

Heteroscedastic Causal Structure Learning

Bao Duong and Thin Nguyen

Applied Artificial Intelligence Institute (A²I²), Deakin University, Australia
 {duongng,thin.nguyen}@deakin.edu.au

Abstract.

Heretofore, learning the directed acyclic graphs (DAGs) that encode the cause-effect relationships embedded in observational data is a computationally challenging problem. A recent trend of studies has shown that it is possible to recover the DAGs with polynomial time complexity under the equal variances assumption. However, this prohibits the heteroscedasticity of the noise, which allows for more flexible modeling capabilities, but at the same time is substantially more challenging to handle. In this study, we tackle the heteroscedastic causal structure learning problem under Gaussian noises. By exploiting the normality of the causal mechanisms, we can recover a valid causal ordering, which can uniquely identify the causal DAG using a series of conditional independence tests. The result is **HOST** (Heteroscedastic causal STructure learning), a simple yet effective causal structure learning algorithm that scales polynomially in both sample size and dimensionality. In addition, via extensive empirical evaluations on a wide range of both controlled and real datasets, we show that the proposed **HOST** method is competitive with state-of-the-art approaches in both the causal order learning and structure learning problems.

1 Introduction

Causal structure learning provides the crucial knowledge to answer not only the statistically descriptive questions of the systems of interest, but also interventional and counterfactual queries. Thus, it is at the heart of interest in many sciences where cause-effect relationships are of utmost concerns, such as bioinformatics [27], econometric [13], and neural sciences [3]. However, this problem poses serious challenges when no prior knowledge is available, not to mention the limitation of its practical solvability with respect to the scale of data.

Constraint-based [30, 31, 6] and score-based methods [5, 22, 21, 37, 38] are conventionally two main approaches in causal structure learning. The former branch leverages a series of statistical tests to eliminate implausible connections between the variables of interest, but the number of tests can grow exponentially with the number of variables. Meanwhile, the latter group assigns dedicated scores to each candidate causal graph based on data fitness, then performs greedy search or continuous optimization for the optimal answer with respect to the defined score, which is a NP-Hard problem since the space of possible DAGs is super-exponentially large in the number of variables [25]. Both of these approaches cannot eliminate all improbable solutions without further experimentations or prior knowledge, so the result is a class of causal graphs that can induce the same observed data, also known as the Markov Equivalent Class (MEC).

Recently, there is a novel line of works that concerns with additive noise models where the noises have equal variances [2, 14, 4,

11, 26, 28]. This assumption allows for polynomial time algorithms that produces unique graphs with provable guarantees in accuracy. Specifically, they put more focus on finding the causal orderings of the causal graph instead of the graph itself, since a correct causal ordering can be used to efficiently trace back to the true causal graph [35, 32].

However, the equal variances assumption limits the modelling capabilities in practice where the noise variances can fluctuate. Specifically, here we exclusively focus on the heteroscedastic causal structure learning problem, wherein the “heteroscedasticity” reflects the fact that the conditional variance of each variable given its causes is non-constant and depends on the causes, in contrary with existing additive noise models where it is constant across all variables in the equal variances assumption aforementioned, or is constant for each variable in the unequal variances models discussed in [2, 11].

While there have been progresses on this model [16, 36, 33, 14], theoretical identifiability and methods are only specifically studied for the case of two variables. As a result, they often propose to extend to more than two variables by first recovering the skeleton of the graph using conventional methods, then orient the undirected edges with developed methods. This approach severely relies on the quality of the skeleton recovery algorithm, which may scale non-polynomially with dimensionality, while does not ensure acyclicity. A few ordering-based approaches can address the heteroscedasticity, but under very restrictive conditions, e.g., constant expected noise variances [11] or positive-valued noises in the multiplicative noise models [20].

Present study. For those reasons, in this work we propose to handle the generic heteroscedastic causal structure learning problem from the causal ordering standpoint. For simplicity, we first consider Gaussian noises. We introduce **HOST**¹ (Heteroscedastic causal STructure learning), a method that operates in polynomial time and produces unique graphs that are guaranteed to be acyclic. More specifically, we exploit the conditional normality of each variable given its ancestral sets and order the variables based on their conditional normality statistics. Once a causal order is retrieved, an array of conditional independence tests is deployed to restore the causal graph. To the best of our knowledge, **HOST** is the first polynomial-time method that handles generic heteroscedasticity.

Contributions. Our key contributions in this study are summarized as follows:

- We propose to tackle the heteroscedastic causal structure learning problem under Gaussian noises by causal ordering. This is the

¹ Source code is available at <https://github.com/baosws/HOST>.

first attempt to handle the heteroscedasticity in causal models with polynomial-time, to the best of our knowledge.

- We present **HOST**, a causal structure learning method for heteroscedastic Gaussian noise models that scales polynomially in both sample size and dimensionality and finds unique acyclic DAGs, as well as being invariant to data scaling.
- We demonstrate the effectiveness of the proposed **HOST** method on a broad range of both synthetic and real data under multiple crucial aspects. The numerical evaluations confirm the quality of our method against state-of-the-arts on both the tasks of causal ordering and structure learning.

