.63±0.23	0.83±0.24	0.42±0.19	273.35
	GraN-DAG	1.13±0.39	0.40±0.55	0.12±0.16	506.48
	ELCS	0.67±0.19	0.52±0.04	0.59±0.30	0.07
	PCD-by-PCD	1.18±0.33	0.20±0.27	0.14±0.22	**0.03**
	MB-by-MB	0.76±0.44	0.57±0.43	0.42±0.28	0.14
	PSL	0.78±0.39	0.57±0.37	0.38±0.26	0.12
	TCSL	**0.51±0.17**	0.76±0.23	**0.60±0.13**	9.09
100	Notears-MLP	0.89±0.37	0.36±0.24	0.40±0.33	912.72
	DAG-GNN	1.21±0.30	0.27±0.44	0.09±0.12	306.12
	GraN-DAG	1.28±0.30	0.20±0.25	0.05±0.11	1355.52
	ELCS	0.80±0.20	0.41±0.11	0.47±0.20	0.21
	PCD-by-PCD	1.31±0.15	0.10±0.15	0.05±0.07	**0.07**
	MB-by-MB	1.08±0.20	0.25±0.14	0.23±0.14	0.27
	PSL	0.80±0.25	0.47±0.14	0.42±0.22	0.35
	TCSL	**0.54±0.19**	**0.64±0.19**	**0.64±0.17**	8.59
200	Notears-MLP	1.12±0.29	0.16±0.18	0.29±0.32	6188.89
	DAG-GNN	0.74±0.41	**0.67±0.41**	0.38±0.25	410.56
	GraN-DAG	1.41±0.00	0.00±0.00	0.00±0.00	4174.76
	ELCS	1.04±0.38	0.27±0.25	0.27±0.30	0.13
	PCD-by-PCD	1.41±0.00	0.00±0.00	0.00±0.00	**0.05**
	MB-by-MB	1.30±0.26	0.10±0.22	0.07±0.15	0.27
	PSL	1.06±0.39	0.35±0.42	0.20±0.22	0.32
	TCSL	**0.62±0.11**	0.58±0.11	**0.55±0.10**	9.24
300	Notears-MLP	1.00±0.30	0.25±0.22	0.35±0.23	17049.77
	DAG-GNN	1.21±0.31	0.19±0.27	0.12±0.17	1152.56
	GraN-DAG	1.28±0.19	0.13±0.18	0.07±0.09	8245.53
	ELCS	0.79±0.33	0.55±0.31	0.40±0.31	0.43
	PCD-by-PCD	1.28±0.30	0.20±0.45	0.05±0.11	**0.12**
	MB-by-MB	0.98±0.21	0.37±0.19	0.26±0.11	1.23
	PSL	0.77±0.38	0.65±0.39	0.36±0.25	0.61
	TCSL	**0.41±0.08**	**0.76±0.14**	**0.70±0.06**	8.84
500	Notears-MLP	1.25±0.18	0.12±0.15	0.12±0.13	62800.61
	DAG-GNN	1.08±0.33	0.43±0.04	0.14±0.17	2594.90
	GraN-DAG	1.41±0.00	0.00±0.00	0.00±0.00	22251.33
	ELCS	0.97±0.16	0.42±0.17	0.23±0.10	0.94
	PCD-by-PCD	1.38±0.08	0.03±0.07	0.02±0.04	**0.23**
	MB-by-MB	1.36±0.13	0.04±0.09	0.04±0.09	3.48
	PSL	0.88±0.26	0.53±0.31	0.29±0.17	1.46
	TCSL	**0.46±0.11**	**0.74±0.16**	**0.65±0.10**	9.43

Table 5: Experimental results for query hop budget $k = 2$ with $n = 1000$ samples (Nonlinear ANM with GP).

Nodes	Algorithm	Distance	Precision	Recall	Time (s)↓
50	Notears-MLP	0.86±0.14	**0.81±0.21**	0.18±0.11	201.08
	DAG-GNN	1.04±0.11	0.60±0.28	0.09±0.03	278.94
	GraN-DAG	1.22±0.26	0.40±0.55	0.03±0.04	535.24
	ELCS	0.83±0.13	0.57±0.09	**0.29±0.14**	0.11
	PCD-by-PCD	1.11±0.28	0.36±0.36	0.12±0.11	**0.04**
	MB-by-MB	0.89±0.12	0.56±0.12	0.23±0.09	0.14
	PSL	**0.81±0.14**	0.64±0.13	0.28±0.11	0.12
	TCSL	0.90±0.11	0.57±0.11	0.21±0.09	13.87
100	Notears-MLP	0.96±0.29	0.40±0.26	0.26±0.19	903.28
	DAG-GNN	1.28±0.19	0.18±0.25	0.03±0.05	331.61
	GraN-DAG	1.34±0.16	0.10±0.22	0.02±0.04	1367.41
	ELCS	0.87±0.06	0.48±0.04	0.31±0.05	0.45
	PCD-by-PCD	1.00±0.08	0.45±0.12	0.17±0.06	**0.14**
	MB-by-MB	1.01±0.11	0.33±0.10	0.25±0.11	0.89
	PSL	**0.79±0.11**	0.53±0.08	**0.36±0.08**	0.44
	TCSL	0.86±0.10	**0.65±0.09**	0.22±0.08	17.93
200	Notears-MLP	1.29±0.11	0.07±0.08	0.11±0.12	6247.73
	DAG-GNN	**1.00±0.25**	**0.58±0.37**	0.13±0.08	399.01
	GraN-DAG	1.41±0.00	0.00±0.00	0.00±0.00	4431.25
	ELCS	1.12±0.27	0.34±0.31	0.11±0.11	0.25
	PCD-by-PCD	1.07±0.22	0.38±0.24	0.14±0.10	**0.10**
	MB-by-MB	1.15±0.16	0.23±0.14	0.15±0.10	0.97
	PSL	1.05±0.29	0.40±0.33	**0.16±0.13**	0.33
	TCSL	1.04±0.06	0.53±0.12	0.08±0.03	30.95
300	Notears-MLP	1.26±0.06	0.03±0.02	0.20±0.08	17622.13
	DAG-GNN	1.27±0.21	0.15±0.22	0.06±0.09	1202.68
	GraN-DAG	1.38±0.07	0.02±0.05	0.03±0.05	8132.66
	ELCS	0.95±0.19	0.59±0.30	0.18±0.13	0.95
	PCD-by-PCD	1.05±0.24	0.44±0.37	0.15±0.10	**0.35**
	MB-by-MB	1.15±0.10	0.20±0.10	0.19±0.07	5.57
	PSL	**0.92±0.16**	0.54±0.22	**0.22±0.08**	1.08
	TCSL	1.00±0.05	**0.62±0.16**	0.09±0.03	40.41
500	Notears-MLP	1.34±0.04	0.03±0.04	0.08±0.04	64310.72
	DAG-GNN	1.10±0.18	**0.50±0.36**	0.06±0.05	2741.22
	GraN-DAG	1.41±0.00	0.00±0.00	0.00±0.00	22514.62
	ELCS	1.00±0.10	0.44±0.13	0.18±0.07	14.08
	PCD-by-PCD	1.01±0.15	0.47±0.18	0.14±0.07	**0.84**
	MB-by-MB	1.16±0.12	0.16±0.07	0.19±0.11	23.31
	PSL	**0.97±0.10**	0.42±0.10	**0.22±0.07**	7.95
	TCSL	1.04±0.05	0.49±0.10	0.09±0.03	70.25

