

THE ROBUSTNESS OF DIFFERENTIABLE CAUSAL DISCOVERY IN MISSPECIFIED SCENARIOS

Anonymous authors

Paper under double-blind review

ABSTRACT

Causal discovery aims to learn causal relationships between variables from targeted data, making it a fundamental task in machine learning. However, causal discovery algorithms often rely on unverifiable causal assumptions, which are usually difficult to satisfy in real-world data, thereby limiting the broad application of causal discovery in practical scenarios. Inspired by these considerations, this work extensively benchmarks the empirical performance of various mainstream causal discovery algorithms, which assume i.i.d. data, under eight model assumption violations. Our experimental results show that differentiable causal discovery methods exhibit robustness under the metrics of Structural Hamming Distance and Structural Intervention Distance of the inferred graphs in challenging scenarios, except for scale variation. We also provide the theoretical explanations for the performance of differentiable causal discovery methods. Finally, our work aims to comprehensively benchmark the performance of recent differentiable causal discovery methods under model assumption violations, and provide the standard for reasonable evaluation of causal discovery, as well as to further promote its application in real-world scenarios.

1 INTRODUCTION

In the realm of modern science, numerous endeavors hinge upon the elucidation of underlying causal mechanisms. However, owing to practical constraints including costs, risks, and ethical implications, the execution of randomized experiments frequently proves unviable. Consequently, mining causal relationships from purely observational data, known as causal discovery, plays a crucial role in addressing causal questions such as intervention and counterfactual (Peters et al., 2017; Spirtes et al., 2001; Pearl, 2009; Pearl & Mackenzie, 2018).

Causal discovery encompasses a comprehensive suite of methodologies, primarily categorized into constraint-based, score-based, functional causal model-driven, and gradient-based approaches. These methods often rely on unverifiable causal assumptions as their foundation (Peters et al., 2017; Vowels et al., 2022). Constraint-based methods, notably PC (Spirtes & Glymour, 1991) and FCI (Spirtes et al., 1995), meticulously reconstruct causal graphs to the Markov equivalence class (MEC) through rigorous statistical independence tests, guided by the faithfulness assumption. Score-based techniques, such as GES (Chickering, 2002), employ a scoring function to quantify the congruence between an equivalence class and observed data, optimally searching the vast landscape of directed acyclic graphs (DAGs) to identify the MEC.

To transcend the limitation of solely identifying MECs from observational data, functional causal model-based methods, exemplified by LiNGAM (Shimizu et al., 2006), leverage precise assumptions regarding the functional class and noise distribution within structural equation models (SEMs), enabling the unambiguous identification of a unique DAG. Recently, Zheng et al. (2018) introduced gradient-based techniques (e.g. NOTEARS (Zheng et al., 2018)), which convert combinatorial acyclic constraints into smooth equality constraints and solve the optimization by transforming equality-constrained optimization into unconstrained optimization through the augmented Lagrangian method (ALM) (Nemirovsky, 1999). In some literature (Zhang et al., 2023; Liu et al., 2023), gradient-based methods are also referred to as differentiable causal discovery.

Apart from the various assumptions of the methods above, traditional approaches typically rely on causal sufficiency and no measurement error assumptions to simplify the problem (Peters et al., 2017;

Zhang et al., 2018). Real-world data often fail to meet all of these assumptions, and these are also impossible to verify adequately (Peters et al., 2017). Although some studies have considered the complexity of real data and developed algorithms targeted at latent confounders (Spirtes et al., 1995; Xie et al., 2020; Salehkaleybar et al., 2020; Cai et al., 2019; Kong et al., 2023; Cai et al., 2023), measurement error (Zhang et al., 2018; Dai et al., 2022), heterogeneity (Huang et al., 2020; Cai et al., 2020; Ghassami et al., 2017; 2018), scale variation (Shimizu et al., 2011; Reisach et al., 2023; Deng et al., 2024), and missing data (Tu et al., 2019a; Gao et al., 2022), the true mechanisms remain unclear when applied to real data. These specifically designed algorithms also cannot be effectively employed for real data. Therefore, the robustness of causal discovery algorithms in scenarios where model hypotheses are violated is of great importance.

Previous research (Heinze-Deml et al., 2018) mainly evaluated various constraint-based and score-based algorithms under different scenarios, only focusing on linear SEM. The work (Mooij et al., 2016) benchmarked causal discovery for nonlinear additive noise models and information-geometric approaches, limiting to bivariate scenarios. The previous study (Singh et al., 2017) primarily assessed algorithms that use only observational data, a mix of observational and interventional data, and active learning, but their algorithm outputs were restricted to MEC. Also, those works (Glymour et al., 2019; Vowels et al., 2022) reviewed the advancements in traditional causal discovery (constraint-based, score-based, and functional causal model-based) and differentiable causal discovery, respectively, but lacked experimental support. The recent work (Ng et al., 2024) conducted an experimental assessment of the advancements in differentiable causal discovery, illuminating the shortcomings of current methods. However, their evaluation overlooked the ubiquitous violation of model assumptions that characterize real-world applications. Conversely, another work (Montagna et al., 2023) benchmarked the efficacy of traditional causal discovery algorithms, encompassing score-matching techniques, under scenarios where model assumptions were violated. Nevertheless, their analysis did not encompass the recent strides made in differentiable causal discovery, and the misspecified conditions they evaluated were constrained in scope. Given that the application of causal discovery methods to real data inevitably entails the violation of one or more unidentified assumptions, and that algorithms premised on specific assumptions may falter in practical use, the robustness of causal discovery in such misspecified contexts assumes paramount importance.

Our study undertakes an exhaustive empirical evaluation of both established and cutting-edge causal discovery methodologies, comprehensively examining their performance under diverse scenarios with assumption violations. The misspecified scenarios encompassed in our analysis represent the most extensive coverage in the current research landscape. We meticulously evaluate mainstream causal discovery approaches, spanning constraint-based, score-based, functional causal model-based, gradient-based methodologies, among others, ensuring a holistic view of the field. Notably, our work fills a crucial gap in the literature by being the first to assess the performance of gradient-based methods across a wide array of misspecified scenarios. **Considering their practical implementation potential, it is important to evaluate their performance.** Our contributions can be summarized as follows:

- We conduct extensive large-scale experimental evaluations of twelve prominent causal discovery algorithms across eight pivotal model assumption violation scenarios. Our rigorous research endeavor involves executing over 70,000 experiments on more than 2,400 synthetic datasets, ensuring a comprehensive assessment of the algorithm capabilities.
- We delve into challenging scenarios such as heterogeneity, scale variation, missing data, and mechanism violation, thereby enriching the benchmark data landscape for model assumption violations. This also aims to foster more comprehensive benchmark testing and foster a more rational evaluation framework for future causal discovery algorithms.
- Our analysis of the experimental outcomes offers theoretical insights into the performance of linear differentiable causal discovery methods under certain misspecified scenarios. Recognizing the robustness demonstrated by differentiable methods in these challenging settings, we underscore the significant value of further in-depth research into differentiable causal discovery, as it holds the promise of advancing the field in novel and impactful directions.

2 BACKGROUND

In this section, we introduce the definition of causal discovery (Section 2.1), functional causal model-based (Section 2.2), score-based (Section 2.3) and gradient-based (Section 2.3) methods. For constraint-based methods, see Appendix A.2 and B.1.

2.1 TASK FORMULATION

A structural causal model \mathcal{M} (Pearl, 2009) consists of the set of endogenous variables $X = (X_1, \dots, X_d) \in \mathbb{R}^d$, exogenous variables $U = (U_1, \dots, U_d) \in \mathbb{R}^d$, and functional mechanisms $\mathcal{F} = (f_1, \dots, f_d)$. Each variable X_i is defined by a structural equation:

$$X_i = f_i(X_{pa(X_i)}, U_i), \forall i = 1, \dots, d, \quad (1)$$

where X_i is the i -th node variable, $pa(X_i)$ denote the parents of X_i , $f_i : \mathbb{R}^{|X_{pa(X_i)}|+1} \rightarrow \mathbb{R}$ is the causal structure function, and $U = (U_1, \dots, U_d)$ are jointly independent noise variables with covariance matrix $\Omega = \text{cov}(U) = \text{diag}(\sigma_1^2, \dots, \sigma_d^2)$.

The task of causal discovery is to infer a DAG \mathcal{G} that describes the causal relationships among variables from n independent and identically distributed (i.i.d.) observational data $\mathbf{X} \in \mathbb{R}^{n \times d}$, which are drawn from the joint probability distribution $P(X)$.

2.2 STRUCTURE IDENTIFIABILITY

To uniquely identify a DAG \mathcal{G} from purely observational data \mathbf{X} sampled from $P(X)$, we need to make assumptions about the SEM in (1). Considering a set of assumptions \mathcal{A} on a structural causal model (SCM) $\mathcal{M}_{\mathcal{A}} = (P(X), \mathcal{G})$, the graph \mathcal{G} is identifiable from $P(X)$ if there is no other SCM $\tilde{\mathcal{M}}_{\mathcal{A}} = (\tilde{P}(X), \tilde{\mathcal{G}})$ satisfying the same \mathcal{A} such that $\tilde{\mathcal{G}} \neq \mathcal{G}$ and $\tilde{P}(X) = P(X)$. Existing identifiable causal models include: linear non-Gaussian acyclic models (Shimizu et al., 2006), linear Gaussian models with equal noise variances (Peters & Bühlmann, 2014), post-nonlinear models (Zhang & Hyvarinen, 2012) and nonlinear additive noise models (Hoyer et al., 2008; Peters et al., 2014).

2.3 DIFFERENTIABLE SCORE-BASED CAUSAL DISCOVERY

Traditional score-based causal discovery defines a combinatorial optimization problem:

$$\min_{\mathcal{G}} F(\mathcal{G}; \mathbf{X}) = \mathcal{L}_{\text{rec}}(\mathcal{G}; \mathbf{X}) + \lambda \mathcal{L}_{\text{sparse}}(\mathcal{G}) \quad \text{s.t.} \quad \mathcal{G} \in \text{DAG}, \quad (2)$$

where F is a score function, $\mathcal{L}_{\text{rec}}(\mathcal{G}; \mathbf{X})$ represents the goodness-of-fit between the estimated DAG \mathcal{G} and the true DAG, $\mathcal{L}_{\text{sparse}}(\mathcal{G})$ denotes the sparsity regularization term and λ is a hyperparameter that controls the strength of regularization.

As the number of nodes rises, the total count of possible DAGs expands super-exponentially (Robinson, 1973). Consequently, most conventional score-based approaches utilize local heuristic search techniques, including greedy search (Chickering, 2002; Hauser & Bühlmann, 2012) and hill-climbing (Gámez et al., 2011; Tsamardinos et al., 2006).

In addition to search strategies, the design of score functions is also crucial. Commonly, score functions are classified into two categories: Bayesian-based scores and information-theoretic scores. Bayesian-based scores emphasize goodness-of-fit and enable the integration of prior knowledge, such as the Bayesian Dirichlet equivalent (Heckerman et al., 1995) and the K2 score (Kayaalp & Cooper, 2012). Information-theoretic scores, on the other hand, account for both model goodness-of-fit and complexity, including the Bayesian information criterion (Neath & Cavanaugh, 2012) and the Akaike information criterion (Akaike, 1998).

To overcome the challenges of combinatorial optimization, NOTEARS (Zheng et al., 2018) formulates the DAG structure learning task as:

$$\min_{\mathcal{G}} F(\mathcal{G}; \mathbf{X}) \quad \text{s.t.} \quad h(W(\mathcal{G})) = 0, \quad (3)$$

where $W(\mathcal{G}) \in \mathbb{R}^{d \times d}$ is a weighted adjacency matrix, d is the number of nodes, $h(W(\mathcal{G})) = 0$ is a differentiable equality DAG constraint.

$h(W(\mathcal{G})) = 0$ if and only if $W(\mathcal{G})$ is a DAG. Commonly used DAG constraints include $h(W(\mathcal{G})) = \text{Tr}(e^{W \circ W}) - d$ (Zheng et al., 2018), $h(W(\mathcal{G})) = \text{Tr}[(I + \alpha W \circ W)^d] - d$ ($\alpha > 0$) (Yu et al., 2019) and $h^s(W(\mathcal{G})) = -\log \det(sI - W \circ W) + d \log s$ ($s > 0$) (Bello et al., 2022). Furthermore, we can transform the equality-constrained optimization (3) into unconstrained optimization (4) using the ALM (Nemirovsky, 1999):

$$\min_{\mathcal{G}} F(\mathcal{G}; \mathbf{X}) + \alpha_t h(W(\mathcal{G})) + \frac{\mu_t}{2} |h(W(\mathcal{G}))|^2, \quad (4)$$

where α_t and u_t are the Lagrange multiplier and penalty parameter at the t -th iteration, respectively.

3 EXPERIMENTAL DESIGN

In this section, we introduce the generation of synthetic datasets with violated model assumptions, the tested causal discovery algorithms, the algorithm hyperparameters, and the evaluation metrics.

3.1 SYNTHETIC DATASETS

Many prevalent causal discovery approaches hinge upon unverifiable assumptions. Our study primarily scrutinizes the efficacy of these methodologies in circumstances where their underlying assumptions are breached. To achieve this, we commence by elucidating the baseline model under both linear and nonlinear frameworks, subsequently delving into scenarios where these fundamental assumptions fail to hold.

	PC	GES	DirectLINGAM	CAM	Var-SortnRegress	R ² -SortnRegress	NOTEARS	GOLEM	NOTEARS-MLP	Grain-DAG	NoCurl	DAGMA
Gaussian noise	✓	✓	✗	✓	✓	✓	✓	✓	✓	✓	✓	✓
Non-Gaussian noise	✓	✗	✓	✗	✓	✓	✓	✗	✓	✗	✓	✓
Linear mechanisms	✓	✓	✓	✗	✓	✓	✓	✓	✗	✗	✓	✓
Nonlinear mechanisms	✓	✓	✗	✓	✗	✗	✗	✗	✓	✓	✓	✓
Unfaithful distribution	✗	✗	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Confounding effects	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗
Measurement errors	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗
Autoregressive effects	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗
Heterogeneous effects	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗
Scale-variant effects	✗	✗	✓	✗	✗	✓	✗	✗	✗	✗	✗	✗
Missing mechanisms	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗
Output	CPDAG	CPDAG	DAG	DAG	DAG	DAG	DAG	DAG	DAG	DAG	DAG	DAG

Table 1: Summary of the algorithm assumptions and their corresponding output graph types. The content within the cells indicates whether an algorithm supports (✓) or does not support (✗) the specific condition in the corresponding row. The table style is adjusted from Montagna et al. (2023).

Linear vanilla model. In linear SCM, following the settings of Zheng et al. (2018), coefficients are sampled from $U(-2, -0.5) \cup U(0.5, 2)$ with additive standard Gaussian noise. We refer to this model as the linear vanilla model, which satisfies both identifiability and the assumptions of most linear benchmark methods (see Table 1).

Nonlinear vanilla model. In nonlinear settings, following the settings of Zheng et al. (2020), the SEM in equation (1) is generated under the Gaussian process with radial basis function kernel of bandwidth one, where f_i is additive noise models with U_i as a standard Gaussian noise. We refer to this model as the nonlinear vanilla model, which satisfies both identifiability and the assumptions of nonlinear benchmark methods (see Table 1).

To eliminate the impact of Gaussian noise in the vanilla model on experimental results, we also consider cases where the vanilla model uses non-Gaussian noise (see Appendix G).

3.1.1 MODEL ASSUMPTION VIOLATION SCENARIOS

Four scenarios of model assumption violations are defined below. The other four cases, i.e., confounded, measurement error, unfaithful and autoregressive model, follow the same settings as Montagna et al. (2023). These eight misspecified scenarios can be applied to both the linear vanilla and nonlinear vanilla model to generate datasets.

Heterogeneous model. Existing causal discovery algorithms typically rely on the assumption of i.i.d. data. However, real data often exhibit distribution shifts (Huang et al., 2020). The heterogeneous multi-domain data considered in this paper primarily refers to scenarios where the underlying causal generation process remains unchanged, but the distribution of noise terms varies (Huang et al., 2020; Zhang et al., 2023; Wang et al., 2022). Specifically, we consider data from two domains e_1 and e_2 . The proportion of data from e_1 is $P_1 \in \{0.1, 0.3, 0.5, 0.7, 0.9\}$, and the proportion from e_2 is $1 - P_1$. The noise variance in e_1 is set the same as the vanilla model, while variance in e_2 is set to $\gamma \in \{0.01, 0.05, 0.1, 0.5\}$.

Scale-variant model. Reisach et al. (2021) observed a significant performance decline in linear gradient-based methods, such as NOTEARS (Zheng et al., 2018) and GOLEM (Ng et al., 2020), when applied to data with scale variation. However, there has been a notable absence of research investigating the performance of nonlinear methods in the context of scale variation. Therefore, we also consider scale variation as a misspecified scenario. The structural equations considered are:

$$\overline{X_i} = \frac{X_i - u_i}{\sqrt{\text{Var}(X_i)}}, \forall i = 1, \dots, d, \quad (5)$$

where u_i and $\text{Var}(X_i)$ are the mean and variance of X_i , respectively. The input data are standardized, while the ground truth graph remains consistent with the causal graph that generates the original data.

Missing model. Missing data is a prevalent challenge in real-world datasets, necessitating that causal discovery algorithms effectively address this issue (Tu et al., 2019b). In our study, we adopt the Missing Completely At Random (MCAR) mechanism (Tu et al., 2019a), where the occurrence of missing values follows a Bernoulli distribution with a missingness probability of $\beta \in \{0.005, 0.01, 0.05, 0.1\}$. Given that the algorithms under consideration are incapable of directly processing datasets with missing values, we eliminate any records containing such gaps. To mitigate the influence of data quantity on the experimental outcomes, we ensure the volume of data remains consistent before and after the removal of incomplete records.

Mechanism violation. Most current causal discovery algorithms presuppose either linear or nonlinear mechanism, especially methods based on functional causal models (Shimizu et al., 2006; Peters & Bühlmann, 2014; Zhang & Hyvarinen, 2012; Hoyer et al., 2008; Peters et al., 2014). These methods necessitate specific assumptions about the SEM mechanism to guarantee identifiability. Given that the SEM in real-world data is typically unknown, the robustness of algorithms in the face of mechanism violation becomes critically important. In mechanism violation, the input data for algorithms designed for linear SEM will adhere to the nonlinear vanilla model, whereas the input data for algorithms tailored to nonlinear SEM will conform to the linear vanilla model.

3.1.2 DATA GENERATION

Following the data generation of Zheng et al. (2018; 2020) and Liu et al. (2023), different datasets are generated for both linear and nonlinear vanilla model. We simulate ER and SF graphs based on the number of nodes $d \in \{10, 20, 50\}$, average degree of nodes $k \in \{2, 4\}$. In addition, we consider scenarios with Gaussian Random Partitions (GRP) (Brandes et al., 2003) graph and an average node degree of 6. For each experimental configuration and scenario, 10 datasets of 2000 samples are generated. The mean and standard deviations of the evaluation metrics (Section 3.4) is reported to ensure a fair comparison.