2 Background

Let $X = (X_1, \dots, X_d)$ be the d -dimensional random vector of interest, and $x = (x_1, \dots, x_d) \in \mathbb{R}^d$ be an observation of X . We use subscript indices for dimensions and superscript indices for samples. For instance, $x_i^{(k)}$ indicates the k -th sample of x_i . $P(\cdot)$ represents distributions and $p(\cdot)$ represents probability density functions.

A causal structure describes the cause-effect amongst all variables and can be represented by a directed acyclic graph (DAG) $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ where each vertex $i \in \mathcal{V} = [1, \dots, d]$ represents a random variable X_i , and each edge $(j \rightarrow i) \in \mathcal{E}$ indicates that X_j is a direct cause of X_i . In this graph, the *parents* of a variable is defined as the set of its direct causes, i.e., $\text{pa}_i^{\mathcal{G}} := \{j \in \mathcal{V} \mid (j \rightarrow i) \in \mathcal{E}\}$. Similarly, we denote $\text{nd}_i^{\mathcal{G}}$ as the non-descendants set of X_i (excluding itself). Whenever it is clear from context, we drop the superscript \mathcal{G} to reduce notational clutter.

2.1 Heteroscedastic Causal Models

We consider the Structural Causal Model as follows

Definition 1. (Heteroscedastic Causal Model (HCM)). In a HCM, each variable is generated by

$$X_i := \mu_i(X_{\text{pa}_i}) + \sigma_i(X_{\text{pa}_i})E_i \quad (1)$$

where $E \sim \mathcal{N}(\mathbf{0}; \mathbf{I}_d)$ is the exogenous Gaussian noise vector.

This model is also referred to as Location-Scale Noise Model (LSNM) [14] or Heteroscedastic Noise Model (HNM) [33] (except that in HNM, σ models the conditional mean absolute deviation instead of the conditional standard deviation).

In addition, for ease of analysis, we suppose the following regularity constraints on the functions governing the causal processes:

Assumption 2. (Regularity conditions). We assume the following for all $i = 1, \dots, d$:

- μ_i and σ_i are deterministic and differentiable.
- $\sigma_i > 0$.
- $\frac{\partial \mu_i}{\partial X_j} \neq 0$ and $\frac{\partial \sigma_i}{\partial X_j} \neq 0$ (i.e., they are non-zero somewhere in their domains) for all $j \in \text{pa}_i$.

Assumption 2 outlines the conditions required to eliminate any trivial degenerate cases. To be more specific, the first condition favors deterministic functions over stochastic functions so the source of randomness is explicitly captured in the noise variables. Additionally, we restrict the functions to be differentiable for modelling simplicity. The second condition ensures that the conditional variance of a variable depends on its parent variables, thereby guaranteeing the

heteroscedasticity of the model. Lastly, the third condition stipulates that the functions are non-constant with respect to any parent $j \in \text{pa}_i$.

Following this model, it is clear that $P(X_i \mid X_{\text{pa}_i}) = \mathcal{N}(\mu_i(X_{\text{pa}_i}), \sigma_i^2(X_{\text{pa}_i}))$, where $\mu_i(x_{\text{pa}_i})$ and $\sigma_i^2(x_{\text{pa}_i})$ respectively models the conditional mean and variance of each variable given its parents.

This model nicely follows the stable causal mechanism postulation [15] in the sense that the causal mechanism, i.e., the conditional distribution of an effect given its causes, is simple and easy to be described, as it is always a canonical normal distribution. However, this is not enough for identifiability since the vice versa may not hold, i.e., a simple conditional distribution does not make it a causal mechanism. Therefore, to ensure the HCM model is identifiable, we assume that normality is only achieved for the sufficient parents set, i.e., a superset of the parents of a node. More formally:

Assumption 3. Let $C \subseteq \text{nd}_i$. Denote $X_i \mid X_C \sim \mathcal{N}$ when $\frac{X_i - \mathbb{E}[X_i \mid X_C]}{\sqrt{\mathbb{V}[X_i \mid X_C]}} \sim \mathcal{N}(0, 1)$. We assume that

$$X_i \mid X_C \sim \mathcal{N} \Leftrightarrow \text{pa}_i \subseteq C \quad (2)$$

Assumption 3 can be explained as follows: At first, each noise variable E_i is a Gaussian variable, and when combined with the parents through the structural assignment $X_i := \mu_i(X_{\text{pa}_i}) + \sigma_i(X_{\text{pa}_i})E_i$, the normality is disrupted. Essentially, this assumption implies that we can only recover normality by conditioning on all parent variables. The reason for this is that, in certain cases, the influences from the parents may not affect the normality of E_i , resulting in X_i being a Gaussian variable and potentially being incorrectly detected as a node without any parent. In brief, the aim of this assumption is to exclude such cases.

We now argue that this is not a strict condition and can be achieved in general.