Table 6: Experimental results for query hop budget $k = 3$ with $n = 1000$ samples (Nonlinear ANM with GP).

Nodes	Algorithm	Distance	Precision	Recall	Time (s)↓
50	Notears-MLP	0.94±0.09	0.81±0.21	0.10±0.07	174.26
	DAG-GNN	1.13±0.09	0.42±0.18	0.03±0.02	291.43
	GraN-DAG	1.24±0.24	0.40±0.55	0.01±0.01	551.21
	ELCS	0.92±0.07	0.53±0.05	0.21±0.07	0.16
	PCD-by-PCD	1.04±0.08	0.53±0.16	0.08±0.03	**0.06**
	MB-by-MB	1.01±0.03	0.36±0.04	0.22±0.05	0.23
	PSL	0.86±0.07	0.57±0.06	0.26±0.07	0.13
	TCSL	**0.56±0.10**	**0.84±0.04**	**0.46±0.10**	16.67
100	Notears-MLP	0.97±0.16	0.45±0.18	0.21±0.10	944.87
	DAG-GNN	1.31±0.14	0.14±0.19	0.02±0.02	327.53
	GraN-DAG	1.38±0.06	0.05±0.08	0.01±0.01	1305.82
	ELCS	0.92±0.04	0.48±0.04	0.24±0.05	0.78
	PCD-by-PCD	0.99±0.03	0.47±0.06	0.17±0.04	**0.27**
	MB-by-MB	1.07±0.05	0.24±0.10	0.25±0.07	1.44
	PSL	0.86±0.05	0.47±0.05	0.33±0.03	0.49
	TCSL	**0.66±0.08**	**0.82±0.03**	**0.36±0.08**	71.08
200	Notears-MLP	1.30±0.07	0.03±0.02	0.13±0.09	6724.61
	DAG-GNN	1.12±0.17	0.45±0.26	0.04±0.03	434.61
	GraN-DAG	1.41±0.00	0.00±0.00	0.00±0.00	4392.61
	ELCS	1.09±0.18	0.48±0.27	0.06±0.05	0.58
	PCD-by-PCD	1.06±0.20	0.49±0.28	0.09±0.07	**0.25**
	MB-by-MB	1.20±0.14	0.16±0.10	**0.14±0.10**	3.32
	PSL	1.10±0.22	0.35±0.23	0.13±0.12	0.82
	TCSL	**0.92±0.04**	**0.78±0.08**	0.11±0.03	70.06
300	Notears-MLP	1.28±0.02	0.02±0.01	**0.18±0.02**	18390.40
	DAG-GNN	1.31±0.15	0.14±0.20	0.02±0.03	1142.63
	GraN-DAG	1.39±0.02	0.01±0.01	0.02±0.02	8341.17
	ELCS	1.10±0.08	0.38±0.15	0.09±0.03	2.53
	PCD-by-PCD	1.04±0.11	0.49±0.18	0.10±0.03	**0.91**
	MB-by-MB	1.23±0.04	0.10±0.04	0.17±0.02	19.23
	PSL	1.00±0.08	0.42±0.08	0.18±0.05	2.57
	TCSL	**0.91±0.08**	**0.73±0.04**	0.13±0.07	100.17
500	Notears-MLP	1.32±0.03	0.01±0.01	0.12±0.04	62538.27
	DAG-GNN	1.18±0.16	0.43±0.37	0.02±0.01	2563.52
	GraN-DAG	1.41±0.00	0.00±0.00	0.00±0.01	22334.57
	ELCS	1.06±0.06	0.35±0.04	0.17±0.09	20.46
	PCD-by-PCD	1.01±0.07	0.44±0.08	0.11±0.03	**2.86**
	MB-by-MB	1.21±0.01	0.07±0.01	0.23±0.02	94.08
	PSL	0.97±0.02	0.37±0.02	**0.36±0.03**	11.31
	TCSL	**0.91±0.02**	**0.72±0.03**	0.13±0.02	1003.63

Table 7: Experimental results for query hop budget $k = 1$ with $n = 5000$ samples (Nonlinear ANM with GP).