3.2 METHODS

We select 12 mainstream causal discovery algorithms, including constraint-based, score-based, functional causal model-based, gradient-based and other representative methods. For a more detailed introduction to the various methods, see the Appendix B.

3.3 HYPERPARAMETER SETTINGS

PC (Spirtes & Glymour, 1991), CAM (Bühlmann et al., 2014), and GraN-DAG require adjustment of the significance level α in the statistical independence tests. NOTEARS, GOLEM, NOTEARS-MLP, and DAGMA need adjustment of the sparsity coefficient λ_1 for the l_1 -norm regularization term. Typically, the ground truth of real data is unknown, making it difficult to effectively select hyperparameters for various algorithms. Thus, to ensure a fair comparison of various methods, we tune λ_1 in $\{0.005, 0.01, 0.05, 0.5, 2, 5\}$ and tune α in $\{0.001, 0.01, 0.05, 0.1\}$.

3.4 EVALUATION METRICS

We employ Structural Hamming Distance (SHD) and Structural Intervention Distance (SID) (Peters & Bühlmann, 2015) to evaluate performance. SHD counts the number of edge insertions, deletions, and reversals necessary to transform the estimated graph into the true graph. SID is used to assess the distinctions in intervention distribution between the estimated and the true graph. Intuitively, SHD focuses on differences in graph structure, while SID focuses on differences in causal ordering. Lower SHD and SID values indicate better estimation of the target causal graph by the algorithm. For cases where the output is a MEC, we follow the same approach as Zheng et al. (2018), evaluating them favorably by assuming the undirected edges in the MEC are in the correct direction.

4 CRITICAL EXPERIMENTAL RESULTS AND INSIGHTS

In this section, we first present the experimental results of the misspecified datasets generated according to Section 3.1.1, comparing them with the findings from the vanilla scenario to draw conclusions. Finally, we provide a more in-depth discussion on the performance of CAM (Section 4.1.1) and offer theoretical insights into the performance of differentiable causal discovery (Section 4.1.2). Due to space limitations, the main text focuses on the experimental results for ER-2 graphs, whereas similar conclusions apply to different graph types and graph densities (see Appendix E and J). To visually and concisely present the results, the outcomes of the 10 nodes graph under linear, nonlinear, and MLP settings (Section 4.1.1) are summarized in Figure 1. Results for various nodes in the linear scenarios are displayed in Table 2, and results for nonlinear scenarios are shown in Table 3. [We also consider the real-world Sachs \(Sachs et al., 2005\) dataset \(see Appendix I\), combined misspecified scenarios \(see Appendix F\), vanilla model with non-Gaussian noise \(see Appendix G\) and runtime of the benchmark methods \(see Appendix D\).](#) For each scenario, we generate datasets with 10 different random seeds, each time drawing 2000 samples. [We report the mean and standard deviations of the metrics over 10 trials.](#) To guarantee a fair comparison of various methods, the hyperparameters for each method are determined as the optimal values relative to the specific dataset.

4.1 CURRENT METHODS’ PERFORMANCE IN MISSPECIFIED SCENARIOS

Our experiments show that differentiable causal discovery algorithms almost always achieve optimal or competitive performance in scenarios other than scale variation. In this paper, robustness refers to the ability of the model to perform well in misspecified scenarios, consistent with the understanding of Montagna et al. (2023).

Confounded, measurement error, autoregressive and heterogeneous model. Table 2 and 3 indicate that under the confounded, measurement error, autoregressive and heterogeneous ($P_1 = 0.5, \gamma = 0.1$) scenarios, differentiable causal discovery achieves optimal or competitive performance compared to other methods. For the nonlinear Gaussian process mechanism, although CAM (Bühlmann et al., 2014) demonstrates better performance, the discussion in Section 4.1.1 reveals that CAM still has limitations compared to differentiable causal discovery.

Missing model. We generate missing data that are MCAR with the missingness probability $\alpha = 0.01$. Table 2 and 3 indicate that under MCAR, the result of various algorithms is close to that in the vanilla model. Our experiments show that the performance of differentiable causal discovery under missing data is consistent with traditional methods considered by Tu et al. (2019b;a), including PC and GES.

Mechanism violation. For mechanism violation in the linear setting of Table 2, despite PC and GES’s ability to handle nonlinear mechanisms, we are surprised to observe that linear differentiable

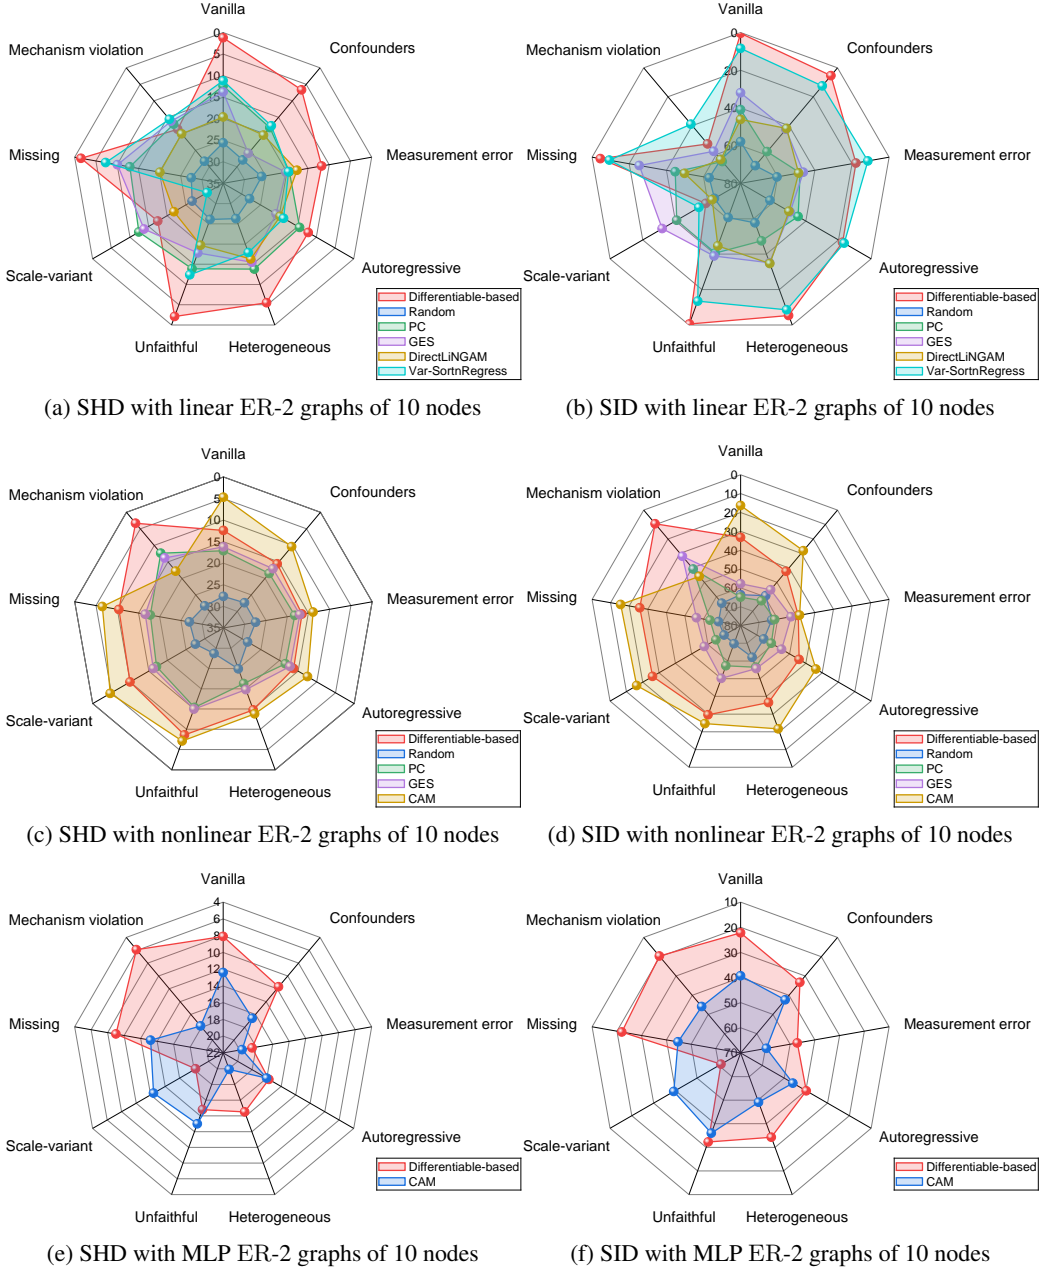


Figure 1: Experimental results under the linear, nonlinear, and mlp settings for both the vanilla scenario and the eight misspecified scenarios. SHD (the lower the better) and SID (the lower the better) are evaluated over 10 trials on the 10 nodes ER-2 graphs. For the differentiable causal discovery method, we present only the optimal results. As the nonlinear settings in Figure 1c and Figure 1d are more favorable to CAM, we conduct a more reasonable evaluation of CAM and differentiable causal discovery under the MLP setting (Section 4.1.1).

causal discovery algorithms achieve competitive performance. In the nonlinear setting of Table 3, we discover that nonlinear differentiable causal discovery algorithms, such as NOTEARS-MLP (Zheng et al., 2020), GraN-DAG (Lachapelle et al., 2019), and DAGMA (Bello et al., 2022), outperform other types of algorithms. From Table 3, we also see that CAM does not perform well under mechanism violation, although it excels in other scenarios. We speculate that this is because the Gaussian process mechanism used in the nonlinear vanilla model aligns well with CAM’s assumptions about SEM. For a more detailed discussion on CAM performance, see Section 4.1.1.

Table 2: Linear Setting, for ER-2 graphs of 10, 20, 50 nodes.

10 nodes	Vanilla model	Latent confounders	Measurement error	Autoregressive	Heterogeneous	Unfaithful	Scale-variant	Missing	Mechanism violation
	SHD ₁	SHD ₁	SHD ₁	SHD ₁	SHD ₁	SHD ₁	SHD ₁	SHD ₁	SHD ₁
Random	25.6±3.1	67.9±8.5	27.9±2.3	62.0±11.1	27.9±3.2	67.7±7.6	26.1±3.7	60.7±12.0	64.0±8.7
PC	12.4±3.1	40.9±13.4	18.1±4.7	48.0±13.1	14.5±2.0	47.5±10.3	13.9±3.2	40.6±10.4	38.3±10.0
GES	13.8±7.8	32.0±13.6	25.9±7.7	42.6±14.0	20.8±5.5	49.7±11.5	15.5±6.1	35.2±12.2	17.8±6.4
DirectLiNGAM	19.6±2.3	46.1±10.6	20.4±5.0	42.0±16.0	19.7±4.2	50.4±8.4	16.3±3.6	34.7±9.9	18.7±4.3
Var-SortnRegress	11.2±3.5	8.4±8.5	17.6±5.8	12.6±9.9	18.8±2.4	16.5±10.6	17.9±3.3	8.6±9.3	12.4±3.1
SortnRegress	20.2±4.8	32.4±14.0	25.7±4.1	37.6±13.0	25.6±4.9	39.2±16.0	26.0±5.4	37.0±14.4	29.0±4.8
NOTEARS	1.5±1.6	1.8±4.2	8.5±3.9	1.5±2.0	19.6±8.0	5.5±2.7	5.4±5.1	2.7±3.4	3.1±5.1
GOLEM	1.4±1.4	0.4±1.2	6.7±2.8	14.2±9.8	17.8±2.5	43.1±13.3	16.6±4.0	34.9±16.9	6.5±4.5
NoCurl	2.0±1.8	5.1±5.8	9.1±4.3	5.4±3.9	11.8±1.8	17.0±8.4	14.8±2.5	17.5±10.7	6.9±7.9
DAGMA	1.2±1.2	3.3±3.3	8.4±3.9	8.8±7.7	12.6±2.5	18.5±8.6	12.3±3.6	28.4±15.3	5.5±2.3
20 nodes	SHD ₁	SHD ₁	SHD ₁	SHD ₁	SHD ₁	SHD ₁	SHD ₁	SHD ₁	SHD ₁
Random	107.9±7.0	253.2±26.3	105.2±8.0	239.8±22.1	109.1±8.0	254.9±18.8	107.8±5.7	247.6±26.1	106.7±5.3
PC	31.6±6.5	108.5±27.6	92.5±4.3	213.0±28.5	48.0±7.8	201.3±30.4	33.5±5.7	181.3±26.0	29.0±7.7
GES	34.3±2.6	104.5±11.1	133.5±12.5	153.7±32.0	57.6±8.2	203.7±40.1	49.4±18.3	114.6±29.4	37.8±17.0
DirectLiNGAM	55.7±9.1	166.4±10.0	72.5±9.7	152.6±28.2	40.0±6.7	132.3±22.5	55.0±9.0	166.7±17.4	50.5±11.2
Var-SortnRegress	39.0±16.3	26.8±9.1	117.3±6.0	34.0±12.2	74.8±1.7	36.6±16.4	74.6±13.0	27.9±14.5	45.3±9.4
SortnRegress	88.7±20.8	117.3±38.6	136.6±10.3	96.7±9.1	89.4±15.1	156.6±42.3	106.9±15.4	140.8±45.6	107.7±13.9
NOTEARS	6.5±4.1	23.4±17.0	26.3±7.5	108.7±4.0	38.8±8.0	114.8±24.0	39.5±8.2	102.1±29.5	14.8±6.1
GOLEM	4.4±2.2	18.0±13.5	24.5±8.2	98.0±36.9	39.5±1.1	212.7±24.2	39.6±3.5	203.2±21.2	14.5±6.6
NoCurl	8.0±6.0	13.6±8.0	60.1±10.1	39.8±13.1	38.3±8.9	69.7±27.8	60.0±9.8	59.1±16.2	23.7±8.0
DAGMA	5.4±3.9	14.2±10.3	20.7±4.6	111.9±23.3	34.2±7.0	115.7±22.0	37.7±7.9	148.6±35.4	12.7±5.7
50 nodes	SHD ₁	SHD ₁	SHD ₁	SHD ₁	SHD ₁	SHD ₁	SHD ₁	SHD ₁	SHD ₁
Random	637.2±24.3	1393.2±102.2	652.6±21.5	1416.5±109.4	642.9±16.1	1436.3±96.8	635.5±15.6	1431.1±147.4	641.7±17.0
DirectLiNGAM	146.4±27.3	781.7±103.0	300.9±18.0	884.4±175.3	104.4±13.1	77.6±122.6	178.8±33.9	683.8±126.1	130.2±16.8
Var-SortnRegress	100.0±16.3	113.1±60.8	920.0±16.1	145.7±19.3	20.6±22.1	208.3±59.2	427.3±36.8	207.3±83.8	309.5±48.3
SortnRegress	388.3±72.4	555.8±142.2	946.6±22.1	401.3±104.6	346.0±138.3	903.4±138.6	552.5±56.8	853.3±215.0	549.4±69.8
NOTEARS	15.2±6.2	71.4±35.5	46.5±1.4	266.5±90.3	81.0±10.2	619.0±18.1	122.4±32.7	944.2±144.6	23.5±3.4
GOLEM	11.8±4.7	58.9±38.1	52.7±1.6	49.4±8.6	97.5±1.0	102.8±14.1	99.2±4.0	102.8±14.1	16.4±3.7
NoCurl	22.8±12.1	60.9±14.0	312.1±30.5	199.8±105.7	109.1±15.7	402.7±76.7	294.3±54.6	287.8±166.4	74.4±15.1
DAGMA	12.4±4.3	47.5±38.7	42.7±10.1	284.2±107.8	81.2±9.4	635.0±172.0	103.0±12.3	932.1±91.9	24.4±7.6

Table 3: Nonlinear Setting, for ER-2 graphs of 10, 20, 50 nodes.

10 nodes	Vanilla model	Latent confounders	Measurement error	Autoregressive	Heterogeneous	Unfaithful	Scale-variant	Missing	Mechanism violation
	SHD ₁	SHD ₁	SHD ₁	SHD ₁	SHD ₁	SHD ₁	SHD ₁	SHD ₁	SHD ₁
Random	27.7±3.2	63.6±11.2	27.4±2.5	59.3±9.4	27.4±3.6	63.2±7.7	28.5±2.3	65.8±8.6	24.9±3.2
PC	17.1±2.3	64.9±10.9	62.8±0.1	91.2±1.1	18.4±1.3	61.2±9.7	21.7±3.2	58.5±10.4	15.4±1.1
GES	16.2±2.2	57.8±10.7	17.1±2.1	55.2±15.5	17.0±1.0	52.8±9.1	19.8±2.5	55.7±8.4	15.0±4.5
CAM	4.7±1.9	16.3±8.5	15.8±2.8	28.2±6.3	13.9±1.9	48.2±7.4	12.5±1.0	33.9±16.3	13.8±2.9
NOTEARS-MLP	12.4±3.2	36.3±7.1	17.0±1.7	49.2±8.6	16.5±0.8	40.9±4.8	17.0±1.7	47.7±11.0	11.4±1.1
GOLEM	12.7±2.4	33.2±10.6	15.6±2.1	42.4±8.8	20.0±1.1	63.8±11.3	16.2±3.2	44.2±10.0	11.2±3.1
DAGMA	13.5±2.0	40.7±8.1	17.1±1.0	62.0±13.3	17.3±1.3	54.9±8.3	19.0±2.0	56.6±10.5	16.3±3.3
20 nodes	SHD ₁	SHD ₁	SHD ₁	SHD ₁	SHD ₁	SHD ₁	SHD ₁	SHD ₁	SHD ₁
Random	103.7±7.8	251.0±26.8	102.7±5.0	237.5±33.7	103.5±5.2	230.9±25.5	102.7±5.2	249.0±31.2	108.1±7.1
PC	34.5±3.4	226.3±26.7	77.8±1.8	202.2±22.8	36.2±2.0	213.2±33.9	38.8±2.8	226.8±35.4	49.9±6.3
GES	32.6±3.3	190.1±42.6	77.5±2.5	199.2±20.2	34.9±2.4	191.0±41.3	36.9±2.3	191.0±28.1	51.0±5.4
CAM	15.0±2.4	70.8±24.7	34.5±2.8	166.4±24.5	31.0±2.1	178.9±32.0	27.9±4.6	125.4±28.5	16.5±2.2
NOTEARS-MLP	27.8±4.0	133.5±26.4	37.8±1.5	204.4±24.8	33.0±2.4	174.1±33.6	39.9±0.3	215.2±26.4	40.0±0.0
GOLEM	20.8±4.2	148.5±24.2	48.1±2.3	218.2±24.4	37.2±1.0	218.2±24.4	37.2±1.0	218.2±24.4	37.2±1.0
DAGMA	27.3±4.4	141.9±36.4	36.2±3.3	175.0±21.9	35.1±2.1	196.0±31.5	39.1±1.3	215.3±27.5	40.0±0.0
50 nodes	SHD ₁	SHD ₁	SHD ₁	SHD ₁	SHD ₁	SHD ₁	SHD ₁	SHD ₁	SHD ₁
Random	646.2±13.6	1455.0±112.9	633.9±11.3	1427.1±145.2	651.7±18.8	1428.8±149.6	626.4±21.2	1412.3±115.2	641.6±18.3
DirectLiNGAM	89.1±3.2	1236.8±191.1	131.2±1.2	1238.3±87.9	105.3±6.7	1170.7±145.2	105.3±6.7	1211.6±15.5	83.0±4.7
Var-SortnRegress	88.6±9.2	100.6±98.8	143.7±4.8	134.1±30.9	84.4±2.0	83.6±105.2	120.4±9.0	114.2±14.2	25.8±10.3
CAM	41.4±4.2	460.0±180.1	75.6±4.2	843.1±84.5	80.6±3.5	864.8±190.8	89.1±11.6	729.9±127.3	127.4±11.1
NOTEARS-MLP	71.1±5.9	73.4±4.8	80.2±9.3	936.4±97.2	87.2±1.0	911.1±108.9	89.1±1.8	907.4±112.3	69.5±4.1
GOLEM	86.0±3.8	832.5±115.8	92.4±7.3	903.8±94.5	101.0±0.0	1041.8±8.5	94.8±1.2	964.7±180.7	96.0±2.3
DAGMA	69.9±5.0	737.2±99.2	78.6±7.3	873.4±82.3	86.6±3.1	924.5±111.9	100.0±0.0	1039.4±88.3	55.8±7.6

Scale-variant model. In Table 2, we observe that the results of linear differentiable causal discovery algorithms, such as NOTEARS (Zheng et al., 2018), GOLEM (Ng et al., 2020), NoCurl (Yu et al., 2021), and DAGMA, significantly decline under scale-variant data, performing worse than PC and GES, which is consistent with the observations of Reisach et al. (2021). For nonlinear differentiable methods, performance under scale-variant data has not been explored in previous research. Table 3 and 4 indicate that nonlinear differentiable methods also show performance degradation under scale variation scenarios, and their results almost always lower than CAM. However, unlike the linear scenarios, the result of nonlinear differentiable algorithms is almost always superior to PC and GES.