Nodes	Algorithm	Distance	Precision	Recall	Time (s)↓
50	Notears-MLP	**0.42±0.26**	**0.90±0.22**	**0.62±0.22**	211.34
	DAG-GNN	0.90±0.48	0.60±0.55	0.27±0.25	1459.52
	GraN-DAG	0.64±0.51	0.80±0.45	0.45±0.36	1388.45
	ELCS	0.72±0.17	0.58±0.27	0.45±0.08	0.18
	PCD-by-PCD	1.17±0.40	0.25±0.43	0.14±0.22	0.11
	MB-by-MB	1.08±0.22	0.18±0.12	0.30±0.19	0.43
	PSL	0.88±0.34	0.37±0.24	0.38±0.24	**0.10**
	TCSL	0.46±0.03	0.86±0.12	0.57±0.04	8.54
100	Notears-MLP	0.48±0.20	0.79±0.22	0.64±0.26	1059.24
	DAG-GNN	0.70±0.26	0.58±0.28	0.50±0.23	1525.42
	GraN-DAG	0.84±0.35	**0.80±0.45**	0.24±0.20	4693.19
	ELCS	0.68±0.17	0.47±0.13	0.59±0.13	0.40
	PCD-by-PCD	0.82±0.05	0.41±0.03	0.44±0.04	**0.38**
	MB-by-MB	0.92±0.26	0.29±0.16	0.42±0.25	0.92
	PSL	0.81±0.14	0.42±0.16	0.44±0.06	0.39
	TCSL	**0.45±0.13**	0.78±0.20	**0.66±0.11**	10.08
200	Notears-MLP	0.74±0.39	0.56±0.38	0.46±0.31	4740.64
	DAG-GNN	1.21±0.33	0.24±0.43	0.10±0.15	2035.18
	GraN-DAG	1.08±0.46	0.40±0.55	0.17±0.24	18250.83
	ELCS	0.66±0.43	0.48±0.29	0.65±0.41	0.36
	PCD-by-PCD	1.22±0.29	0.13±0.21	0.15±0.20	**0.10**
	MB-by-MB	1.05±0.25	0.24±0.18	0.27±0.19	1.48
	PSL	0.73±0.30	0.46±0.23	0.54±0.30	0.30
	TCSL	**0.58±0.14**	**0.57±0.06**	**0.65±0.20**	8.57
300	Notears-MLP	**0.53±0.08**	**0.84±0.15**	**0.52±0.13**	13164.28
	DAG-GNN	0.90±0.31	0.50±0.37	0.29±0.16	4832.60
	GraN-DAG	1.17±0.34	0.40±0.55	0.08±0.12	48776.95
	ELCS	0.84±0.25	0.40±0.16	0.42±0.22	0.99
	PCD-by-PCD	1.28±0.19	0.08±0.11	0.12±0.17	**0.28**
	MB-by-MB	1.05±0.36	0.23±0.23	0.30±0.31	3.82
	PSL	0.96±0.36	0.33±0.28	0.32±0.24	1.71
	TCSL	0.63±0.28	0.67±0.28	0.46±0.18	10.50
500	Notears-MLP	0.99±0.31	0.47±0.39	0.21±0.16	60628.49
	DAG-GNN	1.28±0.30	0.15±0.34	0.06±0.13	12719.34
	GraN-DAG	-	-	-	-
	ELCS	0.56±0.15	0.65±0.21	0.59±0.11	1.97
	PCD-by-PCD	1.28±0.23	0.10±0.19	0.09±0.14	**0.46**
	MB-by-MB	0.91±0.19	0.30±0.13	0.43±0.17	10.31
	PSL	0.57±0.22	**0.73±0.25**	0.54±0.18	6.19
	TCSL	**0.51±0.24**	0.71±0.17	**0.62±0.24**	12.03

Table 8: Experimental results for query hop budget $k = 2$ with $n = 5000$ samples (Nonlinear ANM with GP).

Nodes	Algorithm	Distance	Precision	Recall	Time (s)↓
50	Notears-MLP	0.78±0.15	0.74±0.25	0.29±0.11	217.30
	DAG-GNN	1.10±0.29	0.60±0.55	0.07±0.08	1530.23
	GraN-DAG	0.88±0.33	**0.80±0.45**	0.21±0.17	1345.67
	ELCS	0.97±0.18	0.38±0.18	0.26±0.14	0.43
	PCD-by-PCD	0.99±0.19	0.33±0.14	0.27±0.13	**0.21**
	MB-by-MB	0.93±0.06	0.21±0.03	0.51±0.06	1.18
	PSL	0.84±0.06	0.39±0.05	0.43±0.05	0.38
	TCSL	**0.64±0.10**	0.56±0.05	**0.53±0.12**	15.14
100	Notears-MLP	0.70±0.14	0.72±0.08	0.37±0.18	1109.41
	DAG-GNN	0.97±0.10	0.54±0.26	0.18±0.09	1566.13
	GraN-DAG	0.99±0.24	**0.80±0.45**	0.09±0.08	4788.56
	ELCS	0.82±0.10	0.51±0.05	0.36±0.15	0.89
	PCD-by-PCD	0.91±0.05	0.44±0.16	0.31±0.09	**0.38**
	MB-by-MB	1.00±0.05	0.21±0.11	0.40±0.09	2.84
	PSL	0.72±0.11	0.55±0.12	0.45±0.12	0.73
	TCSL	**0.68±0.18**	0.58±0.08	**0.46±0.18**	22.51
200	Notears-MLP	0.74±0.19	0.64±0.18	0.37±0.15	4755.23
	DAG-GNN	1.22±0.21	0.30±0.42	0.05±0.05	2138.77
	GraN-DAG	1.20±0.29	0.40±0.55	0.04±0.07	18393.28
	ELCS	0.76±0.29	0.52±0.23	0.41±0.20	0.91
	PCD-by-PCD	0.83±0.28	0.42±0.19	0.40±0.21	**0.29**
	MB-by-MB	0.89±0.16	0.27±0.08	**0.51±0.19**	6.35
	PSL	0.71±0.24	0.51±0.18	0.49±0.18	1.61
	TCSL	**0.62±0.11**	**0.65±0.08**	0.50±0.12	13.9
300	Notears-MLP	0.71±0.09	**0.84±0.13**	0.31±0.10	12977.43
	DAG-GNN	1.13±0.18	0.35±0.23	0.09±0.06	4801.35
	GraN-DAG	1.22±0.27	0.40±0.55	0.03±0.04	48032.13
	ELCS	0.85±0.20	0.42±0.13	0.39±0.18	2.93
	PCD-by-PCD	0.83±0.17	0.45±0.10	0.38±0.16	**0.87**
	MB-by-MB	1.04±0.17	0.20±0.10	0.34±0.15	22.08
	PSL	0.87±0.10	0.42±0.09	0.36±0.08	3.42
	TCSL	**0.59±0.17**	0.63±0.07	**0.56±0.20**	21.06
500	Notears-MLP	1.02±0.14	0.47±0.16	0.14±0.07	62437.54
	DAG-GNN	1.32±0.17	0.15±0.26	0.02±0.03	12311.54
	GraN-DAG	-	-	-	-
	ELCS	0.88±0.17	0.53±0.16	0.27±0.15	6.97
	PCD-by-PCD	1.00±0.18	0.36±0.14	0.23±0.12	**2.44**
	MB-by-MB	1.04±0.06	0.16±0.06	**0.38±0.04**	61.70
	PSL	0.92±0.07	0.39±0.04	0.31±0.06	12.77
	TCSL	**0.73±0.09**	**0.63±0.05**	0.38±0.11	26.45

Table 9: Experimental results for query hop budget $k = 1$ with $n = 1000$ samples (Nonlinear Additive Model with GP).

Nodes	Algorithm	Distance	Precision	Recall	Time (s)↓
50	Notears-MLP	1.05±0.20	0.21±0.17	0.31±0.12	219.84
	DAG-GNN	1.16±0.08	0.08±0.03	0.31±0.12	260.35
	GraN-DAG	1.29±0.28	0.20±0.45	0.04±0.09	471.50
	ELCS	0.93±0.47	0.35±0.34	0.33±0.33	0.14
	PCD-by-PCD	1.20±0.32	0.17±0.24	0.14±0.22	**0.04**
	MB-by-MB	0.84±0.38	0.38±0.27	**0.43±0.27**	0.23
	PSL	0.93±0.47	0.35±0.34	0.33±0.33	0.12
	TCSL	**0.84±0.31**	**0.59±0.38**	0.34±0.21	8.42
100	Notears-MLP	0.95±0.32	0.30±0.25	0.37±0.24	841.98
	DAG-GNN	1.09±0.24	0.13±0.13	0.36±0.24	242.02
	GraN-DAG	1.30±0.17	0.09±0.15	0.07±0.09	1241.80
	ELCS	0.75±0.26	**0.52±0.22**	0.45±0.23	0.33
	PCD-by-PCD	0.99±0.36	0.25±0.19	0.39±0.37	**0.09**
	MB-by-MB	**0.65±0.44**	0.50±0.33	**0.58±0.30**	0.60
	PSL	0.77±0.34	0.46±0.22	0.46±0.28	0.23
	TCSL	0.84±0.43	0.45±0.35	0.39±0.29	9.86
200	Notears-MLP	1.30±0.16	0.05±0.08	0.11±0.15	7422.12
	DAG-GNN	1.22±0.20	0.13±0.14	0.14±0.14	472.83
	GraN-DAG	1.41±0.00	0.00±0.00	0.00±0.00	4086.95
	ELCS	0.58±0.41	**0.69±0.34**	0.55±0.32	0.26
	PCD-by-PCD	1.09±0.45	0.30±0.45	0.23±0.35	**0.07**
	MB-by-MB	0.90±0.40	0.47±0.42	0.32±0.24	0.71
	PSL	**0.58±0.19**	0.68±0.21	**0.59±0.28**	0.24
	TCSL	0.81±0.30	0.44±0.23	0.42±0.24	9.35
300	Notears-MLP	0.92±0.14	0.24±0.12	0.49±0.12	14330.69
	DAG-GNN	0.90±0.34	0.36±0.25	0.38±0.25	968.23
	GraN-DAG	1.20±0.29	0.27±0.44	0.10±0.15	8596.39
	ELCS	**0.57±0.36**	0.61±0.25	**0.59±0.28**	0.53
	PCD-by-PCD	1.17±0.23	0.18±0.19	0.17±0.16	**0.15**
	MB-by-MB	1.02±0.08	0.28±0.07	0.27±0.06	1.80
	PSL	0.63±0.42	0.57±0.30	0.55±0.32	0.61
	TCSL	0.73±0.30	**0.66±0.35**	0.42±0.19	9.54
500	Notears-MLP	1.06±0.11	0.18±0.05	0.32±0.13	60280.11
	DAG-GNN	1.07±0.22	0.20±0.13	0.29±0.19	3138.33
	GraN-DAG	1.01±0.24	0.50±0.35	0.16±0.09	22067.13
	ELCS	0.79±0.27	0.48±0.20	0.41±0.23	0.81
	PCD-by-PCD	1.10±0.37	0.21±0.23	0.26±0.34	**0.25**
	MB-by-MB	1.09±0.27	0.19±0.16	0.28±0.25	5.10
	PSL	0.81±0.09	0.44±0.07	**0.41±0.11**	0.98
	TCSL	**0.78±0.18**	**0.69±0.28**	0.34±0.19	15.10

Table 10: Experimental results for query hop budget $k = 2$ with $n = 1000$ samples (Nonlinear Additive Model with GP).