Unfaithful model. In the linear setting of Table 2, we see a significant performance drop for Var-SortnRegress (Reisach et al., 2021) and R^2 -SortnRegress (Reisach et al., 2023) under unfaithful distributions. The explanation for this is that for each triplet $X_i \rightarrow X_j \leftarrow X_k \leftarrow X_i$ in the graph, after the causal direct effect of $X_i \rightarrow X_k$ cancels out, the variance of node X_k changes significantly. This reduces the Var-SortnStability, further leading to a performance decline in the two SortnRegress algorithms and linear differentiable methods. In the nonlinear setting of Table 3, the SHD of various algorithms generally decline to some extent under unfaithful path cancellations. This is consistent with the experimental results of Montagna et al. (2023), which indicate that the cancellation of causal effects in unfaithful nonlinear scenarios makes structural inference of sparse graphs easier.

4.1.1 DISCUSSION ON CAM PERFORMANCE

Motivations. The nonlinear vanilla model follows a Gaussian process that is consistent with the assumptions of the CAM. To provide a fair benchmark for CAM, we consider the nonlinear vanilla model following different functional mechanism. We compare CAM with the representative differentiable causal discovery method: NOTEARS-MLP.

Simulations. We simulate ER-2 graphs based on the number of nodes $d \in \{10, 20, 50\}$. Following the data generation mechanisms of Zheng et al. (2020), we consider f_i in nonlinear vanilla model (Section 3.1) is modified to be parameterized by a neural network with one hidden layer of size 100.

Results. From Table 4, we observe that NOTEARS-MLP achieves better performance under almost all model assumption violations. Considering that the functional mechanisms of data in real-world

scenarios are usually unknown, we believe that differentiable causal discovery has a significant advantage over CAM in all types of assumption violation scenarios except for scale variation.

4.1.2 THEORY ON DIFFERENTIABLE CAUSAL DISCOVERY IN MISSPECIFIED SCENARIOS

We analyze the performance of linear differentiable causal discovery under measurement error, unfaithful and missing scenarios by introducing the theories (Theorem 7 and Theorem 9) from Loh & Bühlmann (2014). Theorem 7 in Loh & Bühlmann (2014) states that for a linear model with equal noise variance, minimizing the least squares score will return the true DAG in the large sample limit. For a linear model with non-equal noise variances, we define the noise ratio

$$r = \frac{\max(\sigma_1^2, \dots, \sigma_d^2)}{\min(\sigma_1^2, \dots, \sigma_d^2)}. \quad (6)$$

Theorem 9 in Loh & Bühlmann (2014) states that if $r < 1 + \frac{\xi}{d}$, where $\xi > 0$ is the gap between the score of the true DAG and the next best DAG, minimizing the least squares score will return the true DAG in the large sample limit.

Measurement error. In measurement error model considered by Montagna et al. (2023), the observed variables are:

$$\tilde{X}_i = X_i + \epsilon_i, \forall i = 1, \dots, d, \quad (7)$$

where $X_i = f_i(X_{pa(X_i)}) + U_i$, f_i is a linear mechanism, $U_i \sim N(0, 1)$, $\epsilon_i \sim N(0, \delta * \text{Var}(X_i))$ with $\delta \in \{0.2, 0.4, 0.6, 0.8\}$. In the vanilla model, $r = 1$. However, in the measurement error model, the noise ratio becomes

$$\tilde{r} = \frac{\max(1 + \delta * \text{Var}(X_i), \dots, 1 + \delta * \text{Var}(X_d))}{\min(1 + \delta * \text{Var}(X_i), \dots, 1 + \delta * \text{Var}(X_d))}. \quad (8)$$

Due to the increasing trend of marginal variances of nodes along the causal direction (Reisach et al., 2021), we infer that $\tilde{r} > r = 1$. In this scenario, there is no guarantee that $\tilde{r} < 1 + \frac{\xi}{d}$ and linear differentiable causal discovery based on least squares cannot guarantee obtaining the true DAG, which can explain their performance decline in Table 2.

Unfaithful model. In unfaithful model considered by Montagna et al. (2023), for each triplet $X_i \rightarrow X_j \rightarrow X_k \leftarrow X_i$ in the graph, the causal mechanisms are adjusted such that the direct effect of X_i on X_k cancels out. To illustrate the change in noise ratio after path cancellation, we consider a DAG \mathcal{G} with variable set $V(\mathcal{G}) = \{X_1, X_2, X_3\}$ and edge set $E(\mathcal{G}) = \{X_1 \rightarrow X_2, X_2 \rightarrow X_3, X_1 \rightarrow X_3\}$. The structural equations is defined as:

$$\begin{aligned} X_1 &= U_1, \\ X_2 &= f_1(X_1) + U_2, \\ X_3 &= f_1(X_1) - X_2 + U_3, \end{aligned} \quad (9)$$

where f_1 is a linear mechanism. After the direct causal effect of $X_1 \rightarrow X_3$ cancels out, the noise term of X_3 is $U_3 - U_2$ with the distribution of $N(0, 2)$. Similarly, in the unfaithful datasets with nodes $d \in \{10, 20, 50\}$ considered by our experiments, for each triplet $X_i \rightarrow X_j \rightarrow X_k \leftarrow X_i$ in the graph, once unfaithful path cancellation occurs, the noise term of X_k is $U_k - U_j$ with the distribution of $N(0, 2)$. In this case, the noise ratio becomes $r' = 2 > r = 1$ (vanilla model). Due to the increasing of the noise ratio, there is no guarantee that $r' < 1 + \frac{\xi'}{d}$ and linear differentiable causal discovery based on least squares cannot guarantee obtaining the true DAG, which can also explain their performance decline in Table 2.

Missing model. Under the MCAR case, we deleted the rows with missing values and regenerated the data under the i.i.d. assumption to ensure the unchanged sample size. The noise ratio $r = 1$ remains constant before and after data imputation. This explains the superior performance of differentiable causal discovery methods observed in Table 2.

Table 4: MLP Setting, for ER-2 graphs of 10, 20, 50 nodes.

10 nodes	Vanilla model		Latent confounders		Measurement error		Autoregressive		Heterogeneous		Unfaithful		Scale-variant		Missing		Mechanism violation	
	SHD ₁	SD ₁	SHD ₁	SD ₁	SHD ₁	SD ₁	SHD ₁	SD ₁	SHD ₁	SD ₁	SHD ₁	SD ₁	SHD ₁	SD ₁	SHD ₁	SD ₁	SHD ₁	SD ₁
CAM	12.4±3.6	39.3±16.5	16.6±4.2	42.4±17.7	19.7±4.7	39.6±12.0	16.0±3.4	46.0±16.6	19.9±3.6	49.1±7.2	13.0±3.2	36.0±14.0	12.4±3.6	39.3±16.5	13.2±3.6	44.7±16.7	17.8±4.4	45.9±17.0
NOTEARS-MLP	8.1±2.7	22.2±10.6	11.7±5.5	33.3±17.0	18.5±3.7	47.1±11.9	15.7±4.3	39.8±9.6	14.5±2.3	34.3±10.4	14.8±3.9	32.5±12.8	18.2±3.6	61.0±11.0	9.0±3.3	22.0±10.0	5.8±2.5	19.7±8.7
20 nodes	SHD ₁	SD ₁	SHD ₁	SD ₁	SHD ₁	SD ₁	SHD ₁	SD ₁	SHD ₁	SD ₁	SHD ₁	SD ₁	SHD ₁	SD ₁	SHD ₁	SD ₁	SHD ₁	SD ₁
	SHD ₁	SD ₁	SHD ₁	SD ₁	SHD ₁	SD ₁	SHD ₁	SD ₁	SHD ₁	SD ₁	SHD ₁	SD ₁	SHD ₁	SD ₁	SHD ₁	SD ₁	SHD ₁	SD ₁
CAM	25.2±8.3	151.1±62.3	32.5±7.6	189.0±34.6	50.3±7.3	230.8±30.1	36.4±9.3	164.0±55.5	48.3±8.0	205.5±38.1	28.5±5.9	149.3±47.7	25.2±8.3	151.1±62.3	26.3±8.6	158.1±63.8	41.2±8.1	213.4±47.0
NOTEARS-MLP	15.6±3.6	74.3±17.2	25.8±3.8	160.6±39.6	37.3±2.1	197.4±25.0	36.9±5.0	187.8±23.3	29.4±6.0	98.6±26.8	28.4±4.2	105.8±27.7	34.2±3.5	209.1±21.2	15.9±3.7	74.9±20.2	19.2±4.9	87.2±17.7
50 nodes	SHD ₁	SD ₁	SHD ₁	SD ₁	SHD ₁	SD ₁	SHD ₁	SD ₁	SHD ₁	SD ₁	SHD ₁	SD ₁	SHD ₁	SD ₁	SHD ₁	SD ₁	SHD ₁	SD ₁
	SHD ₁	SD ₁	SHD ₁	SD ₁	SHD ₁	SD ₁	SHD ₁	SD ₁	SHD ₁	SD ₁	SHD ₁	SD ₁	SHD ₁	SD ₁	SHD ₁	SD ₁	SHD ₁	SD ₁
CAM	59.7±10.9	837.5±199.7	83.4±12.1	942.5±164.2	128.8±10.3	1220.2±124.7	101.3±17.7	855.8±172.9	126.7±18.0	1008.1±135.9	65.2±9.3	717.1±103.0	59.7±10.9	837.5±199.7	61.4±11.4	849.0±158.3	77.8±11.8	877.2±165.6
NOTEARS-MLP	41.5±11.3	378.3±104.5	67.4±9.7	788.6±148.5	93.7±2.6	958.2±115.7	113.7±21.6	1003.5±131.0	74.1±4.0	826.4±106.6	67.2±4.5	724.8±67.8	87.4±6.6	1052.0±107.0	41.0±9.9	392.6±90.5	44.7±5.9	543.8±72.1

4.2 SUMMARY AND IMPLICATIONS FOR PRACTICE

In Appendix H, we summarize the results of the most competitive methods under misspecified scenarios. Differentiable causal discovery methods demonstrate optimal or competitive performance in scenarios other than scale variation. Notably, the recent work by Deng et al. (2024) shows that for linear differentiable methods, scale invariance can be achieved by appropriately choosing the loss function. This further reinforces our conclusion regarding the robustness of differentiable methods. In our benchmarks, the results in Table 4, Table 14 and Table 18 indicate that the performance of nonlinear differentiable methods under scale variation remains challenging and warrants further investigation. In practice, the misspecified scenarios are inevitably encountered, making the robustness of algorithms critically important. Based on the summarized results on eight misspecified synthetic datasets (see Appendix H), real-world data results (see Appendix I) and runtime results of benchmark methods (see Appendix D), we observe that differentiable causal discovery methods have the potential to achieve optimal or competitive performance on real-world data with an almost negligible time cost. The fast and robust characteristics of differentiable methods enable them to better address the challenges of applying causal discovery algorithms to real-world data, demonstrating their practical implementation potential.

5 CONCLUSION

This work assesses the efficacy of twelve preeminent causal discovery methods across eight scenarios involving violations of model assumptions. These methods encompass approaches grounded in independence constraints, scoring criteria, functional causal models, and differentiable causal discovery. Our experimental results show that differentiable causal discovery methods exhibit remarkable resilience in scenarios of model assumption violations, except for scale variation. It is not our intention to assert that differentiable causal discovery will achieve optimal performance across all circumstances, rather, we aim to underscore its substantial potential within the benchmarks we have evaluated, thereby emphasizing the necessity for further exploration in this direction. In future work, more causal discovery methods for semi-synthetic data and more real-world scenarios will be explored. Finally, our study confines itself to non-temporal causal discovery algorithms. Equally crucial is the conduct of benchmark assessments for causal discovery in time series and event sequences under model assumption violations.

REFERENCES

- Hiroto Akaike. Information theory and an extension of the maximum likelihood principle. In *Selected papers of hirotugu akaike*, pp. 199–213. Springer, 1998.
- Kevin Bello, Bryon Aragam, and Pradeep Ravikumar. Dagma: Learning dags via m-matrices and a log-determinant acyclicity characterization. In *Advances in Neural Information Processing Systems*, volume 35, pp. 8226–8239, 2022.
- Ulrik Brandes, Marco Gaertler, and Dorothea Wagner. Experiments on graph clustering algorithms. In *European symposium on algorithms*, pp. 568–579. Springer, 2003.
- Philippe Brouillard, Sébastien Lachapelle, Alexandre Lacoste, Simon Lacoste-Julien, and Alexandre Drouin. Differentiable causal discovery from interventional data. In *Advances in Neural Information Processing Systems*, volume 33, pp. 21865–21877, 2020.
- Peter Bühlmann, Jonas Peters, and Jan Ernest. Cam: Causal additive models, high-dimensional order search and penalized regression. *The Annals of Statistics*, 42(6):2526–2556, 2014. doi: 10.1214/14-AOS1260.

- Ruichu Cai, Feng Xie, Clark Glymour, Zhifeng Hao, and Kun Zhang. Triad constraints for learning causal structure of latent variables. In *Advances in Neural Information Processing Systems*, volume 32, 2019.
- Ruichu Cai, Jincheng Ye, Jie Qiao, Huiyuan Fu, and Zhifeng Hao. Fom: Fourth-order moment based causal direction identification on the heteroscedastic data. *Neural Networks*, 124:193–201, 2020.
- Ruichu Cai, Zhiyi Huang, Wei Chen, Zhifeng Hao, and Kun Zhang. Causal discovery with latent confounders based on higher-order cumulants. In *International Conference on Machine Learning*, pp. 3380–3407. PMLR, 2023.
- David Maxwell Chickering. Optimal structure identification with greedy search. *Journal of Machine Learning Research*, 3(Nov):507–554, 2002.
- Haoyue Dai, Peter Spirtes, and Kun Zhang. Independence testing-based approach to causal discovery under measurement error and linear non-gaussian models. *Advances in Neural Information Processing Systems*, 35:27524–27536, 2022.
- Chang Deng, Kevin Bello, Pradeep Ravikumar, and Bryon Aragam. Likelihood-based differentiable structure learning. *arXiv preprint arXiv:2410.06163*, 2024.
- José A Gámez, Juan L Mateo, and José M Puerta. Learning bayesian networks by hill climbing: efficient methods based on progressive restriction of the neighborhood. *Data Mining and Knowledge Discovery*, 22:106–148, 2011.
- Erdun Gao, Junjia Chen, Li Shen, Tongliang Liu, Mingming Gong, and Howard Bondell. Feddag: Federated dag structure learning. *arXiv preprint arXiv:2112.03555*, 2021.
- Erdun Gao, Ignavier Ng, Mingming Gong, Li Shen, Wei Huang, Tongliang Liu, Kun Zhang, and Howard Bondell. Missdag: Causal discovery in the presence of missing data with continuous additive noise models. In *Advances in Neural Information Processing Systems*, volume 35, pp. 5024–5038, 2022.
- AmirEmad Ghassami, Saber Salehkaleybar, Negar Kiyavash, and Kun Zhang. Learning causal structures using regression invariance. In *Advances in Neural Information Processing Systems*, volume 30, 2017.
- AmirEmad Ghassami, Negar Kiyavash, Biwei Huang, and Kun Zhang. Multi-domain causal structure learning in linear systems. In *Advances in Neural Information Processing Systems*, volume 31, 2018.
- Clark Glymour, Kun Zhang, and Peter Spirtes. Review of causal discovery methods based on graphical models. *Frontiers in Genetics*, 10:524, 2019.
- Alain Hauser and Peter Bühlmann. Characterization and greedy learning of interventional markov equivalence classes of directed acyclic graphs. *Journal of Machine Learning Research*, 13(1): 2409–2464, 2012.
- Yue He, Peng Cui, Zheyang Shen, Renzhe Xu, Furui Liu, and Yong Jiang. Daring: Differentiable causal discovery with residual independence. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pp. 596–605, 2021.
- David Heckerman, Dan Geiger, and David M Chickering. Learning bayesian networks: The combination of knowledge and statistical data. *Machine learning*, 20:197–243, 1995.
- Christina Heinze-Deml, Marloes H Maathuis, and Nicolai Meinshausen. Causal structure learning. *Annual Review of Statistics and Its Application*, 5(1):371–391, 2018.
- Patrik Hoyer, Dominik Janzing, Joris M Mooij, Jonas Peters, and Bernhard Schölkopf. Nonlinear causal discovery with additive noise models. In *Advances in Neural Information Processing Systems*, volume 21, 2008.
- Biwei Huang, Kun Zhang, Jiji Zhang, Joseph Ramsey, Ruben Sanchez-Romero, Clark Glymour, and Bernhard Schölkopf. Causal discovery from heterogeneous/nonstationary data. *Journal of Machine Learning Research*, 21(89):1–53, 2020.