Nodes	Algorithm	Distance	Precision	Recall	Time (s)↓
50	Notears-MLP	1.20±0.08	0.12±0.05	0.18±0.07	220.71
	DAG-GNN	1.25±0.03	0.06±0.01	0.18±0.06	271.83
	GraN-DAG	1.31±0.23	0.20±0.45	0.02±0.04	483.14
	ELCS	1.07±0.25	0.34±0.23	0.17±0.13	0.18
	PCD-by-PCD	0.95±0.32	0.43±0.29	0.25±0.19	**0.10**
	MB-by-MB	0.88±0.10	0.26±0.08	**0.53±0.08**	0.70
	PSL	**0.79±0.16**	0.48±0.10	0.40±0.13	0.21
	TCSL	0.92±0.07	**0.49±0.10**	0.24±0.05	11.21
100	Notears-MLP	1.05±0.12	0.25±0.13	0.28±0.09	849.75
	DAG-GNN	1.19±0.08	0.06±0.04	0.27±0.08	251.37
	GraN-DAG	1.35±0.05	0.02±0.02	0.07±0.07	1293.24
	ELCS	0.82±0.02	0.49±0.09	0.38±0.07	1.11
	PCD-by-PCD	0.89±0.10	0.47±0.06	0.29±0.09	**0.24**
	MB-by-MB	0.97±0.06	0.20±0.05	0.45±0.06	2.12
	PSL	**0.73±0.14**	0.50±0.09	**0.48±0.12**	0.60
	TCSL	0.77±0.09	**0.64±0.09**	0.33±0.12	16.90
200	Notears-MLP	1.34±0.07	0.03±0.03	0.07±0.07	7134.63
	DAG-GNN	1.35±0.06	0.06±0.06	0.03±0.03	493.11
	GraN-DAG	1.41±0.00	0.00±0.00	0.00±0.00	4211.60
	ELCS	**0.77±0.13**	**0.68±0.15**	0.31±0.10	0.53
	PCD-by-PCD	0.85±0.20	0.56±0.20	0.28±0.13	**0.18**
	MB-by-MB	0.89±0.12	0.37±0.22	**0.40±0.14**	3.83
	PSL	0.78±0.07	0.56±0.09	0.36±0.06	0.98
	TCSL	0.92±0.07	0.59±0.11	0.18±0.03	18.69
300	Notears-MLP	1.22±0.11	0.05±0.02	0.26±0.18	14852.31
	DAG-GNN	1.08±0.22	0.30±0.22	0.19±0.13	934.77
	GraN-DAG	1.25±0.19	0.28±0.43	0.04±0.04	8447.23
	ELCS	0.81±0.33	0.56±0.25	0.34±0.27	1.73
	PCD-by-PCD	0.88±0.12	0.49±0.09	0.29±0.12	**0.54**
	MB-by-MB	1.06±0.09	0.19±0.06	0.32±0.09	11.97
	PSL	**0.60±0.23**	**0.63±0.17**	**0.54±0.17**	4.16
	TCSL	0.90±0.08	0.62±0.21	0.22±0.11	28.49
500	Notears-MLP	1.25±0.12	0.06±0.03	0.18±0.14	60482.26
	DAG-GNN	1.24±0.12	0.13±0.09	0.12±0.09	3206.83
	GraN-DAG	1.29±0.19	0.21±0.44	0.04±0.02	22141.77
	ELCS	1.00±0.14	0.44±0.16	0.18±0.10	2.36
	PCD-by-PCD	0.94±0.15	0.50±0.17	0.21±0.09	**1.00**
	MB-by-MB	1.08±0.10	0.14±0.04	0.35±0.12	43.29
	PSL	**0.79±0.06**	0.48±0.06	**0.41±0.04**	7.65
	TCSL	0.88±0.08	**0.58±0.08**	0.23±0.07	50.65

Table 11: Experimental results for query hop budget $k = 3$ with $n = 1000$ samples (Nonlinear Additive Model with GP).