- Diviyan Kalainathan and Olivier Goudet. Causal discovery toolbox: Uncover causal relationships in python. *arXiv preprint arXiv:1903.02278*, 2019.
- Diviyan Kalainathan, Olivier Goudet, Isabelle Guyon, David Lopez-Paz, and Michèle Sebag. Structural agnostic modeling: Adversarial learning of causal graphs. *Journal of Machine Learning Research*, 23(219):1–62, 2022.
- Mehmet Kayaalp and Gregory F Cooper. A bayesian network scoring metric that is based on globally uniform parameter priors. *arXiv preprint arXiv:1301.0576*, 2012.
- Lingjing Kong, Biwei Huang, Feng Xie, Eric Xing, Yuejie Chi, and Kun Zhang. Identification of nonlinear latent hierarchical models. In *Advances in Neural Information Processing Systems*, volume 36, pp. 2010–2032, 2023.
- Trent Kyono, Yao Zhang, and Mihaela van der Schaar. Castle: Regularization via auxiliary causal graph discovery. In *Advances in Neural Information Processing Systems*, volume 33, pp. 1501–1512, 2020.
- Sébastien Lachapelle, Philippe Brouillard, Tristan Deleu, and Simon Lacoste-Julien. Gradient-based neural dag learning. *arXiv preprint arXiv:1906.02226*, 2019.
- Fangfu Liu, Wenchang Ma, An Zhang, Xiang Wang, Yueqi Duan, and Tat-Seng Chua. Discovering dynamic causal space for dag structure learning. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pp. 1429–1440, 2023.
- Po-Ling Loh and Peter Bühlmann. High-dimensional learning of linear causal networks via inverse covariance estimation. *Journal of Machine Learning Research*, 15(1):3065–3105, 2014.
- Francesco Montagna, Atalanti Mastakouri, Elias Eulig, Nicoletta Noceti, Lorenzo Rosasco, Dominik Janzing, Bryon Aragam, and Francesco Locatello. Assumption violations in causal discovery and the robustness of score matching. In *Advances in Neural Information Processing Systems*, volume 36, 2023.
- Joris M Mooij, Jonas Peters, Dominik Janzing, Jakob Zscheischler, and Bernhard Schölkopf. Distinguishing cause from effect using observational data: methods and benchmarks. *Journal of Machine Learning Research*, 17(32):1–102, 2016.
- Joris M Mooij, Sara Magliacane, and Tom Claassen. Joint causal inference from multiple contexts. *Journal of Machine Learning Research*, 21(99):1–108, 2020.
- Andrew A Neath and Joseph E Cavanaugh. The bayesian information criterion: background, derivation, and applications. *Wiley Interdisciplinary Reviews: Computational Statistics*, 4(2): 199–203, 2012.
- AS Nemirovsky. Optimization ii. numerical methods for nonlinear continuous optimization. 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, AmirEmad Ghassami, and Kun Zhang. On the role of sparsity and dag constraints for learning linear dags. In *Advances in Neural Information Processing Systems*, volume 33, pp. 17943–17954, 2020.
- 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.
- Roxana Pamfil, Nisara Sriwattanaworachai, Shaan Desai, Philip Pilgerstorfer, Konstantinos Georgatzis, Paul Beaumont, and Bryon Aragam. Dynotears: Structure learning from time-series data. In *International Conference on Artificial Intelligence and Statistics*, pp. 1595–1605. PMLR, 2020.
- Judea Pearl. *Causality*. Cambridge university press, 2009.
- Judea Pearl and Dana Mackenzie. *The book of why: the new science of cause and effect*. Basic books, 2018.

- Jonas Peters and Peter Bühlmann. Identifiability of gaussian structural equation models with equal error variances. *Biometrika*, 101(1):219–228, 2014.
- Jonas Peters and Peter Bühlmann. Structural intervention distance for evaluating causal graphs. *Neural Computation*, 27(3):771–799, 2015.
- Jonas Peters, Joris M. Mooij, Dominik Janzing, and Bernhard Schölkopf. Causal discovery with continuous additive noise models. *Journal of Machine Learning Research*, 15(58):2009–2053, 2014. URL <http://jmlr.org/papers/v15/peters14a.html>.
- Jonas Peters, Dominik Janzing, and Bernhard Schölkopf. *Elements of causal inference: foundations and learning algorithms*. The MIT Press, 2017.
- Alexander Reisach, Christof Seiler, and Sebastian Weichwald. Beware of the simulated dag! causal discovery benchmarks may be easy to game. In *Advances in Neural Information Processing Systems*, volume 34, pp. 27772–27784, 2021.
- Alexander Reisach, Myriam Tami, Christof Seiler, Antoine Chambaz, and Sebastian Weichwald. A scale-invariant sorting criterion to find a causal order in additive noise models. In *Advances in Neural Information Processing Systems*, volume 36, 2023.
- Robert W Robinson. Counting labeled acyclic digraphs. *New directions in the theory of graphs*, pp. 239–273, 1973.
- 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.
- Saber Salehkaleybar, AmirEmad Ghassami, Negar Kiyavash, and Kun Zhang. Learning linear non-gaussian causal models in the presence of latent variables. *Journal of Machine Learning Research*, 21(39):1–24, 2020.
- Shohei Shimizu, Patrik O Hoyer, Aapo Hyvärinen, Antti Kerminen, and Michael Jordan. A linear non-gaussian acyclic model for causal discovery. *Journal of Machine Learning Research*, 7(10), 2006.
- Shohei Shimizu, Takanori Inazumi, Yasuhiro Sogawa, Aapo Hyvarinen, Yoshinobu Kawahara, Takashi Washio, Patrik O Hoyer, Kenneth Bollen, and Patrik Hoyer. Directlingam: A direct method for learning a linear non-gaussian structural equation model. *Journal of Machine Learning Research*, 12(Apr):1225–1248, 2011.
- Karamjit Singh, Garima Gupta, Vartika Tewari, and Gautam Shroff. Comparative benchmarking of causal discovery techniques. *arXiv preprint arXiv:1708.06246*, 2017.
- Peter Spirtes. Introduction to causal inference. *Journal of Machine Learning Research*, 11(5), 2010.
- Peter Spirtes and Clark Glymour. An algorithm for fast recovery of sparse causal graphs. *Social science computer review*, 9(1):62–72, 1991.
- Peter Spirtes, Christopher Meek, and Thomas Richardson. Causal inference in the presence of latent variables and selection bias. In *Proceedings of the Eleventh Conference on Uncertainty in Artificial Intelligence*, pp. 499–506, 1995.
- Peter Spirtes, Clark Glymour, and Richard Scheines. *Causation, prediction, and search*. MIT press, 2001.
- Ioannis Tsamardinos, Laura E Brown, and Constantin F Aliferis. The max-min hill-climbing bayesian network structure learning algorithm. *Machine learning*, 65:31–78, 2006.
- Ruibin Tu, Cheng Zhang, Paul Ackermann, Karthika Mohan, Hedvig Kjellström, and Kun Zhang. Causal discovery in the presence of missing data. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pp. 1762–1770. PMLR, 2019a.

- Ruibo Tu, Kun Zhang, Bo Bertilson, Hedvig Kjellstrom, and Cheng Zhang. Neuropathic pain diagnosis simulator for causal discovery algorithm evaluation. In *Advances in Neural Information Processing Systems*, volume 32, 2019b.
- Matthew J Vowels, Necati Cihan Camgoz, and Richard Bowden. D’ya like dags? a survey on structure learning and causal discovery. *ACM Computing Surveys*, 55(4):1–36, 2022.
- Yu Wang, An Zhang, Xiang Wang, Yancheng Yuan, Xiangnan He, and Tat-Seng Chua. Differentiable invariant causal discovery. *arXiv preprint arXiv:2205.15638*, 2022.
- Dennis Wei, Tian Gao, and Yue Yu. Dags with no fears: A closer look at continuous optimization for learning bayesian networks. *Advances in Neural Information Processing Systems*, 33:3895–3906, 2020.
- Feng Xie, Ruichu Cai, Biwei Huang, Clark Glymour, Zhifeng Hao, and Kun Zhang. Generalized independent noise condition for estimating latent variable causal graphs. In *Advances in Neural Information Processing Systems*, volume 33, pp. 14891–14902, 2020.
- Shuai Yang, Kui Yu, Fuyuan Cao, Lin Liu, Hao Wang, and Jiuyong Li. Learning causal representations for robust domain adaptation. *IEEE Transactions on Knowledge and Data Engineering*, 35(3): 2750–2764, 2021.
- 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.
- Yue Yu, Tian Gao, Naiyu Yin, and Qiang Ji. Dags with no curl: An efficient dag structure learning approach. In *International Conference on Machine Learning*, pp. 12156–12166. PMLR, 2021.
- An Zhang, Fangfu Liu, Wenchang Ma, Zhibo Cai, Xiang Wang, and Tat-Seng Chua. Boosting differentiable causal discovery via adaptive sample reweighting. *arXiv preprint arXiv:2303.03187*, 2023.
- Keli Zhang, Shengyu Zhu, Marcus Kalander, Ignavier Ng, Junjian Ye, Zhitang Chen, and Lujia Pan. gcastle: A python toolbox for causal discovery. *arXiv preprint arXiv:2111.15155*, 2021.
- Kun Zhang and Aapo Hyvarinen. On the identifiability of the post-nonlinear causal model. *arXiv preprint arXiv:1205.2599*, 2012.
- Kun Zhang, Mingming Gong, Joseph Ramsey, K. Batmanghelich, Peter Spirtes, and Clark Glymour. Causal discovery with linear non-gaussian models under measurement error: Structural identifiability results. In *Conference on Uncertainty in Artificial Intelligence*, 2018. URL <https://api.semanticscholar.org/CorpusID:54058643>.
- Muhan Zhang, Shali Jiang, Zhicheng Cui, Roman Garnett, and Yixin Chen. D-vae: A variational autoencoder for directed acyclic graphs. In *Advances in Neural Information Processing Systems*, volume 32, 2019.
- Xun Zheng, Bryon Aragam, Pradeep K Ravikumar, and Eric P Xing. Dags with no tears: Continuous optimization for structure learning. In *Advances in Neural Information Processing Systems*, volume 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.
- Yujia Zheng, Biwei Huang, Wei Chen, Joseph Ramsey, Mingming Gong, Ruichu Cai, Shohei Shimizu, Peter Spirtes, and Kun Zhang. Causal-learn: Causal discovery in python. *Journal of Machine Learning Research*, 25(60):1–8, 2024.
- Rong Zhu, Andreas Pfadler, Ziniu Wu, Yuxing Han, Xiaoke Yang, Feng Ye, Zhenping Qian, Jingren Zhou, and Bin Cui. Efficient and scalable structure learning for bayesian networks: Algorithms and applications. In *2021 IEEE 37th International Conference on Data Engineering (ICDE)*, pp. 2613–2624. IEEE, 2021.
- Shengyu Zhu, Ignavier Ng, and Zhitang Chen. Causal discovery with reinforcement learning. *arXiv preprint arXiv:1906.04477*, 2019.

CONTENTS

1	Introduction	1
2	Background	3
2.1	Task formulation	3
2.2	Structure identifiability	3
2.3	Differentiable score-based causal discovery	3
3	Experimental design	4
3.1	Synthetic datasets	4
3.1.1	Model assumption violation scenarios	5
3.1.2	Data generation	5
3.2	Methods	5
3.3	Hyperparameter settings	6
3.4	Evaluation metrics	6
4	Critical experimental results and insights	6
4.1	Current methods' performance in misspecified scenarios	6
4.1.1	Discussion on CAM performance	8
4.1.2	Theory on differentiable causal discovery in misspecified scenarios	9
4.2	Summary and implications for practice	10
5	Conclusion	10
A	Causal Assumptions	17
A.1	Causal Markov Property	17
A.2	Faithfulness	17
A.3	Causal Sufficiency	17
A.4	Independent and identically distributed	17
A.5	Equal Noise Variances	17
B	Benchmark methods	18
B.1	PC	18
B.2	GES	18
B.3	DirectLiNGAM	18
B.4	CAM	18
B.5	SortnRegress	18
B.6	NOTEARS	19
B.7	GOLEM	19
B.8	NOTEARS-MLP	20

810	B.9 GraN-DAG	21
811	B.10 NOCURL	21
812	B.11 DAGMA	22
813		
814		
815	C Related work	22
816		
817	D Table results for runtime of the benchmark methods	23
818		
819	E Table results across nodes, graph types, and graph densities	24
820		
821		
822	F Table results for combined misspecified scenarios	26
823		
824	G Table results for non-Gaussian noise	26
825		
826	H Summary of the most competitive methods	27
827		
828	I Table Results on real-world data	28
829		
830	J Figure results across nodes, graph types, and graph densities	29
831		
832		
833		
834		
835		
836		
837		
838		
839		
840		
841		
842		
843		
844		
845		
846		
847		
848		
849		
850		
851		
852		
853		
854		
855		
856		
857		
858		
859		
860		
861		
862		
863		

A CAUSAL ASSUMPTIONS

In this section, we introduce the assumptions frequently used in the causal discovery literature.

A.1 CAUSAL MARKOV PROPERTY

The joint probability distribution $P(X)$ satisfies the global Markov property (Peters et al., 2017) with respect to the DAG \mathcal{G} if

$$X_A \perp\!\!\!\perp_{\mathcal{G}} X_B \mid X_C \Rightarrow X_A \perp\!\!\!\perp_{P(X)} X_B \mid X_C, \quad (10)$$

where X_A , X_B and X_C are the disjoint subsets of $X = (X_1, \dots, X_d) \in \mathbb{R}^d$, $\perp\!\!\!\perp_{\mathcal{G}}$ denotes d -separation in the causal graph \mathcal{G} , and $\perp\!\!\!\perp_{P(X)}$ represents independence in the joint probability distribution $P(X)$.

In the causal graph \mathcal{G} of SCM, each variable is independent of its non-descendant nodes when its parents are known, which is referred to as the local Markov property (Peters et al., 2017). Causal Markov property also implies that the joint probability distribution $P(X)$ can be factorized in the following form:

$$P(X) = \prod_i^d P(X_i \mid pa(X_i)). \quad (11)$$

A.2 FAITHFULNESS

The joint probability distribution $P(X)$ is faithful (Peters et al., 2017) to the DAG \mathcal{G} if

$$X_A \perp\!\!\!\perp_{P(X)} X_B \mid X_C \Rightarrow X_A \perp\!\!\!\perp_{\mathcal{G}} X_B \mid X_C. \quad (12)$$

Faithfulness assumption implies that the conditional independence in the $P(X)$ can be used to infer the graph structure. Constraint-based and traditional score-based causal discovery are typically founded on the faithfulness assumption. In a causal graph, the cancellation of effects along multiple causal paths can lead to a violation of faithfulness assumption.

A.3 CAUSAL SUFFICIENCY

Causal sufficiency assumption is also referred to as no latent confounder. A set of variables X is said to satisfy the causal sufficiency if there is no unobserved common cause variable C that influences more than one variable in X (Spirtes, 2010). This assumption is also frequently considered in causal discovery literature. However, since we cannot always observe all variables in the real world, causal sufficiency assumption is inevitably violated.

A.4 INDEPENDENT AND IDENTICALLY DISTRIBUTED

Non-temporal causal discovery algorithms typically also require the i.i.d. assumption. In the main text, heterogeneous multi-domain data and autoregressive scenarios are two special cases where the i.i.d. assumption is violated. Heterogeneous multi-domain data is closely related to the non-stationary time series data considered in the literature on temporal causal discovery (Huang et al., 2020). Below, we introduce the connection between them. We consider the distribution of X_i changing with domain or time index, where the mechanism for the t -th data point is as follows:

$$X_{i,t} = f_{i,t}(pa(X_{i,t}), \epsilon_{i,t}), \quad (13)$$

where $\epsilon_{i,t}$ is the noise term of $X_{i,t}$. In heterogeneous multi-domain data, t represents the domain index, whereas in non-stationary time series data, t denotes the time index.

A.5 EQUAL NOISE VARIANCES

Ng et al. (2024) observe that the performance of linear differentiable causal discovery methods significantly declines in data with non-equal noise variances. They hypothesize that this may be due to the optimization problem becoming severely non-convex under non-equal noise variances, leading to local optimal solutions. Although differentiable methods do not explicitly assume equal noise variance, the performance decline suggests treating equal noise variance as a causal assumption.

B BENCHMARK METHODS

B.1 PC

PC (Spirtes & Glymour, 1991) algorithm is a representative constraint-based causal discovery method. In the first step, under the faithfulness assumption, the global causal skeleton is determined based on conditional independence tests. In the second step, some edge directions in the skeleton are determined by identifying collider structures. Finally, the remaining edge directions are determined using orientation rules, resulting in the MEC. We use the implementation of the PC algorithm in causal-learn (Zheng et al., 2024) python package, available at <https://github.com/py-why/causal-learn>.

B.2 GES

GES (Chickering, 2002) algorithm is a classical score-based causal discovery method. GES mainly includes two stages. In the first stage, starting from an empty graph, edges are added through greedy equivalence search, and the structures in the equivalence class of the new graph are scored. The graph with the highest score is selected, and the edge-adding process is repeated until the score reaches a local maximum. In the second stage, starting from the graph obtained in the first stage, edges are removed through greedy equivalence search, and the structures in the equivalence class of the new graph are scored. The graph with the highest score is selected, and the edge-removal process is repeated until the score reaches a local maximum. We use the implementation of the GES algorithm in causal-learn (Zheng et al., 2024) python package, available at <https://github.com/py-why/causal-learn>.

B.3 DIRECTLiNGAM

DirectLiNGAM (Shimizu et al., 2011) is a classical linear method based on the functional causal model. To address the issues of slow convergence and large errors in the ICA-based LiNGAM (Shimizu et al., 2006) algorithm, Shimizu et al. (2011) proposed DirectLiNGAM based on the principle of residual independence. Although the solving speed became slower, the accuracy and convergence improved. We use the implementation of DirectLiNGAM algorithm in gCastle (Zhang et al., 2021) python package, available at <https://github.com/huawei-noah/trustworthyAI/tree/master/gcastle>.

B.4 CAM

CAM (Bühlmann et al., 2014) is a method used for high-dimensional additive structural equation models. CAM separates the search for the order of variables from the selection of edges. It performs the variable order search through nonregularized maximum likelihood estimation and uses sparse regression techniques for edge selection. We use the implementation of CAM algorithm in Causal Discovery Toolbox (Kalainathan & Goudet, 2019) python package, available at <https://github.com/FenTechSolutions/CausalDiscoveryToolbox>.

B.5 SORTNREGRESS

SortnRegress algorithm includes Var-SortnRegress (Reisach et al., 2021) and R^2 -SortnRegress (Reisach et al., 2023).

Reisach et al. (2021) emphasized that in the bivariate linear case, causal direction inferred by minimizing mean squared error (MSE) loss is from the variable with smaller variance to the variable with larger variance. They further hypothesized that, in the multivariate case, there is a consistency between the underlying causal direction of the data and the increasing order of the marginal variances of the variables. They provided a general definition of sortability:

$$\mathbf{v}_\tau(X, \mathcal{G}) = \frac{\sum_{i=1}^d \sum_{(s \rightarrow t) \in A(\mathcal{G})^i} \text{incr}(\tau(X, s), \tau(X, t))}{\sum_{i=1}^d \sum_{(s \rightarrow t) \in A(\mathcal{G})^i} 1} \text{ where } \text{incr}(a, b) = \begin{cases} 1 & a < b \\ 1/2 & a = b \\ 0 & a > b \end{cases}, \quad (14)$$

τ represents a function with $\tau(X) \in [0, 1]$, $A(\mathcal{G})$ is the adjacency matrix of \mathcal{G} , $A(\mathcal{G})^i$ is the i -th matrix power, $(s \rightarrow t) \in A(\mathcal{G})^i$ if and only if at least one directed path of length i from X_s to X_t .