Nodes	Algorithm	Distance	Precision	Recall	Time (s)↓
50	Notears-MLP	1.07±0.07	0.22±0.04	0.26±0.06	234.12
	DAG-GNN	1.14±0.06	0.13±0.04	0.27±0.05	266.52
	GraN-DAG	1.32±0.21	0.20±0.45	0.01±0.02	476.21
	ELCS	1.05±0.22	0.38±0.22	0.16±0.12	0.23
	PCD-by-PCD	0.85±0.06	0.50±0.08	0.32±0.06	**0.18**
	MB-by-MB	1.03±0.04	0.13±0.02	0.44±0.04	1.13
	PSL	0.76±0.03	0.48±0.02	**0.45±0.03**	0.26
	TCSL	**0.73±0.06**	**0.66±0.03**	0.35±0.06	16.29
100	Notears-MLP	1.12±0.06	0.19±0.06	0.23±0.06	820.44
	DAG-GNN	1.15±0.05	0.09±0.03	0.30±0.04	257.60
	GraN-DAG	1.28±0.03	0.05±0.01	0.15±0.03	1175.21
	ELCS	0.80±0.08	0.49±0.01	0.39±0.11	1.74
	PCD-by-PCD	0.88±0.04	0.47±0.04	0.30±0.04	**0.49**
	MB-by-MB	1.01±0.04	0.13±0.02	0.49±0.05	3.79
	PSL	**0.70±0.03**	0.49±0.03	**0.52±0.03**	0.76
	TCSL	0.77±0.05	**0.70±0.06**	0.30±0.05	32.10
200	Notears-MLP	1.32±0.08	0.04±0.04	0.10±0.07	7492.31
	DAG-GNN	1.33±0.05	0.09±0.04	0.04±0.03	476.33
	GraN-DAG	1.41±0.02	0.06±0.01	0.01±0.02	4051.62
	ELCS	0.93±0.11	0.57±0.17	0.18±0.06	1.15
	PCD-by-PCD	0.99±0.13	0.45±0.15	0.18±0.08	**0.53**
	MB-by-MB	1.04±0.08	0.14±0.04	**0.42±0.08**	14.60
	PSL	**0.76±0.07**	0.56±0.11	0.38±0.02	1.53
	TCSL	0.99±0.03	**0.61±0.04**	0.09±0.02	68.83
300	Notears-MLP	1.24±0.03	0.03±0.01	0.22±0.06	14302.14
	DAG-GNN	1.22±0.09	0.20±0.10	0.08±0.05	1011.42
	GraN-DAG	1.30±0.18	0.23±0.43	0.01±0.01	8490.17
	ELCS	0.89±0.17	0.55±0.15	0.25±0.16	4.29
	PCD-by-PCD	0.91±0.17	0.47±0.12	0.27±0.12	**1.84**
	MB-by-MB	1.06±0.05	0.11±0.04	0.44±0.05	36.59
	PSL	**0.67±0.17**	0.56±0.13	**0.50±0.13**	5.08
	TCSL	0.94±0.02	**0.69±0.09**	0.11±0.05	113.10
500	Notears-MLP	1.28±0.06	0.04±0.01	0.16±0.08	61007.21
	DAG-GNN	1.31±0.06	0.09±0.04	0.06±0.04	3194.21
	GraN-DAG	1.29±0.16	0.16±0.33	0.05±0.02	22093.81
	ELCS	1.04±0.10	0.43±0.10	0.13±0.08	6.18
	PCD-by-PCD	1.03±0.08	0.40±0.06	0.16±0.06	**3.60**
	MB-by-MB	1.14±0.01	0.06±0.01	0.36±0.02	137.44
	PSL	**0.83±0.04**	0.45±0.03	**0.37±0.07**	10.73
	TCSL	0.99±0.02	**0.57±0.02**	0.11±0.02	1153.79

Table 12: Experimental results for query hop budget $k = 1$ with $n = 5000$ samples (Nonlinear Additive Model with GP).

Nodes	Algorithm	Distance	Precision	Recall	Time (s)↓
50	Notears-MLP	1.24 ± 0.24	0.12 ± 0.16	0.13 ± 0.19	283.34
	DAG-GNN	1.26 ± 0.23	0.13 ± 0.22	0.09 ± 0.12	975.47
	GraN-DAG	0.64 ± 0.51	$\mathbf{0.80 \pm 0.45}$	0.45 ± 0.36	1388.45
	ELCS	0.67 ± 0.28	0.58 ± 0.30	$\mathbf{0.50 \pm 0.16}$	0.30
	PCD-by-PCD	1.05 ± 0.33	0.21 ± 0.20	0.31 ± 0.30	**0.12**
	MB-by-MB	0.89 ± 0.24	0.29 ± 0.15	0.49 ± 0.25	1.01
	PSL	0.87 ± 0.37	0.34 ± 0.23	0.44 ± 0.32	0.61
	TCSL	$\mathbf{0.62 \pm 0.16}$	0.67 ± 0.12	0.49 ± 0.18	11.06
100	Notears-MLP	1.04 ± 0.29	0.26 ± 0.17	0.29 ± 0.28	2395.69
	DAG-GNN	1.14 ± 0.21	0.18 ± 0.13	0.22 ± 0.19	1431.50
	GraN-DAG	0.84 ± 0.35	$\mathbf{0.80 \pm 0.45}$	0.24 ± 0.20	4693.19
	ELCS	0.56 ± 0.41	0.57 ± 0.29	$\mathbf{0.68 \pm 0.33}$	0.91
	PCD-by-PCD	1.16 ± 0.29	0.14 ± 0.16	0.23 ± 0.26	**0.21**
	MB-by-MB	0.85 ± 0.17	0.31 ± 0.12	0.52 ± 0.18	4.91
	PSL	0.65 ± 0.27	0.47 ± 0.21	0.62 ± 0.20	1.00
	TCSL	$\mathbf{0.48 \pm 0.11}$	0.78 ± 0.20	0.62 ± 0.10	11.13
200	Notears-MLP	1.18 ± 0.23	0.20 ± 0.22	0.14 ± 0.14	9620.03
	DAG-GNN	1.28 ± 0.20	0.13 ± 0.22	0.07 ± 0.10	2461.21
	GraN-DAG	1.29 ± 0.28	0.20 ± 0.45	0.04 ± 0.09	18017.41
	ELCS	$\mathbf{0.42 \pm 0.29}$	$\mathbf{0.72 \pm 0.28}$	$\mathbf{0.75 \pm 0.23}$	0.83
	PCD-by-PCD	0.94 ± 0.32	0.26 ± 0.16	0.49 ± 0.40	**0.21**
	MB-by-MB	1.01 ± 0.29	0.24 ± 0.16	0.35 ± 0.28	4.46
	PSL	0.46 ± 0.30	0.63 ± 0.24	0.75 ± 0.23	1.80
	TCSL	0.57 ± 0.09	0.69 ± 0.17	0.55 ± 0.07	9.93
300	Notears-MLP	1.13 ± 0.33	0.22 ± 0.27	0.19 ± 0.21	19349.28
	DAG-GNN	1.09 ± 0.30	0.22 ± 0.20	0.25 ± 0.23	4845.27
	GraN-DAG	1.41 ± 0.00	0.00 ± 0.00	0.00 ± 0.00	37602.27
	ELCS	0.56 ± 0.20	0.58 ± 0.26	$\mathbf{0.66 \pm 0.06}$	1.54
	PCD-by-PCD	1.16 ± 0.25	0.18 ± 0.19	0.18 ± 0.18	**0.45**
	MB-by-MB	1.01 ± 0.30	0.22 ± 0.20	0.36 ± 0.25	18.94
	PSL	0.65 ± 0.32	0.48 ± 0.25	0.63 ± 0.24	4.74
	TCSL	$\mathbf{0.51 \pm 0.14}$	$\mathbf{0.70 \pm 0.12}$	0.61 ± 0.12	11.27
500	Notears-MLP	1.20 ± 0.20	0.14 ± 0.12	0.17 ± 0.12	68673.19
	DAG-GNN	1.33 ± 0.20	0.10 ± 0.22	0.03 ± 0.08	21492.67
	GraN-DAG	-	-	-	-
	ELCS	0.40 ± 0.35	0.68 ± 0.30	0.78 ± 0.23	4.95
	PCD-by-PCD	1.26 ± 0.18	0.08 ± 0.10	0.14 ± 0.17	**1.20**
	MB-by-MB	0.77 ± 0.28	0.32 ± 0.18	0.67 ± 0.28	176.48
	PSL	0.87 ± 0.32	0.30 ± 0.21	0.48 ± 0.26	9.49
	TCSL	$\mathbf{0.38 \pm 0.06}$	$\mathbf{0.70 \pm 0.05}$	$\mathbf{0.78 \pm 0.08}$	13.69

Table 13: Experimental results for query hop budget $k = 2$ with $n = 5000$ samples (Nonlinear Additive Model with GP).

Nodes	Algorithm	Distance	Precision	Recall	Time (s)↓
50	Notears-MLP	1.08 ± 0.16	0.28 ± 0.14	0.21 ± 0.12	289.50
	DAG-GNN	1.29 ± 0.03	0.13 ± 0.05	0.06 ± 0.02	1087.71
	GraN-DAG	0.88 ± 0.33	$\mathbf{0.80 \pm 0.45}$	0.21 ± 0.17	1345.67
	ELCS	0.99 ± 0.11	0.42 ± 0.09	0.20 ± 0.09	**0.30**
	PCD-by-PCD	0.81 ± 0.22	0.44 ± 0.14	0.41 ± 0.17	0.47
	MB-by-MB	0.98 ± 0.06	0.15 ± 0.03	0.51 ± 0.07	4.49
	PSL	0.74 ± 0.15	0.43 ± 0.10	0.52 ± 0.13	0.83
	TCSL	$\mathbf{0.63 \pm 0.15}$	0.62 ± 0.21	$\mathbf{0.52 \pm 0.08}$	21.55
100	Notears-MLP	1.15 ± 0.11	0.22 ± 0.11	0.16 ± 0.09	2167.75
	DAG-GNN	1.20 ± 0.11	0.19 ± 0.11	0.12 ± 0.05	1382.61
	GraN-DAG	0.99 ± 0.24	$\mathbf{0.80 \pm 0.45}$	0.09 ± 0.08	4788.56
	ELCS	0.73 ± 0.07	0.50 ± 0.05	0.47 ± 0.06	2.78
	PCD-by-PCD	0.81 ± 0.15	0.45 ± 0.16	0.42 ± 0.11	**0.97**
	MB-by-MB	0.95 ± 0.05	0.13 ± 0.04	$\mathbf{0.62 \pm 0.02}$	23.18
	PSL	0.65 ± 0.13	0.48 ± 0.08	0.61 ± 0.11	2.55
	TCSL	$\mathbf{0.61 \pm 0.10}$	0.60 ± 0.09	0.55 ± 0.06	26.37
200	Notears-MLP	1.14 ± 0.13	0.23 ± 0.15	0.17 ± 0.10	10201.36
	DAG-GNN	1.09 ± 0.28	0.36 ± 0.31	0.14 ± 0.10	2631.55
	GraN-DAG	1.31 ± 0.23	0.20 ± 0.45	0.02 ± 0.05	18122.07
	ELCS	0.56 ± 0.23	$\mathbf{0.68 \pm 0.24}$	0.58 ± 0.20	3.37
	PCD-by-PCD	0.80 ± 0.19	0.44 ± 0.15	0.44 ± 0.14	**1.14**
	MB-by-MB	0.93 ± 0.07	0.20 ± 0.08	0.53 ± 0.08	38.41
	PSL	$\mathbf{0.54 \pm 0.08}$	0.58 ± 0.10	$\mathbf{0.67 \pm 0.05}$	5.06
	TCSL	0.66 ± 0.16	0.56 ± 0.16	0.52 ± 0.12	18.07
300	Notears-MLP	1.16 ± 0.17	0.24 ± 0.16	0.12 ± 0.09	19048.23
	DAG-GNN	1.18 ± 0.06	0.23 ± 0.09	0.12 ± 0.04	5112.81
	GraN-DAG	1.41 ± 0.00	0.00 ± 0.00	0.00 ± 0.00	36928.71
	ELCS	0.70 ± 0.16	$\mathbf{0.58 \pm 0.12}$	0.46 ± 0.16	5.96
	PCD-by-PCD	0.97 ± 0.18	0.33 ± 0.11	0.31 ± 0.15	**2.28**
	MB-by-MB	0.96 ± 0.06	0.17 ± 0.08	0.53 ± 0.08	139.39
	PSL	0.65 ± 0.13	0.50 ± 0.11	$\mathbf{0.59 \pm 0.10}$	12.75
	TCSL	$\mathbf{0.63 \pm 0.09}$	0.57 ± 0.05	0.54 ± 0.08	28.43
500	Notears-MLP	1.25 ± 0.05	0.13 ± 0.05	0.10 ± 0.03	69442.31
	DAG-GNN	1.31 ± 0.14	0.14 ± 0.19	0.02 ± 0.03	18962.73
	GraN-DAG	-	-	-	-
	ELCS	0.73 ± 0.18	0.61 ± 0.17	0.40 ± 0.17	14.15
	PCD-by-PCD	0.86 ± 0.20	0.36 ± 0.12	0.42 ± 0.16	**6.59**
	MB-by-MB	0.97 ± 0.04	0.11 ± 0.02	$\mathbf{0.63 \pm 0.09}$	1285.01
	PSL	0.69 ± 0.17	0.43 ± 0.12	0.61 ± 0.14	28.14
	TCSL	$\mathbf{0.53 \pm 0.06}$	$\mathbf{0.69 \pm 0.04}$	0.57 ± 0.06	45.78

Table 14: Experimental results for query hop budget $k = 3$ with $n = 5000$ samples (Nonlinear Additive Model with GP).