In Var-Sortability, $\tau(X, s)$ denotes the variance of X_s . Var-Sortability measures the consistency between the causal structure order and the increasing order of marginal variances of nodes. Intuitively, the greater Var-Sortability of the data, the better performance of methods based on MSE loss. Subsequently, they proposed the Var-SortnRegress algorithm to discover causality using only variance. In the first step, the nodes are ranked according to the increasing order of marginal variances. In the second step, linear and Lasso regression are used for estimation. We use the implementation of Var-SortnRegress algorithm in CausalDisco python package provided by the authors, available at <https://github.com/CausalDisco/CausalDisco>.

Reisach et al. (2021) demonstrated that the performance of linear differentiable causal discovery algorithms is greatly affected by Var-Sortability. Building on previous work, Reisach et al. (2023) pointed out that the coefficient of determination R^2 remains unchanged after scaling the data, and proposed the R^2 -SortnRegress algorithm, which achieves better performance on scale-variant data. The definition of R^2 :

$$R^2 = 1 - \frac{\text{Var}(X_t - \mathbb{E}[X_t | X_{\{1, \dots, d\} \setminus \{t\}}])}{\text{Var}(X_t)}. \quad (15)$$

In R^2 -Sortability, $\tau(X, s)$ denotes the coefficient of determination R^2 of X_s . R^2 -Sortability measures the consistency between the causal structure order and the increasing order of R^2 . If $\mathbf{v}_{R^2}(X, \mathcal{G}) = 1$, the causal order can be fully identified by the increasing order of R^2 . If $\mathbf{v}_{R^2}(X, \mathcal{G}) = 0$, the causal order can be fully identified by the decreasing order of R^2 . The only difference between R^2 -SortnRegress and Var-SortnRegress lies in the definition of τ . We use the implementation of R^2 -SortnRegress algorithm in CausalDisco python package provided by the authors, available at <https://github.com/CausalDisco/CausalDisco>.

B.6 NOTEARS

Based on (3), the NOTEARS score function is:

$$\min_{\mathcal{G}} F(\mathcal{G}; \mathbf{X}) = \frac{1}{2n} \|\mathbf{X} - \mathbf{X}W(\mathcal{G})\|_F^2 + \lambda \|W(\mathcal{G})\|_1 \quad \text{s.t.} \quad h(W(\mathcal{G})) = 0, \quad (16)$$

where $\|\cdot\|_F$ is the Frobenius norm, $\|\cdot\|_1$ is the sum of absolute values of all elements in the matrix.

The unconstrained objective function obtained through the ALM is:

$$\min_{\mathcal{G}} L_{\mu}(W(\mathcal{G}), \theta, \alpha) = \frac{1}{2n} \|\mathbf{X} - \mathbf{X}W(\mathcal{G})\|_F^2 + \lambda \|W(\mathcal{G})\|_1 + \alpha_t h(W(\mathcal{G})) + \frac{\mu_t}{2} |h(W(\mathcal{G}))|^2. \quad (17)$$

The update rule for the parameters is:

$$\begin{aligned} W(\mathcal{G})_k, \theta_k &= \arg \min_{W(\mathcal{G}), \theta} L_{\mu}(W(\mathcal{G}), \theta, \alpha) \\ \alpha_{k+1} &= \alpha_k + \mu_k h(W(\mathcal{G})_k) \\ \mu_{k+1} &= \begin{cases} \eta \mu_k, & \text{if } |h(W(\mathcal{G})_k)| > \gamma |h(W(\mathcal{G})_{k-1})| \\ \mu_k, & \text{otherwise} \end{cases}, \end{aligned} \quad (18)$$

where θ represents the parameters of a neural network used to fit a nonlinear function, and θ can be ignored for a linear model. The hyperparameters are usually set as $\eta = 10$ and $\gamma = \frac{1}{4}$.

In practice, the optimization stopping criterion is $h(W(\mathcal{G})_k) < \epsilon \in \{1e^{-6}, 1e^{-8}, 1e^{-10}\}$, which does not guarantee the output to be a DAG. Finally, for values in $W(\mathcal{G})$ with absolute values smaller than a threshold τ , we set them to 0 in order to obtain a DAG as closely as possible. We use the implementation of NOTEARS algorithm provided by the authors, available at <https://github.com/xunzheng/notears>.

B.7 GOLEM

GOLEM (Ng et al., 2020) proposed an improved loss function to address the numerical and ill-conditioned issues that often arise during the multiple iterations of optimization in NOTEARS. The

unconstrained optimization problem formulated by GOLEM is:

$$\min_{W(\mathcal{G}) \in \mathbb{R}^{d \times d}} \mathcal{S}_i(W(\mathcal{G}); \mathbf{X}) = \mathcal{L}_i(W(\mathcal{G}); \mathbf{X}) + \lambda_1 \|W(\mathcal{G})\|_1 + \lambda_2 h(W(\mathcal{G})), \quad (19)$$

where λ_1 and λ_2 are hyperparameters, $i \in \{1, 2\}$.

When assuming linear Gaussian with non-equal variances, that is, GOLEM-NV:

$$\mathcal{L}_1(W(\mathcal{G}); \mathbf{X}) = \frac{1}{2} \sum_{i=1}^d \log \left(\sum_{k=1}^n \left(X_i^{(k)} - W(\mathcal{G})_i^\top X^{(k)} \right)^2 \right) - \log |\det(I - W(\mathcal{G}))|, \quad (20)$$

where $X_i^{(k)}$ denotes k -th data point of X_i .

When assuming linear Gaussian with equal variances, that is, GOLEM-EV:

$$\mathcal{L}_2(W(\mathcal{G}); \mathbf{X}) = \frac{d}{2} \log \left(\sum_{i=1}^d \sum_{k=1}^n \left(X_i^{(k)} - W(\mathcal{G})_i^\top X^{(k)} \right)^2 \right) - \log |\det(I - W(\mathcal{G}))|. \quad (21)$$

The authors proved that in the case of linear Gaussian with equal variances, when the hard DAG constraint is not satisfied, the least-squares optimal solution of NOTEARS returns a cyclic graph, whereas the optimal solution of GOLEM-EV corresponds to the ground-truth. GOLEM combines the maximum likelihood objective function with a soft DAG constraint, replacing the least-squares objective function and hard DAG constraint, making the optimization easier to solve and the results better. We use the implementation of GOLEM-EV algorithm provided by the authors, available at <https://github.com/ignavierng/golem>.

B.8 NOTEARS-MLP

NOTEARS-MLP (Zheng et al., 2020) extends the differentiable causal discovery framework to the nonlinear case. Each variable X_j is defined as follows:

$$X_j = f_j(X_{pa(X_j)}, U_j), \forall j = 1, \dots, d. \quad (22)$$

The authors proved that f_j is independent of X_k if and only if $\|\partial_k f_j\|_{L^2} = 0$, where $\|\cdot\|_{L^2}$ is the L^2 -norm. Next, they define nonlinear causal effects through partial derivatives:

$$[W(f)]_{kj} = \|\partial_k f_j\|_{L^2}, \quad (23)$$

where $[W(f)]_{kj}$ is the causal effects from X_k to X_j , and $W(f) = W(f_1, \dots, f_d) \in \mathbb{R}^{d \times d}$.

In practice, neural networks are used to fitting nonlinear functional relationships f_j :

$$\text{MLP}_j(\mathbf{X}; A^{(1)}, \dots, A^{(h)}) = \sigma \left(A^{(h)} \sigma \left(\dots A^{(2)} \sigma \left(A^{(1)} \mathbf{X} \right) \right) \right), \quad (24)$$

where $A^{(\ell)} \in \mathbb{R}^{m_\ell \times m_{\ell-1}}$, σ is the activation function.

For the convenience of derivative computation, the authors proved that MLP_j are independent of X_k if and only if the k -th column of the first-layer weight matrix $A^{(1)}$ is entirely zero. The parameters of MLP_j are $\theta_j = (A_j^{(1)}, \dots, A_j^{(h)})$. The authors ultimately obtain a weighted adjacency matrix representation that is independent of the depth of the neural network:

$$[W(\theta)]_{kj} = \left\| \text{the } k\text{-th column of } A_j^{(1)} \right\|_2. \quad (25)$$

The objective function of NOTEARS-MLP is:

$$\min_{\theta} \frac{1}{n} \sum_{j=1}^d L(X_j, \text{MLP}_j(\mathbf{X}; \theta_j)) + \lambda \left\| A_j^{(1)} \right\|_1 \quad \text{s.t.} \quad h(W(\theta)) = 0. \quad (26)$$

NOTEARS-MLP trains d neural networks and represents the acyclicity constraint only through the first-layer parameters of neural networks. We use the implementation of NOTEARS-MLP algorithm provided by the authors, available at <https://github.com/xunzheng/notears>.

B.9 GRAN-DAG

GraN-DAG also extends NOTEARS to the nonlinear case and considers the ANM data generation mechanism:

$$X_j = f_j(X_{pa(X_j)}) + U_j, \forall j = 1, \dots, d. \quad (27)$$

The authors state that the parameters of j -th neural network (NN) are:

$$\phi_{(j)} = \{W_{(j)}^{(1)}, \dots, W_{(j)}^{(L+1)}\}, \quad (28)$$

where $W_{(j)}^{(\ell)}$ is the ℓ -th weight matrix of the j -th NN.

They define the j -th connection matrix:

$$C_{(j)} = |W_{(j)}^{(L+1)}| \dots |W_{(j)}^{(2)}| |W_{(j)}^{(1)}|. \quad (29)$$

Based on the above definition, the authors construct a weighted adjacency matrix related to the depth of the NN:

$$(W_\phi)_{ij} = \begin{cases} \sum_{k=1}^m (C_{(j)})_{ki}, & \text{if } j \neq i \\ 0, & \text{otherwise} \end{cases}, \quad (30)$$

where m is the output dimension of the NN.

The objective function of GraN-DAG is:

$$\max_{\phi} \mathbb{E}_{X \sim P(X)} \sum_{j=1}^d \log p_j(X_j | X_{pa(X_j)}; \phi_{(j)}) \quad \text{s.t. } h(\phi) = \text{Tr } e^{W_\phi} - d = 0. \quad (31)$$

The unconstrained objective function of GraN-DAG is:

$$\max_{\phi} \mathcal{L}(\phi, \alpha_t, \mu_t) = \mathbb{E}_{X \sim P(X)} \sum_{j=1}^d \log p_j(X_j | X_{pa(X_j)}; \phi_{(j)}) - \alpha_t h(\phi) - \frac{\mu_t}{2} h(\phi)^2. \quad (32)$$

When the data generation mechanism follows the nonlinear Gaussian additive noise model, it can be proven that the optimal solution of GraN-DAG corresponds to the ground-truth. We use the implementation of GraN-DAG algorithm provided by the authors, available at <https://github.com/kurowasan/GraN-DAG>.

B.10 NOCURL

Since a DAG is related to curl-free functions on its edge set, NoCurl proposed a new representation of a DAG:

$$A = \gamma(W, p), \quad (33)$$

where W is a skew-symmetric matrix with $W = -W^T$, $p \in \mathbb{R}^d$ is the potential function on the vertices of the graph.

The authors further proved that:

$$\gamma(W, p) = W \circ \text{ReLU}(\text{grad}(p)), \quad (34)$$

where grad is the gradient operator.

The optimization problem established by NoCurl is:

$$(W^*, p^*) = \underset{W, p}{\operatorname{argmin}} F(\gamma(W, p), \mathbf{X}), \quad (35)$$

with the optimal DAG $A^* = W^* \circ \text{ReLU}(\text{grad}(p^*))$. NoCurl implicitly ensures the acyclicity constraint, overcoming the shortcomings of the ALM, avoiding multiple iterations, and improving computational efficiency. We use the implementation of NoCurl algorithm provided by the authors, available at <https://github.com/fishmoon1234/DAG-NoCurl>.

B.11 DAGMA

DAGMA (Bello et al., 2022) proposed a log-determinant form of acyclicity representation $h_{\text{ldet}}^s(W) = -\log \det(sI - W \circ W) + d \log s$ ($s > 0$), which has three advantages compared to the exponential acyclicity constraints $h_{\text{expm}}(W) = \text{Tr}(e^{W \circ W}) - d$ (Zheng et al., 2018) and polynomial acyclicity constraints $h_{\text{poly}}(W) = \text{Tr}[(I + \alpha W \circ W)^d] - d$ ($\alpha > 0$) (Yu et al., 2019). The authors proved that:

$$\begin{aligned} h_{\text{expm}}(W) &= \sum_{k=0}^{\infty} \frac{1}{k!} \text{Tr}((W \circ W)^k) - d, \\ h_{\text{poly}}(W) &= \sum_{k=0}^d \frac{\binom{d}{k}}{d^k} \text{Tr}((W \circ W)^k) - d, \end{aligned} \quad (36)$$

where $\text{Tr}((W \circ W)^k)$ represents the information of cycles of length k . The information of cycles of length k in $h_{\text{expm}}(W)$ and $h_{\text{poly}}(W)$ is weakened by $\frac{1}{k!}$ and $\frac{\binom{d}{k}}{d^k}$, respectively. It can be theoretically proven that $h_{\text{ldet}}^s(W)$ is an upper bound of $h_{\text{expm}}(W)$ and $h_{\text{poly}}(W)$, retaining more information about the cycles.

The authors also proved that:

$$\begin{aligned} \nabla h_{\text{expm}}(W) &= 2(e^{W \circ W})^\top \circ W \\ \nabla h_{\text{poly}}(W) &= 2 \left(\left(I + \frac{1}{d} W \circ W \right)^{d-1} \right)^\top \circ W. \\ \nabla h_{\text{ldet}}^s(W) &= 2((sI - W \circ W)^{-1})^\top \circ W \end{aligned} \quad (37)$$

$\nabla h_{\text{expm}}(W)$ and $\nabla h_{\text{poly}}(W)$ are prone to the vanishing gradient problem. It can be theoretically proven that $\nabla h_{\text{ldet}}^s(W)$ is an upper bound of $\nabla h_{\text{expm}}(W)$ and $\nabla h_{\text{poly}}(W)$, retaining more information about the cycles.

The third advantage is that, in practice, $h_{\text{ldet}}^s(W)$ and $\nabla h_{\text{ldet}}^s(W)$ are faster to compute. Because the computation of $h_{\text{ldet}}^s(W)$ and $\nabla h_{\text{ldet}}^s(W)$ involves matrix log-determinant and matrix inverse, both of which have been extensively studied and solved. In contrast, other acyclicity constraints and their partial derivatives involve multiple matrix-to-matrix multiplications, which are slower. We use the implementation of DAGMA algorithm provided by the authors, available at <https://github.com/kevinsbello/dagma>.

C RELATED WORK

Differentiable causal discovery methods. Building on traditional score-based causal discovery algorithms, NOTEARS (Zheng et al., 2018) transformed discrete constrained optimization into smooth equality-constrained optimization. This formulation has been extended to various settings, including more efficient linear models (GOLEM (Ng et al., 2020), NoCurl (Yu et al., 2021), NOFEARS (Wei et al., 2020), LEAST (Zhu et al., 2021)), neural networks (NOTEARS-MLP (Zheng et al., 2020), GraN-DAG (Lachapelle et al., 2019), DAGMA (Bello et al., 2022), DARING (He et al., 2021), CASTLE (Kyono et al., 2020)), generative adversarial networks (SAM (Kalinathan et al., 2022)), variational autoencoders (D-VAE (Zhang et al., 2019)), graph neural network (GAE (Ng et al., 2019), DAG-GNN (Yu et al., 2019)), federated learning (FedDAG (Gao et al., 2021)), reinforcement learning (RL-BIC (Zhu et al., 2019)), interventional data (DCDI (Brouillard et al., 2020)), time series data (DYNOTEARS (Pamfil et al., 2020)), multi-domain data (DICD (Wang et al., 2022), ReScore (Zhang et al., 2023), CASPER (Liu et al., 2023)), and domain adaptation (CAE (Yang et al., 2021)). Although differentiable causal discovery has made significant progress, it is also affected by Var-Sortability (Reisach et al., 2021; 2023) and highly non-convex optimization problems (Ng et al., 2024). Recent research by Deng et al. (2024) shows that differentiable causal discovery methods can achieve scale invariance and global optimization when the correct loss function is used.

D TABLE RESULTS FOR RUNTIME OF THE BENCHMARK METHODS

The results in Table 2, 3, 4, and 5 show that differentiable causal discovery, exemplified by DAGMA, NOTEARS-MLP, and NoCurl, achieve superior performance with almost negligible runtime cost.

Table 5: Results for runtime (in seconds) on degree $k = 2$ graphs of 10 and 20 nodes. The reported results are the mean and standard deviation of the runtime over 10 repetitions across different graph types, vanilla and model assumption violation scenarios.

Method	d	Runtime (seconds)
PC	10	1.29 ± 0.24
	20	1.91 ± 0.35
GES	10	1.64 ± 0.53
	20	7.82 ± 2.03
DirectLiNGAM	10	1.25 ± 0.28
	20	2.06 ± 0.43
Var-SortnRegress	10	1.11 ± 0.25
	20	1.26 ± 0.29
R^2 -SortnRegress	10	1.13 ± 0.16
	20	1.24 ± 0.35
NOTEARS	10	9.35 ± 2.63
	20	35.57 ± 4.71
GOLEM	10	130.23 ± 2.54
	20	178.52 ± 3.41
NoCurl	10	5.68 ± 0.49
	20	10.29 ± 1.84
CAM	10	45.73 ± 2.67
	20	113.37 ± 3.85
NOTEARS-MLP	10	4.70 ± 0.82
	20	5.84 ± 0.97
GraN-DAG	10	119.28 ± 2.15
	20	211.59 ± 4.35
DAGMA	10	2.41 ± 0.32
	20	3.19 ± 0.48

Table 9: Nonlinear Setting, for SF-2 graphs of 10, 20, 50 nodes.

	Vanilla model		Latent confounders		Measurement error		Autoregressive		Heterogeneous		Unfaithful		Scale-variant		Missing		Mechanism violation	
	SHD ₁	SHD ₂	SHD ₁	SHD ₂	SHD ₁	SHD ₂	SHD ₁	SHD ₂	SHD ₁	SHD ₂	SHD ₁	SHD ₂	SHD ₁	SHD ₂	SHD ₁	SHD ₂	SHD ₁	SHD ₂
10 nodes																		
Random	263.3±9	66.6±8.7	25.3±2.5	66.6±7.6	27.8±3.0	70.2±4.4	27.7±2.1	66.7±6.3	26.1±2.8	65.8±7.8	27.4±1.8	66.4±10.8	27.3±2.6	68.1±10.2	27.0±3.0	71.0±4.4	27.7±2.3	61.5±10.6
PC	201.3±30	165.3±40	16.4±1.4	85.7±4.1	16.4±1.4	85.7±4.1	17.7±1.9	78.3±8.3	22.3±3.2	71.1±4.0	17.1±3.1	79.2±3.5	20.1±3.0	80.8±8.5	19.2±3.5	77.7±3.8	15.0±4.2	71.9±6.1
GES	165.3±40	73.3±6.6	13.5±1.3	72.4±11.0	15.1±2.9	78.6±8.7	15.3±1.7	70.6±8.9	22.4±4.1	60.9±8.5	14.9±1.6	50.8±16.6	15.4±1.0	73.3±6.6	16.2±3.7	60.2±10.0	7.7±6.3	32.3±18.4
2.1-16	4.7±1.1	23.1±7	2.5±1.7	9.6±6.1	9.5±1.5	150.4±8.4	3.9±1.1	139.1±6.6	15.2±2.5	6.6±7.6	3.7±1.2	22.5±4.1	2.1±1.6	67.5±8	2.7±1.4	10.3±10.0	59.2±10.5	9.2±10.5
NOTEARS MLP	8.9±4.2	41.7±15.5	12.3±2.4	62.1±10.5	11.4±1.7	52.7±16.4	12.2±2.4	57.7±11.8	16.4±3.1	49.9±7.2	10.0±1.8	58.5±11.0	10.8±3.9	49.5±12.2	8.9±1.9	40.0±10.0	7.3±1.7	27.0±7.9
Grn-DAG	30.6±39	11.2±13.4	9.7±2.0	48.8±11.3	13.5±2.0	64.5±17.9	9.5±4.3	45.8±22.3	53.1±2.0	31.9±15.1	7.4±2.2	47.7±12.3	19.1±9	8.9±12.3	1.4±1.7	7.3±10.2	9.3±2.7	53.1±14.4
DAGMA	102.4±5	45.2±18.2	14.2±1.7	76.2±10.4	13.2±1.8	68.0±9.8	14.1±2.0	69.0±10.0	16.6±1.5	42.5±10.5	8.2±2.7	66.2±15.7	13.5±2.5	66.8±3.3	8.9±1.6	32.7±13.4	2.5±3.1	14.6±21.1
20 nodes																		
Random	106.6±6.1	280.1±10.9	102.6±6.1	273.8±32.1	104.6±7.1	266.8±28.7	104.3±6.3	268.2±16.0	100.1±5.2	270.2±33.8	101.6±6.2	278.9±20.9	101.6±7.9	267.1±36.6	104.2±8.0	259.1±30.5	104.5±5.6	265.1±31.6
PC	433.5±6	346.5±16.2	34.7±1.7	360.6±16.0	38.8±4.2	360.6±16.0	40.3±4.1	361.6±10.0	63.2±4.8	327.3±31.3	38.2±4.1	336.2±25.5	43.3±5.6	346.5±16.2	44.6±5.6	336.8±19.1	45.0±4.9	247.1±31.1
GES	373.6±5	292.9±33.8	33.0±2.7	354.0±18.1	33.0±2.7	354.0±18.1	33.0±2.7	315.3±38.4	67.5±6.4	280.2±26.1	35.2±3.4	252.3±13.7	37.6±6.2	292.9±33.8	39.6±7.6	314.3±18.0	23.5±1.8	323.0±37.7
CAM	6.7±2.3	18.5±12.1	23.1±2.3	24.5±10.2	26.5±4.9	27.9±14.6	12.3±5.1	71.2±40.5	41.5±7.2	44.1±24.8	8.7±2.4	88.7±27.2	6.7±2.3	18.5±12.1	4.8±1.7	9.7±10.8	23.9±4.0	21.9±43.0
NOTEARS MLP	22.6±3.3	21.4±18.1	31.9±2.9	33.6±41.2	29.1±2.5	29.1±2.5	37.7±2.1	36.0±9.0	37.0±1.0	36.1±0.0	20.5±2.2	25.6±7.28	26.6±1.8	29.4±54.4	18.8±3.1	184.5±40.6	13.3±2.8	123.8±25.6
Grn-DAG	19.1±13.9	91.7±79.7	33.3±6.3	345.0±14.4	35.9±13.1	346.3±29.8	36.3±3.7	308.7±15.3	26.5±4.2	291.3±17.0	22.8±2.5	299.9±11.9	16.4±16.4	112.9±90.8	16.7±6.4	210.7±70.0	26.1±7.1	269.0±75.5
DAGMA	23.0±3.8	198.2±42.9	34.1±4.2	347.2±18.2	29.6±1.7	328.8±26.2	37.0±6.8	360.0±11.2	37.0±1.0	361.0±0.0	28.0±8.5	217.4±36.6	26.1±3.0	240.6±47.0	20.8±5.2	189.5±11.1	8.2±3.2	88.1±47.2
50 nodes																		
Random	640.0±12.1	1727.7±57.4	632.7±10.6	1797.8±62.3	632.6±9.9	1697.4±18.8	629.2±13.1	1734.0±62.4	634.9±10.6	1716.2±14.8	637.7±21.5	1724.3±97.4	649.8±12.8	1690.1±138.5	639.5±18.3	1700.5±78.9	639.0±17.1	1683.9±148.1
PC	152.4±2	189.7±102.5	17.2±1.2	152.6±125.9	48.8±7.8	190.7±120.7	54.2±22.2	579.1±102.6	131.2±11.1	355.4±196.8	20.7±5.2	527.0±100.2	15.2±4.2	189.7±102.5	13.4±4.9	125.4±55.6	101.7±10.9	1494.7±139.5
GES	53.4±5.7	139.8±20.1	78.2±7.5	195.7±135.2	82.1±3.7	222.7±19.4	97.6±10.8	200.4±40.8	63.6±4.5	201.9±12.2	54.0±5.7	149.0±17.8	62.2±3.9	182.0±151.9	53.1±9.6	198.2±126.1	39.1±7.9	982.8±79.4
NOTEARS MLP	68.9±9.2	157.2±173.8	89.3±4.2	175.3±127.3	98.0±10.0	239.9±16.6	89.7±19.7	210.7±194.3	90.3±3.0	256.4±38.5	77.1±5.6	250.7±17.3	57.2±9.3	163.5±136.4	62.1±12.4	140.9±330.5	78.8±6.6	267.7±138.2
Grn-DAG	53.7±40.0	134.0±122.1	83.2±2.7	109.1±11.9	85.5±2.5	20.6±15.3	77.6±4.0	234.0±110.9	98.9±3.1	217.9±24.7	38.0±5.4	150.2±120.4	67.6±6.2	175.0±136.4	53.6±10.0	140.7±137.7	29.2±5.2	86.5±132.9

Table 10: Linear Setting, for SF-4 graphs of 10, 20, 50 nodes.

Vanilla model		Latent confounders		Measurement error		Autoregressive		Heterogeneous		Unfaithful		Scale-variant		Missing		Mechanism violation	
10 nodes	SHD ₁	SHD ₁	SHD ₁	SHD ₁	SHD ₁	SHD ₁	SHD ₁	SHD ₁	SHD ₁	SHD ₁	SHD ₁	SHD ₁	SHD ₁	SHD ₁	SHD ₁	SHD ₁	SHD ₁
Random	28.7±2.8	77.2±4.4	30.0±3.8	71.4±9.8	30.6±3.8	74.1±7.6	29.3±1.8	73.2±6.1	28.9±3.1	76.4±3.6	31.1±3.7	78.6±3.8	30.8±3.0	73.1±6.1	30.0±4.6	75.2±5.5	27.0±3.8
PC	27.2±2.8	71.4±4.1	21.4±1.5	78.6±2.6	25.8±4.1	70.8±6.0	27.1±1.1	73.2±5.4	28.8±2.4	75.7±5.2	28.5±2.3	74.2±4.4	27.3±2.8	71.4±4.1	26.3±2.7	71.1±4.2	26.3±2.7
GES	23.0±5.1	62.7±9.8	27.2±6.6	64.3±11.3	24.5±3.7	60.0±14.5	22.8±5.5	53.5±14.0	27.4±5.6	64.7±13.8	24.6±4.5	67.5±16.5	23.0±5.1	62.7±9.8	26.1±4.6	67.0±6.7	26.1±2.2
2.1-16	21.6±5.0	63.8±11.6	24.9±2.8	65.1±6.5	26.1±2.2	67.8±4.3	21.4±4.1	61.6±7.9	23.0±4.2	58.6±8.0	19.8±4.7	61.7±8.5	24.9±5.5	71.4±9.4	23.3±4.9	65.0±6.5	30.0±0.0
Var-Semireg	10.7±2.7	25.5±6.9	13.0±4.7	21.4±40.5	17.8±1.9	27.6±4.0	15.7±1.6	18.6±13.4	15.5±3.3	20.1±7.1	13.2±5.2	24.3±10.8	26.2±4.4	61.8±6.2	9.9±4.0	17.2±12.9	20.6±4.5
NOTEARS	25.3±5.0	60.2±8.5	24.5±4.0	51.1±11.7	27.9±3.3	60.2±7.2	25.8±4.5	49.2±11.5	25.8±4.5	57.4±11.6	26.5±5.9	57.3±11.8	25.3±5.0	60.2±8.5	25.5±4.7	47.7±5.1	29.9±1.5
Grn-DAG	5.2±1.2	16.5±4.2	8.0±2.5	18.5±11.3	9.5±2.7	18.5±11.3	12.6±2.5	21.6±12.3	10.9±4.8	14.7±11.8	2.4±1.6	14.7±11.8	24.7±4.9	40.8±7.7	9.9±2.4	18.6±10.9	29.1±2.2
GOLEM	2.2±2.1	9.0±4.8	9.3±2.7	21.3±4.5	24.6±2.6	69.4±6.2	18.4±1.7	44.1±10.2	9.4±3.0	21.5±9.6	1.6±1.9	7.0±7.8	25.7±4.5	6.9±5.1	1.4±1.4	6.9±5.1	2.8±1.2
NCut	4.0±3.2	17.4±13.1	10.4±2.8	24.9±13.4	15.5±2.7	37.6±8.0	14.1±3.7	24.2±12.8	10.2±3.6	21.3±15.3	2.5±1.4	17.0±10.6	26.8±2.8	75.9±9.2	2.6±2.4	11.7±6.6	20.9±1.1
DAGMA	2.0±2.2	10.9±11.9	8.5±2.5	20.6±8.5	15.2±3.8	37.9±7.6	15.5±1.8	15.7±12.0	9.6±3.5	12.8±10.8	1.3±1.5	7.3±10.6	25.7±2.9	79.3±4.3	1.1±1.1	6.4±4.6	29.1±0.7
20 nodes	SHD ₁	SHD ₁	SHD ₁	SHD ₁	SHD ₁	SHD ₁	SHD ₁	SHD ₁	SHD ₁	SHD ₁	SHD ₁	SHD ₁	SHD ₁	SHD ₁	SHD ₁	SHD ₁	SHD ₁
Random	113.4±6.5	314.7±10.9	109.9±5.4	303.6±17.9	110.9±7.0	308.7±17.0	109.9±8.4	303.2±13.1	114.5±4.7	310.0±14.8	113.3±6.9	309.1±10.0	113.3±6.9	309.1±10.0	111.5±7.2	303.0±15.9	115.3±8.5
PC	77.5±2.8	345.1±10.6	78.9±16.4	349.5±7.8	92.6±5.1	329.8±15.4	82.8±1.5	341.9±7.3	78.9±4.4	347.3±6.1	76.2±4.7	347.6±10.9	77.5±2.8	345.1±10.6	78.1±4.2	350.5±10.8	67.6±2.7
GES	96.1±11.1	235.8±10.8	110.1±10.1	235.8±10.8	110.1±10.1	235.8±10.8	110.1±10.1	235.8±10.8	110.1±10.1	235.8±10.8	110.1±10.1	235.8±10.8	110.1±10.1	235.8±10.8	110.1±10.1	235.8±10.8	110.1±10.1
Direct-INGAM	58.3±10.3	283.4±17.1	76.9±6.6	280.1±11.7	67.1±8.0	310.0±12.8	61.7±8.0	244.8±7.2	72.9±10.8	292.2±36.2	65.6±5.8	296.3±26.4	68.3±10.9	344.1±11.1	60.2±12.1	287.6±22.9	69.0±3.7
Var-Semireg	54.7±10.4	104.0±12.3	103.5±10.3	88.9±21.5	111.2±17.1	119.0±44.7	108.9±7.3	187.1±35.8	106.3±6.9	180.0±18.3	66.8±19.7	118.2±35.5	108.5±20.7	235.3±21.8	57.9±9.9	108.5±12.2	61.3±4.8
NOTEARS	128.5±10.1	251.4±16.7	137.4±9.8	229.1±25.9	135.7±4.9	254.4±20.1	133.3±4.9	225.8±18.4	141.7±5.3	246.0±12.8	137.5±18.1	261.9±13.0	128.5±10.1	251.4±16.7	127.5±10.9	254.7±16.7	177.7±4.2
Grn-DAG	13.5±1.8	42.1±15.9	35.7±8.8	76.2±12.8	45.7±8.8	24.7±18.3	84.5±16.2	185.3±40.5	30.6±10.7	100.3±45.1	10.3±6.6	68.6±16.2	68.2±5.5	35.5±19.5	10.5±6.8	87.8±41.0	68.0±4.9
NCut	8.5±5.4	83.9±50.1	52.3±7.4	136.2±27.6	69.8±0.4	360.4±1.2	70.5±48.8	360.3±0.8	40.1±18.1	190.1±51.8	10.1±5.7	82.3±49.2	66.1±6.2	35.1±81.3	7.9±4.0	72.7±28.9	67.8±0.9
DAGMA	8.3±6.2	18.1±30.5	41.8±6.5	11.1±25.3	85.6±7.0	270.6±15.4	77.3±11.9	207.9±10.5	35.9±9.7	137.0±43.1	14.5±11.9	75.5±100.2	73.9±10.7	336.5±20.6	7.5±5.3	52.5±40.6	68.6±0.9
50 nodes	SHD ₁	SHD ₁	SHD ₁	SHD ₁	SHD ₁	SHD ₁	SHD ₁	SHD ₁	SHD ₁	SHD ₁	SHD ₁	SHD ₁	SHD ₁	SHD ₁	SHD ₁	SHD ₁	SHD ₁
Random	659.9±20.4	1901.2±49.6	670.5±22.8	1899.4±30.8	658.5±18.7	1904.2±40.1	667.6±19.5	1908.0±47.7	668.3±14.9	1914.0±28.9	659.2±28.8	1918.7±22.9	664.4±18.1	1914.4±31.4	669.1±13.1	1911.2±38.4	648.4±21.4
Direct-INGAM	207.1±25.8	2126.2±120.4	318.1±21.1	1975.1±103.8	217.2±2.2	85.3±4.1	27.9±2.0	87.1±7.2	29.1±3.0	72.6±40.9	230.6±21.5	2196±29.1	229.4±5.5	2187.1±17.9	205.5±22.0	2101.7±139.7	190.0±0.0
Var-Semireg	347.1±52.9	422.1±10.6	80.6±3.3	694.0±106.0	695.9±36.7	1112.8±90.2	770.7±42.6	74.7±4.8	744.3±30.4	800.0±149.5	441.3±70.4	800.0±1	441.3±70.4	800.0±1	441.3±70.4	800.0±1	441.3±70.4
NOTEARS	777.6±55.5	1057.5±65.4	986.8±11.7	1489.5±81.2	788.3±19.4	1712.3±35.2	914.7±31.2	1419.7±87.8	937.2±95.6	1305.2±22.4	907.2±96.8	1098.2±49.4	777.6±55.5	1057.5±65.4	777.6±55.5	1057.5±65.4	777.6±55.5
Grn-DAG	46.3±9.1	1062.9±31.6	119.3±1.7	1579.3±61.0	186.2±34.2	2391.1±61.0	199.3±9.1	111.4±24.2	59.0±10.5	17.7±3.891	1740.4±1	289.2±21.2	46.3±9.1	1062.9±31.6	46.3±9.1	1062.9±31.6	46.3±9.1
DAGMA	27.9±11.6	494.2±78.0	113.8±7.7	1937.1±116.1	197.2±47	2387.6±49	277.5±34	2326.0±47	134.2±12.4	132.6±7.6	333.4±23	996.9±239	27.9±11.6	494.2±78.0	27.9±11.6	494.2±78.0	27.9±11.6

F TABLE RESULTS FOR COMBINED MISSPECIFIED SCENARIOS

We consider two (confounding and heterogeneity), three (confounding, measurement error, and heterogeneity) and four (confounding, measurement error, heterogeneity, and autoregression) combined misspecified scenarios. The results in Table 15 and Table 16 show that under combined misspecified scenarios, the performance of various methods is worse compared to single misspecified scenario. However, differentiable causal discovery still achieves optimal or competitive performance.

Table 15: Linear Setting with two (confounding and heterogeneity), three (confounding, measurement error, and heterogeneity) and four (confounding, measurement error, heterogeneity, and autoregression) combined misspecified scenarios, for ER-2 graphs of 10 nodes.

10 nodes	Vanilla model		Latent confounders		Measurement error		Heterogeneous		Autoregressive		Two combined scenarios		Three combined scenarios		Four combined scenarios	
	SHD _↓	SID _↓	SHD _↓	SID _↓	SHD _↓	SID _↓	SHD _↓	SID _↓	SHD _↓	SID _↓	SHD _↓	SID _↓	SHD _↓	SID _↓	SHD _↓	SID _↓
Random	29.1±1.1	68.4±2.7	25.1±4.7	65.2±5.7	24.3±3.9	65.6±6.6	26.2±4.3	70.4±2.3	27.5±3.6	65.8±4.2	27.1±1.3	63.4±4.5	25.7±4.2	61.2±7.2	28.1±2.3	69.2±3.5
PC	12.4±3.1	40.9±13.4	18.1±4.7	58.1±15.6	19.4±4.1	48.0±13.1	13.8±2.6	47.5±10.3	14.5±2.0	44.8±9.5	19.5±3.2	61.0±8.3	21.8±4.4	55.3±13.5	21.6±4.0	59.3±10.0
GES	13.8±7.8	32.0±15.6	25.9±7.7	42.6±14.0	20.2±4.8	46.2±16.7	15.5±6.1	35.2±12.2	20.8±5.5	49.7±11.5	28.1±3.4	41.2±12.8	27.4±1.7	62.0±8.9	24.7±5.0	50.2±11.1
DirectLINGAM	19.6±3.3	46.1±10.6	20.4±5.0	42.0±6.0	17.6±2.4	48.8±12.4	16.3±3.9	34.7±9.9	19.7±4.2	50.4±8.4	22.4±2.4	46.7±12.7	18.0±3.1	43.7±6.3	18.9±3.2	38.9±4.8
Var-SortnRegress	11.2±3.5	8.4±8.5	17.6±5.8	12.6±9.9	19.6±2.8	11.4±8.7	17.9±3.3	8.6±9.3	18.8±2.4	16.5±10.6	21.5±3.9	11.5±10.8	20.7±3.3	13.0±9.9	20.3±3.3	23.7±9.9
R ² -SortnRegress	20.2±4.8	32.4±14.0	25.7±4.1	37.6±13.0	25.6±6.0	39.2±16.0	26.0±5.4	37.6±14.4	25.6±4.9	38.8±19.0	25.5±3.6	31.9±17.6	26.8±3.8	48.3±10.4	26.8±5.2	48.1±12.1
NOTEARS	1.5±1.6	1.8±4.2	8.5±3.9	9.5±8.1	12.5±2.0	19.6±8.6	5.5±2.7	5.4±5.1	12.2±3.6	27.5±14.2	12.3±3.4	33.1±5.7	14.1±3.5	25.2±6.8	18.6±3.9	31.8±10.1
GOLEM	1.4±1.4	0.4±1.2	6.7±2.8	14.2±9.8	17.8±2.5	43.1±13.3	6.5±4.5	9.8±8.1	16.6±4.0	34.9±16.9	10.3±1.7	24.3±6.0	18.6±3.7	46.2±9.1	19.7±0.5	72.7±6.3
NoCurl	2.0±1.8	5.1±5.8	9.1±4.2	5.4±3.9	11.8±1.8	17.9±8.4	6.6±2.9	5.5±5.7	14.8±2.5	17.5±10.8	13.4±3.4	7.4±6.9	15.4±4.3	19.8±7.0	20.2±3.5	21.3±8.2
DAGMA	1.2±1.2	3.3±3.3	8.4±3.9	8.8±7.7	12.6±2.5	18.5±8.6	5.5±2.3	12.0±8.2	12.2±3.6	28.4±15.3	12.0±3.8	12.3±8.9	14.9±4.4	20.4±6.6	17.9±3.1	50.2±14.2

Table 16: MLP Setting with two (confounding and heterogeneity), three (confounding, measurement error, and heterogeneity) and four (confounding, measurement error, heterogeneity, and autoregression) combined misspecified scenarios, for ER-2 graphs of 10 nodes.