Nodes	Algorithm	Distance	Precision	Recall	Time (s)↓
50	Notears-MLP	1.08 ± 0.05	0.28 ± 0.03	0.20 ± 0.04	277.61
	DAG-GNN	1.25 ± 0.03	0.17 ± 0.04	0.07 ± 0.02	906.95
	GraN-DAG	0.99 ± 0.24	$\mathbf{0.80 \pm 0.45}$	0.10 ± 0.06	1401.21
	ELCS	1.03 ± 0.11	0.43 ± 0.08	0.15 ± 0.12	**0.37**
	PCD-by-PCD	0.73 ± 0.08	0.46 ± 0.03	0.51 ± 0.09	0.78
	MB-by-MB	1.01 ± 0.03	0.10 ± 0.01	$\mathbf{0.54 \pm 0.04}$	5.36
	PSL	0.79 ± 0.06	0.39 ± 0.05	0.50 ± 0.04	0.90
	TCSL	$\mathbf{0.62 \pm 0.02}$	0.59 ± 0.03	0.53 ± 0.02	41.18
100	Notears-MLP	1.13 ± 0.09	0.25 ± 0.11	0.16 ± 0.07	2101.44
	DAG-GNN	1.18 ± 0.03	0.21 ± 0.03	0.13 ± 0.04	1473.55
	GraN-DAG	1.05 ± 0.20	$\mathbf{0.80 \pm 0.45}$	0.03 ± 0.03	4592.33
	ELCS	0.71 ± 0.06	0.44 ± 0.05	0.57 ± 0.05	6.61
	PCD-by-PCD	0.77 ± 0.10	0.41 ± 0.07	0.52 ± 0.08	**2.27**
	MB-by-MB	1.01 ± 0.02	0.07 ± 0.00	0.62 ± 0.04	31.15
	PSL	0.67 ± 0.03	0.43 ± 0.03	$\mathbf{0.64 \pm 0.02}$	2.71
	TCSL	$\mathbf{0.56 \pm 0.04}$	0.65 ± 0.05	0.56 ± 0.01	61.65
200	Notears-MLP	1.16 ± 0.04	0.24 ± 0.06	0.12 ± 0.04	10047.07
	DAG-GNN	1.16 ± 0.18	0.30 ± 0.23	0.09 ± 0.07	2498.37
	GraN-DAG	1.33 ± 0.19	0.20 ± 0.45	0.01 ± 0.01	18095.16
	ELCS	0.70 ± 0.13	0.60 ± 0.13	0.45 ± 0.21	8.86
	PCD-by-PCD	0.77 ± 0.11	0.45 ± 0.04	0.46 ± 0.12	**3.74**
	MB-by-MB	1.00 ± 0.04	0.09 ± 0.01	$\mathbf{0.59 \pm 0.08}$	101.06
	PSL	0.73 ± 0.09	0.49 ± 0.02	0.48 ± 0.11	7.21
	TCSL	$\mathbf{0.64 \pm 0.10}$	$\mathbf{0.62 \pm 0.09}$	0.48 ± 0.06	44.76
300	Notears-MLP	1.25 ± 0.10	0.17 ± 0.11	0.06 ± 0.04	19396.77
	DAG-GNN	1.23 ± 0.05	0.22 ± 0.06	0.06 ± 0.02	4936.34
	GraN-DAG	1.41 ± 0.00	0.00 ± 0.00	0.00 ± 0.00	37415.20
	ELCS	0.85 ± 0.16	0.47 ± 0.07	0.35 ± 0.19	15.47
	PCD-by-PCD	0.83 ± 0.06	0.41 ± 0.05	0.42 ± 0.06	**8.77**
	MB-by-MB	1.03 ± 0.02	0.07 ± 0.02	$\mathbf{0.57 \pm 0.04}$	268.77
	PSL	0.70 ± 0.05	0.48 ± 0.05	0.53 ± 0.04	16.66
	TCSL	$\mathbf{0.66 \pm 0.06}$	$\mathbf{0.64 \pm 0.04}$	0.45 ± 0.05	758.34
500	Notears-MLP	1.27 ± 0.03	0.14 ± 0.03	0.08 ± 0.02	72688.03
	DAG-GNN	1.29 ± 0.12	0.18 ± 0.17	0.01 ± 0.01	18130.26
	GraN-DAG	-	-	-	-
	ELCS	0.84 ± 0.70	0.45 ± 0.03	0.38 ± 0.10	42.55
	PCD-by-PCD	0.88 ± 0.03	0.34 ± 0.03	0.41 ± 0.02	**22.98**
	MB-by-MB	1.04 ± 0.02	0.04 ± 0.00	0.58 ± 0.04	2136.64
	PSL	0.71 ± 0.02	0.40 ± 0.01	$\mathbf{0.61 \pm 0.03}$	35.73
	TCSL	$\mathbf{0.67 \pm 0.03}$	$\mathbf{0.56 \pm 0.02}$	0.49 ± 0.03	2433.70