10 nodes	Vanilla model		Latent confounders		Measurement error		Heterogeneous		Autoregressive		Two combined scenarios		Three combined scenarios		Four combined scenarios	
	SHD _↓	SID _↓	SHD _↓	SID _↓	SHD _↓	SID _↓	SHD _↓	SID _↓	SHD _↓	SID _↓	SHD _↓	SID _↓	SHD _↓	SID _↓	SHD _↓	SID _↓
Random	27.5±2.6	62.6±7.3	29.4±1.1	62.8±9.4	29.3±1.2	69.4±3.1	28.9±1.8	70.4±2.8	27.9±2.3	72.2±5.5	28.4±2.0	67.2±4.5	24.2±4.1	66.2±6.3	29.6±2.3	68.5±5.6
PC	16.7±3.2	32.0±9.5	18.3±3.5	54.1±9.3	17.9±5.0	47.8±15.5	20.7±3.5	53.9±13.1	16.2±4.7	49.2±13.1	22.9±2.6	63.7±9.6	24.7±4.2	64.4±11.5	22.6±3.2	65.1±14.3
GES	21.9±4.4	46.9±8.6	28.2±7.5	53.1±18.6	20.2±5.4	54.7±14.3	27.2±4.9	48.6±10.9	18.9±5.2	43.2±15.5	28.6±2.9	53.0±8.4	28.5±3.5	62.9±6.3	24.6±2.3	61.6±6.7
CAM	12.4±3.6	39.3±16.5	16.6±4.2	42.4±17.7	19.7±4.7	59.6±12.0	19.9±3.6	49.1±7.2	16.0±3.4	46.0±16.6	23.4±3.1	60.9±6.1	25.4±3.3	53.8±13.4	24.4±1.9	63.4±8.8
NOTEARS-MLP	8.1±2.7	22.2±10.6	11.7±5.5	33.3±17.0	18.5±3.7	47.1±11.9	14.5±2.2	34.3±10.4	15.7±4.3	39.8±9.6	18.5±2.3	57.4±11.7	19.6±0.7	63.4±10.2	22.5±3.3	61.0±9.2
Grn-DAG	13.3±3.6	32.9±12.3	12.9±3.2	40.2±12.7	16.6±2.0	45.0±10.5	14.3±2.5	38.5±7.7	15.5±2.3	46.0±9.5	16.0±2.5	52.9±11.2	20.4±1.6	63.8±7.2	20.3±2.8	62.2±10.0
DAGMA	6.2±1.7	18.2±8.7	9.3±4.3	27.8±10.8	14.1±2.6	39.2±8.7	12.7±2.9	31.6±7.2	13.6±2.1	41.2±3.6	15.3±2.7	48.2±10.4	18.4±0.9	51.6±10.4	19.6±2.5	60.8±7.3

G TABLE RESULTS FOR NON-GAUSSIAN NOISE

We consider the vanilla model with exponential noise. The results in Table 17 and Table 18 show that differentiable causal discovery still achieve optimal or competitive performance when model assumptions are violated.

Table 17: Linear Setting with exponential noise, for ER-2 graphs of 10 nodes.

10 nodes	Vanilla model		Latent confounders		Measurement error		Autoregressive		Heterogeneous		Unfaithful		Scale-variant		Missing		Mechanism violation	
	SHD _↓	SID _↓	SHD _↓	SID _↓	SHD _↓	SID _↓	SHD _↓	SID _↓	SHD _↓	SID _↓	SHD _↓	SID _↓	SHD _↓	SID _↓	SHD _↓	SID _↓	SHD _↓	SID _↓
Random	29.2±1.4	64.0±8.2	27.5±1.9	69.2±3.0	25.7±2.6	65.3±4.8	25.1±2.8	65.7±8.5	26.5±3.0	64.2±7.6	24.7±3.1	71.0±2.4	28.1±1.6	65.7±5.4	26.6±3.4	66.1±5.4	27.9±3.0	67.2±7.2
PC	11.7±2.3	38.7±7.3	20.1±2.6	58.6±8.7	20.5±3.3	46.8±15.0	17.9±2.2	50.2±8.8	20.7±4.3	58.4±14.5	14.0±3.1	44.6±13.5	11.7±2.3	38.7±7.3	14.4±4.8	42.2±15.3	17.9±3.6	54.4±12.2
GES	19.4±8.1	41.1±19.7	24.8±6.8	41.6±16.0	24.4±4.4	56.7±14.5	22.3±3.8	37.0±11.9	26.4±5.0	45.3±17.3	18.3±5.8	38.2±14.5	18.9±8.3	37.9±20.1	17.3±6.5	34.3±19.7	19.6±4.9	42.5±12.9
DirectLINGAM	0.0±0.0	0.0±0.0	13.6±5.9	20.1±15.4	15.0±3.7	27.1±9.5	12.0±3.4	18.1±12.2	9.1±4.3	10.3±9.2	1.4±2.2	3.9±4.9	4.9±2.5	18.3±9.0	0.0±0.0	0.0±0.0	16.5±4.3	45.9±10.1
Var-SortnRegress	6.8±4.5	9.7±7.8	17.0±5.2	10.0±6.6	18.4±4.0	12.2±9.1	16.2±5.7	10.2±7.9	22.1±2.1	10.9±6.9	8.7±6.1	15.0±12.0	26.9±5.5	85.2±8.3	5.7±3.8	10.6±11.3	10.4±5.0	28.8±7.9
R ² -SortnRegress	15.7±6.7	30.4±12.2	27.9±4.1	35.9±16.1	26.4±5.9	44.5±14.9	23.1±5.0	37.3±13.9	26.6±3.4	32.6±14.3	20.9±8.3	34.3±15.0	15.7±6.7	30.4±12.2	15.8±7.1	27.0±10.0	29.0±4.2	55.8±11.0
NOTEARS	1.5±1.7	6.7±7.7	10.7±3.7	14.0±7.2	13.4±3.5	15.8±6.7	14.4±5.2	18.5±12.1	6.6±2.7	13.2±8.3	0.1±0.3	0.7±2.1	18.9±15	65.3±8.7	0.9±1.4	5.1±7.0	14.1±4.6	29.0±6.2
GOLEM	2.3±2.7	6.6±8.4	10.5±3.4	13.9±7.1	14.3±4.1	18.0±7.3	13.6±5.5	20.1±13.3	6.7±2.7	13.5±8.6	0.2±0.4	0.8±2.7	18.6±2.2	67.7±10.6	0.6±1.7	4.5±6.6	14.5±4.4	28.9±5.7
NoCurl	5.4±3.6	11.5±8.5	14.6±4.0	14.4±8.7	17.6±4.6	14.3±8.6	18.1±6.2	16.8±13.7	11.9±4.7	9.0±8.7	8.6±4.3	9.2±12.1	24.6±6.1	65.6±13.9	8.0±2.7	7.1±7.7	22.2±4.4	27.9±8.5
DAGMA	0.0±0.0	0.0±0.0	10.1±2.9	13.8±6.9	13.2±2.4	24.6±9.3	12.9±4.7	31.0±10.8	6.8±2.9	4.7±5.3	0.1±0.3	0.7±2.1	18.2±4.3	67.5±10.4	0.3±0.9	0.6±1.8	14.6±4.7	28.8±7.3

Table 18: MLP Setting with exponential noise, for ER-2 graphs of 10 nodes.

10 nodes	Vanilla model		Latent confounders		Measurement error		Autoregressive		Heterogeneous		Unfaithful		Scale-variant		Missing		Mechanism violation	
	SHD _↓	SID _↓	SHD _↓	SID _↓	SHD _↓	SID _↓	SHD _↓	SID _↓	SHD _↓	SID _↓	SHD _↓	SID _↓	SHD _↓	SID _↓	SHD _↓	SID _↓	SHD _↓	SID _↓
Random	27.9±2.2	62.6±6.6	26.3±2.5	62.6±5.9	28.0±1.3	66.9±6.9	25.0±3.2	68.1±3.2	29.8±1.1	67.6±4.3	26.1±4.0	60.3±5.7	30.0±1.0	69.9±2.6	25.4±5.0	63.2±8.4	26.7±4.8	64.7±9.2
PC	17.9±3.6	54.4±12.2	17.2±4.7	57.3±14.1	18.4±5.1	49.2±9.7	17.0±3.8	55.5±11.7	22.2±2.5	62.5±11.7	17.9±2.4	48.2±13.7	17.9±3.6	54.4±12.2	17.0±3.9	53.2±10.2	11.7±2.3	38.7±7.3
GES	19.6±4.9	42.5±12.9	22.8±7.0	42.8±18.0	19.6±4.7	46.6±12.4	17.1±4.8	47.0±16.2	30.0±2.3	58.5±8.4	23.6±6.0	44.1±12.3	19.6±4.9	42.2±12.5	20.7±4.3	44.2±14.5	19.4±8.1	41.1±19.7
CAM	11.2±2.6	40.3±17.3	14.1±4.9	35.5±14.1	21.2±3.8	61.8±11.2	13.7±3.5	42.3±16.4	23.3±3.5	63.5±12.6	12.5±5.2	35.7±21.4	11.2±2.6	40.3±17.3	12.3±2.1	41.6±20.6	21.2±7.6	62.8±13.1
NOTEARS-MLP	6.6±2.3	17.1±11.0	13.6±5.3	38.7±12.5	16.2±5.1	31.8±11.2	15.0±5.7	34.3±17.7	12.6±3.2	25.6±10.5	11.2±2.4	21.6±8.5	17.2±1.7	61.7±9.6	6.5±2.5	16.5±11.9	10.4±4.2	43.9±17.5
Grn-DAG	12.2±3.0	38.6±15.1	15.1±2.9	51.5±10.1	16.4±1.7	47.3±7.2	14.9±2.8	48.1±12.8	13.6±3.9	33.8±15.1	14.1±2.0	42.0±9.6	17.3±1.9	58.6±12.4	12.6±3.1	40.5±13.1	14.7±3.7	47.4±12.3
DAGMA	4.5±2.5	13.5±9.3	12.1±5.2	32.0±9.7	14.3±5.0	28.2±9.3	12.2±5.6	28.5±12.5	10.8±3.0	21.8±8.6	9.2±1.7	19.5±7.6	15.5±1.6	52.0±9.7	5.3±2.4	15.2±11.8	9.9±3.1	38.3±14.6

H SUMMARY OF THE MOST COMPETITIVE METHODS

Table 19: Summary of performances of the most competitive methods under linear setting. The reported results are the mean and standard deviation of the metrics over 10 repetitions across different graph types, vanilla and model assumption violation scenarios.

Method	d	SHD	SID
NOTEARS	10	8.51±5.92	23.88±23.32
	20	22.66±14.27	129.96±109.48
	50	61.07±39.27	945.01±774.12
GOLEM	10	9.02±6.90	30.67±27.39
	20	21.79±14.19	155.60±126.41
	50	55.57±35.00	1000.95±842.80
NoCurl	10	9.39±7.06	23.59±23.92
	20	32.22±23.51	99.43±99.13
	50	128.97±113.22	914.22±699.51
DAGMA	10	8.17±6.12	22.42±22.72
	20	20.98±14.10	127.31±112.55
	50	55.19±37.87	882.76±853.40

Table 20: Summary of performances of the most competitive methods under nonlinear setting. The reported results are the mean and standard deviation of the metrics over 10 repetitions across different graph types, vanilla and model assumption violation scenarios.

Method	d	SHD	SID
CAM	10	8.21±5.17	22.66±16.98
	20	22.23±12.36	117.37±79.80
	50	61.92±35.50	696.47±463.94
NOTEARS-MLP	10	12.23±3.63	44.68±10.27
	20	28.86±7.74	217.09±78.45
	50	73.40±18.41	1271.24±560.81
GraN-DAG	10	10.44±5.11	38.19±16.78
	20	30.46±8.23	214.86±75.62
	50	85.34±11.81	1453.35±575.73
DAGMA	10	12.64±4.66	46.99±17.88
	20	28.28±8.83	211.26±85.84
	50	72.07±20.15	1254.57±583.11

Table 21: Summary of performances of the most competitive methods under MLP setting. The reported results are the mean and standard deviation of the metrics over 10 repetitions across different graph types, vanilla and model assumption violation scenarios.

Method	d	SHD	SID
CAM	10	15.67±2.87	44.70±6.53
	20	34.88±9.23	179.14±29.38
	50	84.89±26.26	904.99±134.22
NOTEARS-MLP	10	12.93±4.24	34.63±12.57
	20	26.97±8.02	132.86±52.22
	50	70.08±23.80	740.91±239.08

I TABLE RESULTS ON REAL-WORLD DATA

We test the performance of 12 benchmark methods on the real-world Sachs (Sachs et al., 2005) dataset. Sachs is a bioinformatics dataset used to study the expression levels of various proteins and phospholipids in human cells, and it is a commonly used benchmark in the causal discovery field. We conduct experiments based on 7466 samples. The true graph structure of the Sachs dataset contains 11 nodes and 17 edges, and it is widely accepted by the biological research community.

Table 22: [Results on Sachs dataset.](#)

Method	SHD	SID
Random	33	56
PC	22	49
GES	30	47
DirectLiNGAM	14	50
Var-SortnRegress	19	49
R^2 -SortnRegress	22	51
NOTEARS	17	48
GOLEM	15	58
NoCurl	16	50
CAM	15	51
NOTEARS-MLP	14	46
GraN-DAG	15	45
DAGMA	12	42

The results in Table 22 show that, represented by DAGMA, differentiable causal discovery achieves optimal performance on the real-world Sachs dataset. Considering that Sachs is also regarded as a real-world heterogeneous dataset (Mooij et al., 2020), the results on both Sachs and synthetic datasets further indicate that differentiable causal discovery performs better under model assumption violations.

J FIGURE RESULTS ACROSS NODES, GRAPH TYPES, AND GRAPH DENSITIES

This section presents a comprehensive analysis of the figure results across varying numbers of nodes, graph types, and graph densities, in Figure 2, 3, 4, 5, 6, 7, and 8.

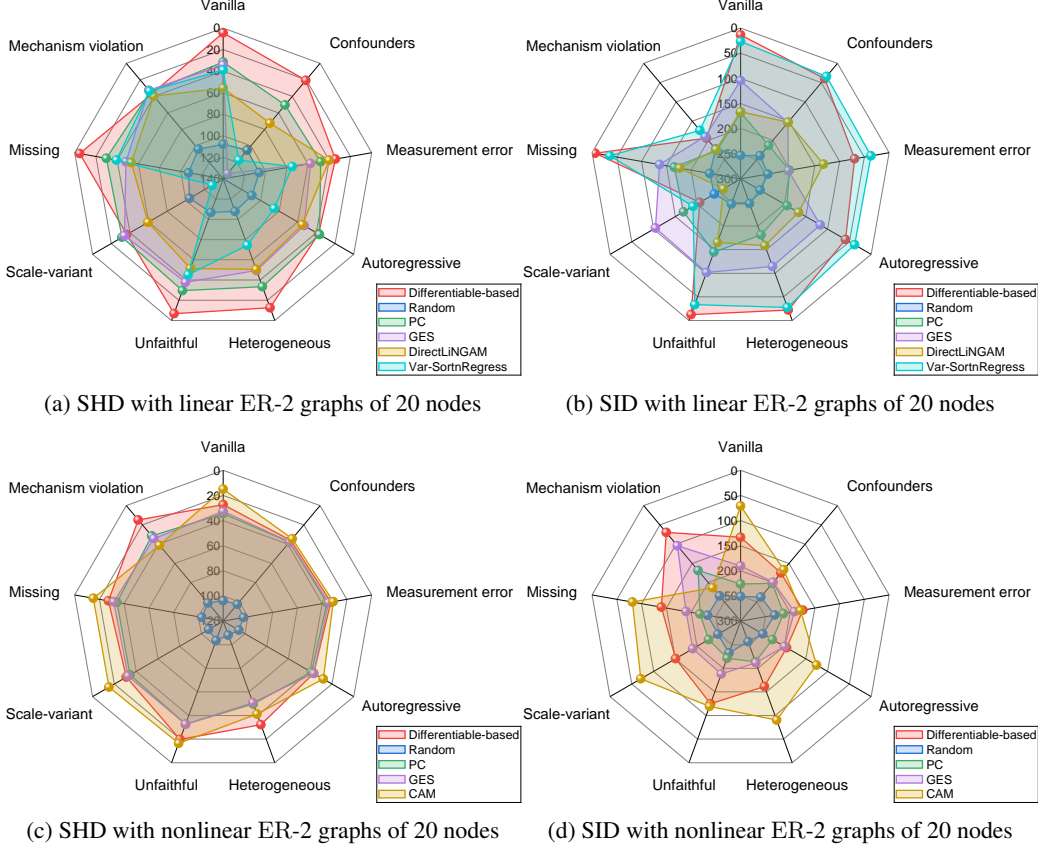


Figure 2: Experimental results under the linear and nonlinear ER-2 graphs of 20 nodes. SHD (the lower the better) and SID (the lower the better) are evaluated over 10 trials. For the differentiable causal discovery method, we present only the optimal results. As the nonlinear settings in Figure 2c and Figure 2d are more favorable to CAM, we conduct a more reasonable evaluation of CAM and differentiable causal discovery under the MLP setting (Section 4.1.1).

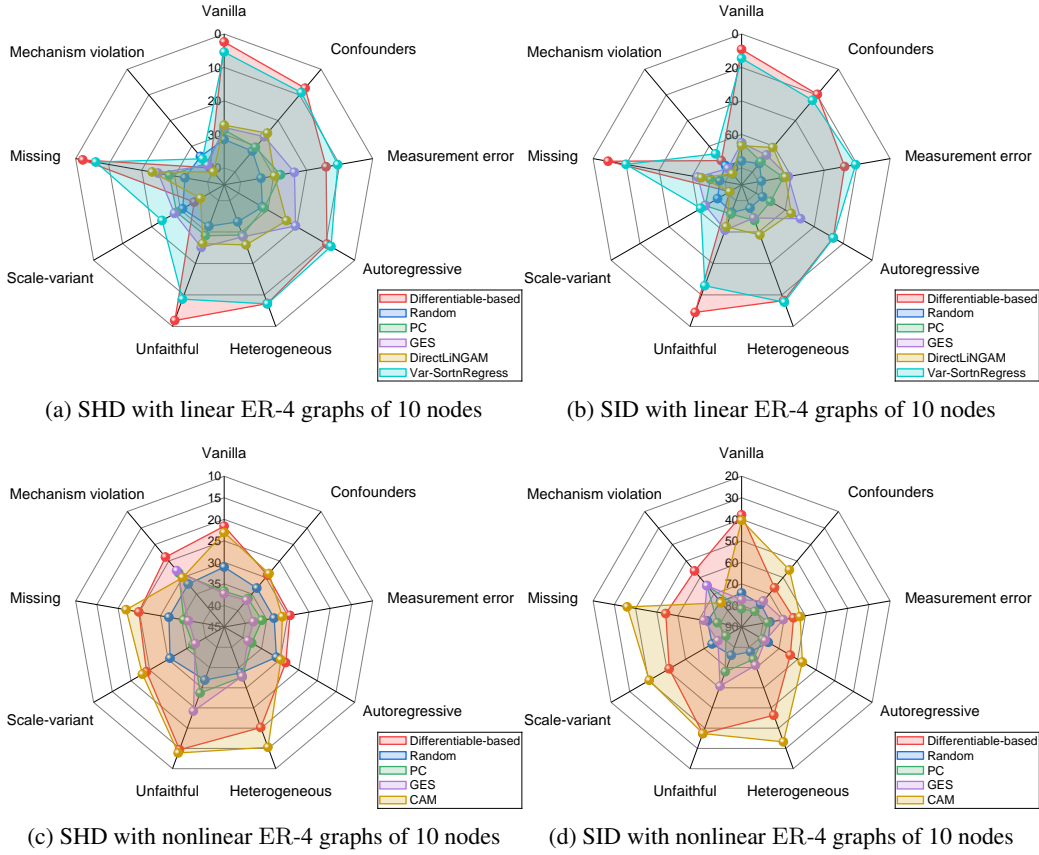


Figure 3: Experimental results under the linear and nonlinear ER-4 graphs of 10 nodes. SHD (the lower the better) and SID (the lower the better) are evaluated over 10 trials. For the differentiable causal discovery method, we present only the optimal results. As the nonlinear settings in Figure 3c and Figure 3d are more favorable to CAM, we conduct a more reasonable evaluation of CAM and differentiable causal discovery under the MLP setting (Section 4.1.1).

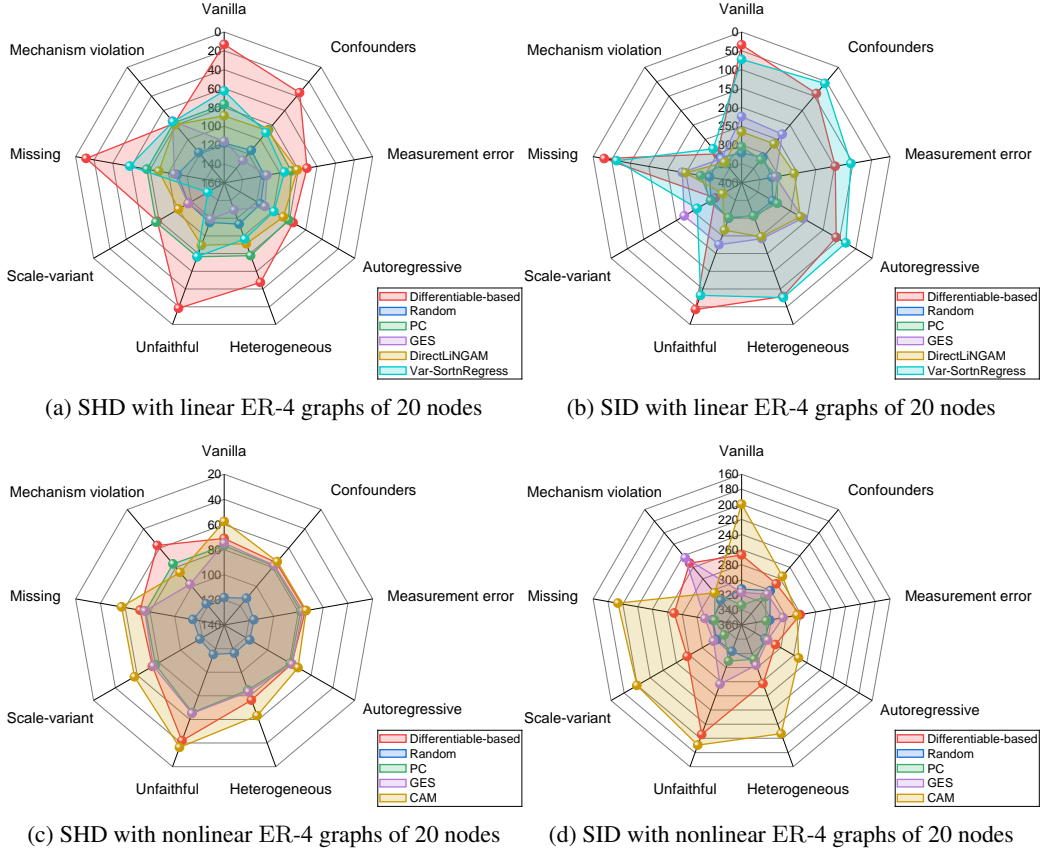


Figure 4: Experimental results under the linear and nonlinear ER-4 graphs of 20 nodes. SHD (the lower the better) and SID (the lower the better) are evaluated over 10 trials. For the differentiable causal discovery method, we present only the optimal results. As the nonlinear settings in Figure 4c and Figure 4d are more favorable to CAM, we conduct a more reasonable evaluation of CAM and differentiable causal discovery under the MLP setting (Section 4.1.1).

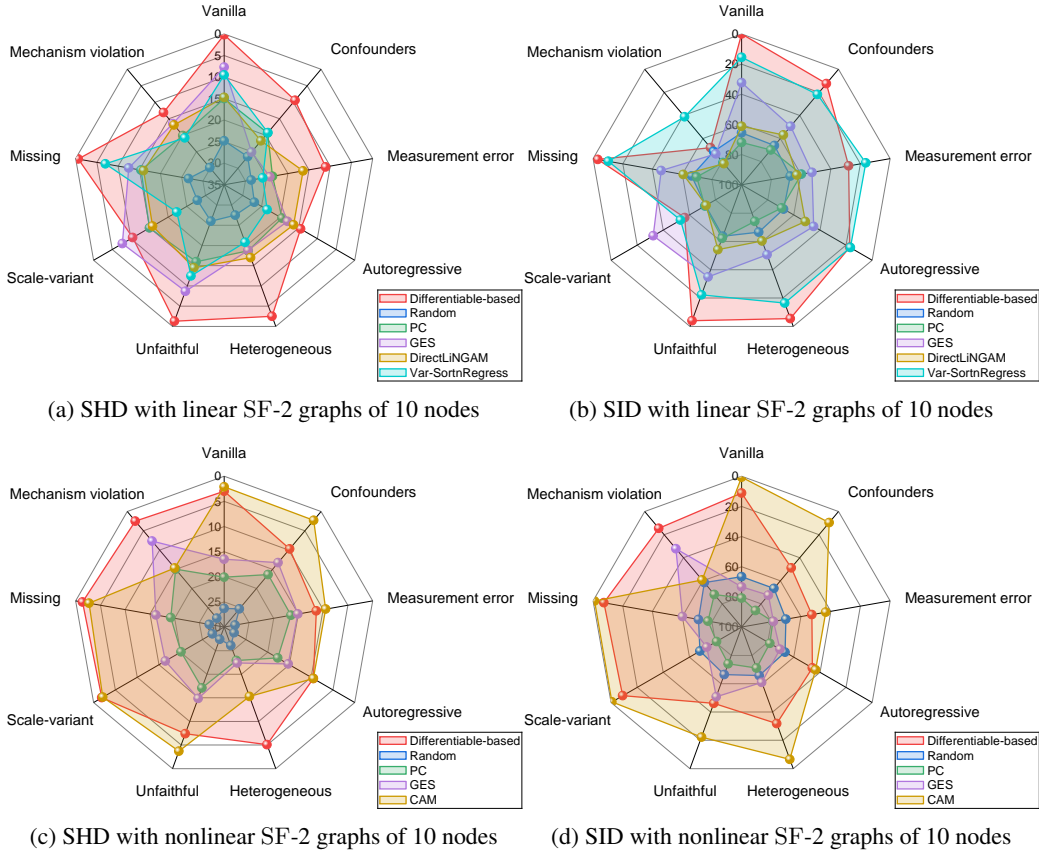


Figure 5: Experimental results under the linear and nonlinear SF-2 graphs of 10 nodes. SHD (the lower the better) and SID (the lower the better) are evaluated over 10 trials. For the differentiable causal discovery method, we present only the optimal results. As the nonlinear settings in Figure 5c and Figure 5d are more favorable to CAM, we conduct a more reasonable evaluation of CAM and differentiable causal discovery under the MLP setting (Section 4.1.1).

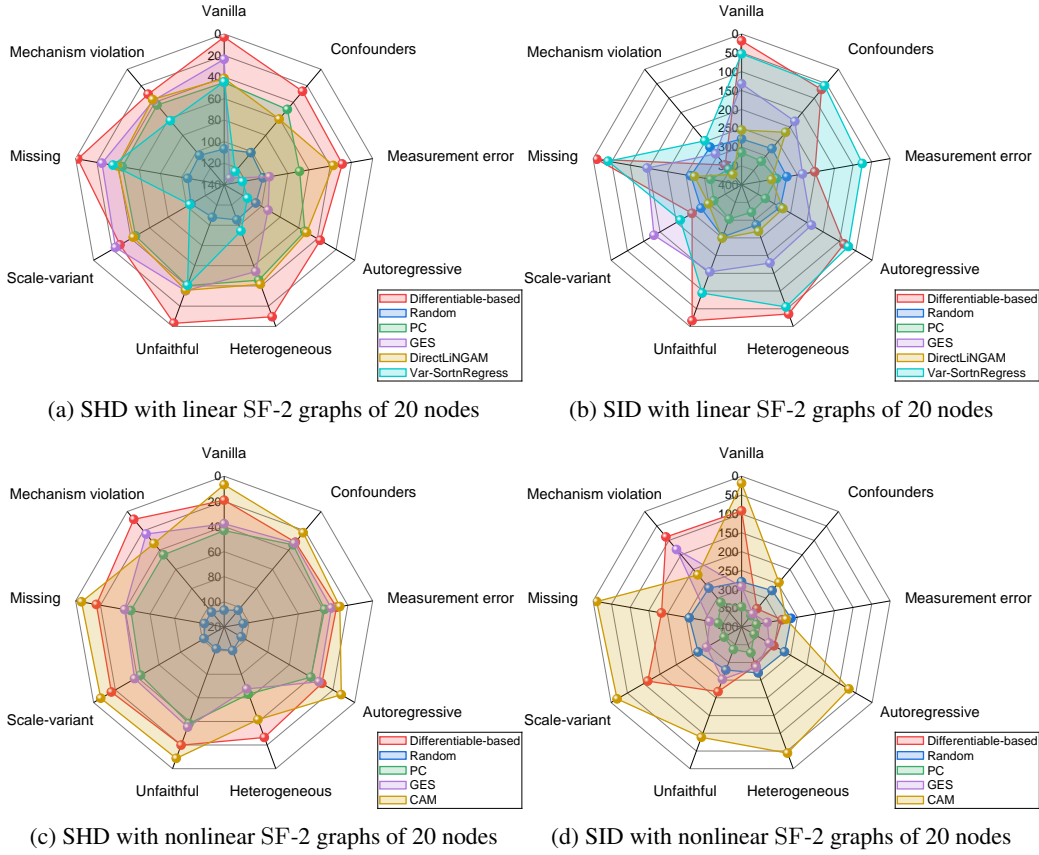


Figure 6: Experimental results under the linear and nonlinear SF-2 graphs of 20 nodes. SHD (the lower the better) and SID (the lower the better) are evaluated over 10 trials. For the differentiable causal discovery method, we present only the optimal results. As the nonlinear settings in Figure 6c and Figure 6d are more favorable to CAM, we conduct a more reasonable evaluation of CAM and differentiable causal discovery under the MLP setting (Section 4.1.1).

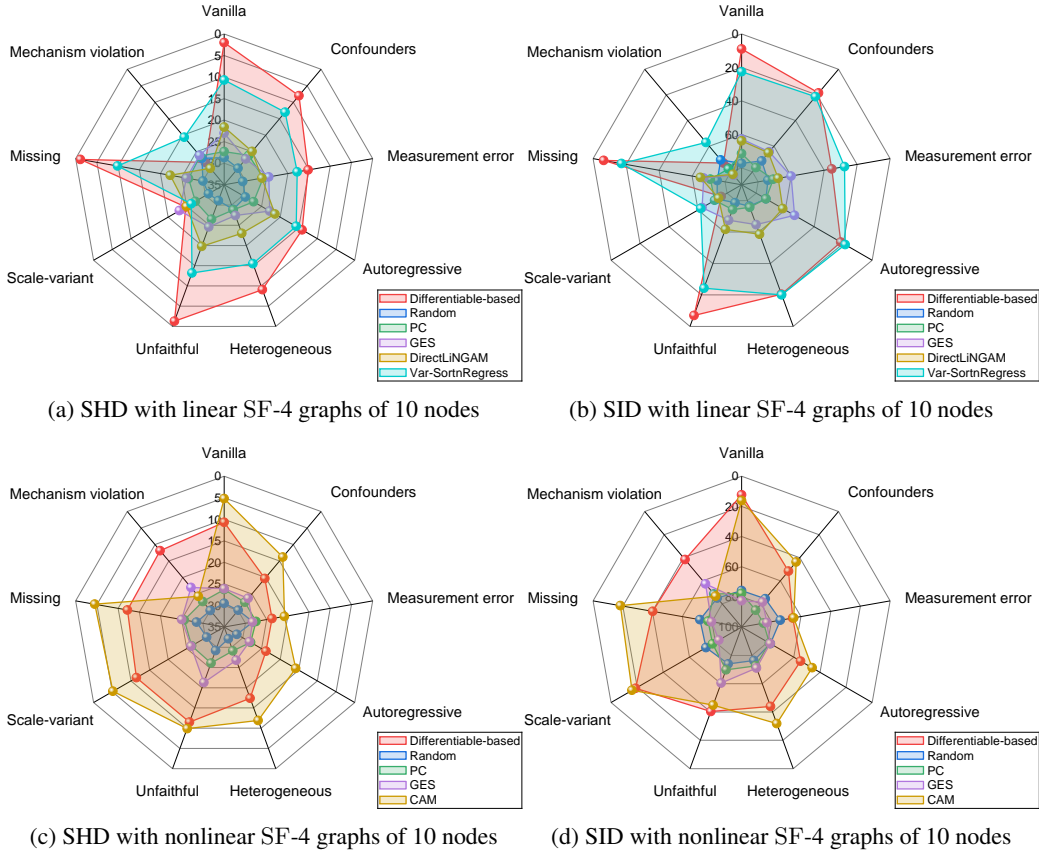


Figure 7: Experimental results under the linear and nonlinear SF-4 graphs of 10 nodes. SHD (the lower the better) and SID (the lower the better) are evaluated over 10 trials. For the differentiable causal discovery method, we present only the optimal results. As the nonlinear settings in Figure 7c and Figure 7d are more favorable to CAM, we conduct a more reasonable evaluation of CAM and differentiable causal discovery under the MLP setting (Section 4.1.1).

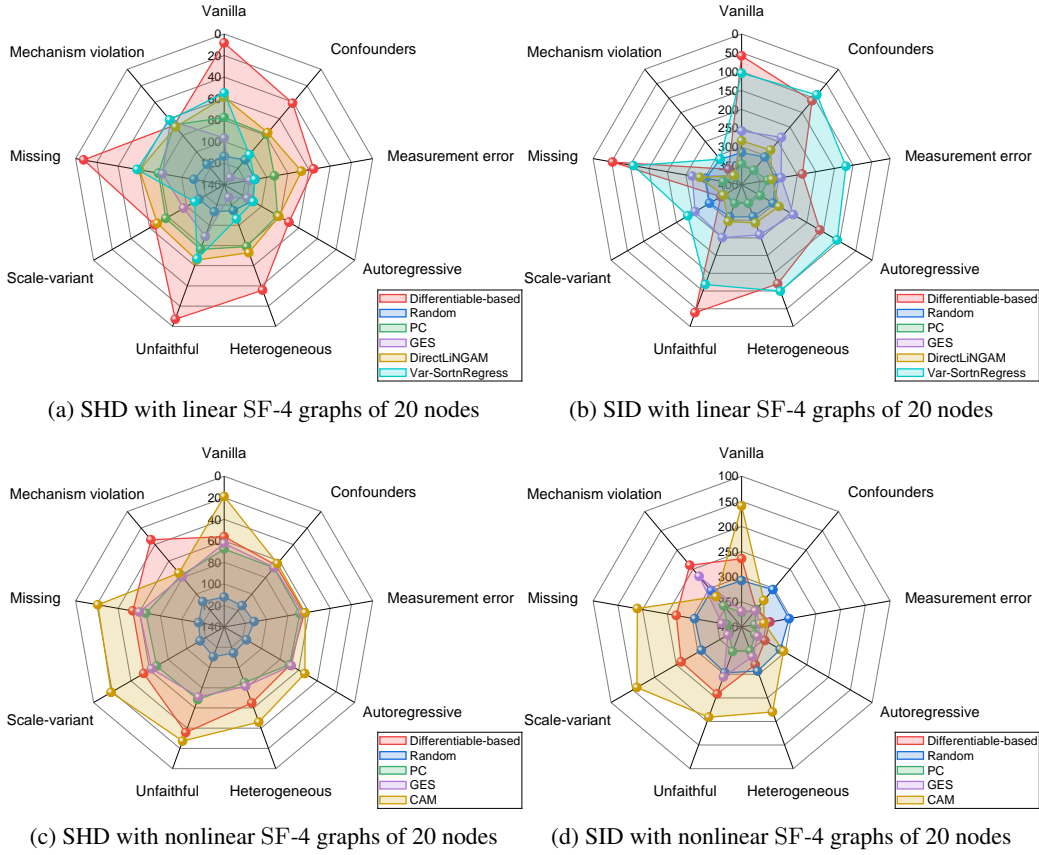


Figure 8: Experimental results under the linear and nonlinear SF-4 graphs of 20 nodes. SHD (the lower the better) and SID (the lower the better) are evaluated over 10 trials. For the differentiable causal discovery method, we present only the optimal results. As the nonlinear settings in Figure 8c and Figure 8d are more favorable to CAM, we conduct a more reasonable evaluation of CAM and differentiable causal discovery under the MLP setting (Section 4.1.1).