Let us condition on a non-descendant set C and define:

$$U_i = \frac{X_i - \mathbb{E}[X_i \mid X_C]}{\sqrt{\mathbb{V}[X_i \mid X_C]}} \quad (3)$$

$$= \underbrace{\frac{\mu_i(X_{\text{pa}_i}) - \mathbb{E}[X_i \mid X_C]}{\sqrt{\mathbb{V}[X_i \mid X_C]}}}_{a_i(X_C)} + \underbrace{\frac{\sigma_i(X_{\text{pa}_i})}{\sqrt{\mathbb{V}[X_i \mid X_C]}}}_{b_i(X_C)} E_i \quad (4)$$

Assumption 3 implicitly excludes all the functions $\mu_i(\cdot)$ and $\sigma_i(\cdot)$ such that $U_i \sim \mathcal{N}(0, 1)$ when $\text{pa}_i \not\subseteq C$.

Given $U_i, E_i \sim \mathcal{N}(0, 1)$, a trivial condition for Eqn. (4) to happen is when $b_i(X_C)$ is a constant t and $a_i(X_C) \sim \mathcal{N}(0, 1 - t^2)$. Here t is constrained to $[0, 1]$.

When $\text{pa}_i \not\subseteq C$ there exists $j \in \text{pa}_i \setminus C$. Since $\mathbb{V}[X_i \mid X_C]$ is not a function of X_j , differentiating $b_i(X_C)$ with respect to X_j gives

$$0 = \frac{1}{\sqrt{\mathbb{V}[X_i \mid X_C]}} \frac{\partial \sigma_i}{\partial X_j} \quad (5)$$

which contradicts Assumption 2. On the other hand, if a_i is constant then it must be zero since $\mathbb{E}[a_i] = \mathbb{E}[U_i] = 0$. This is precisely when $b_i \equiv 1$. In other words, neither $a(X_{\text{pa}_i})$ nor $b(X_{\text{pa}_i})$ can be constant.

Now for all positive integers n , the n -th order moments of U_i and

Algorithm 1 Algorithm for recovering the latents.

Input: Dataset $\mathcal{D} = \left\{ \left(x_C^{(k)}, x_i^{(k)} \right) \right\}_{i=1}^n$ and additional hyperparameters.

Output: The latent variable U_i .

- 1: Initialize two neural networks $\eta_1(X_C; \theta)$ and $\ln(-\eta_2(X_C; \theta))$ with scalar outputs.
- 2: Find the optimal parameter θ^* using, e.g., gradient ascent, by maximizing

$$\theta^* = \operatorname{argmax}_{\theta \in \Theta} \mathbb{E}_{\mathcal{D}} \left[\eta_1(\cdot) X_i + \eta_2(\cdot) X_i^2 + \frac{\eta_1^2(\cdot)}{4\eta_2(\cdot)} + \frac{1}{2} \ln(-2\eta_2(\cdot)) \right]$$

- 3: **return** $U_i = \frac{X_i - \hat{\mu}_i(X_C)}{\hat{\sigma}_i(X_C)}$, where

$$\hat{\mu}_i(X_C) = -\frac{\eta_1(X_C; \theta^*)}{2\eta_2(X_C; \theta^*)}, \text{ and}$$

$$\hat{\sigma}_i(X_C) = \frac{1}{\sqrt{-2\eta_2(X_C; \theta^*)}}$$

E_i must match since they belong to the same distribution:

$$\mathbb{E}[E_i^n] = \mathbb{E}[(a_i + b_i E_i)^n] \quad (6)$$

$$= \mathbb{E} \left[\sum_{k=0}^n \binom{n}{k} a_i^{n-k} b_i^k E_i^k \right] \quad (7)$$

$$= \sum_{k=0}^n \binom{n}{k} \mathbb{E}[a_i^{n-k} b_i^k] \mathbb{E}[E_i^k] \quad (8)$$

Since $\mathbb{E}[E_i^k]$ is a known constant for all positive integers k , if exist, three functions $p(X_C)$, $a_i(X_C)$, $b_i(X_C)$ must be *inextricably interwoven* because the term $\mathbb{E}_{p(X_C)}[a_i^{n-k} b_i^k]$ must satisfy Eqn. (8) for all positive integers n . When limited to non-constant differentiable functions, the solution space of these functions is even more severely limited. Therefore, we expect Assumption 3 to hold in general.

Interestingly, this assumption implies a faithfulness consequence between the DAG \mathcal{G} and the joint distribution of the data, which is characterized by the following Corollary:

Corollary 4. $X_i \not\perp\!\!\!\perp X_j \mid X_{\text{pa}_i \setminus j}$ for all $j \in \text{pa}_i$.

Proof. Take any $j \in \text{pa}_i$. If $X_i \perp\!\!\!\perp X_j \mid X_{\text{pa}_i \setminus j}$ then $P(X_i \mid X_{\text{pa}_i \setminus j}) = P(X_i \mid X_{\text{pa}_i \setminus j}, X_j) = P(X_i \mid X_{\text{pa}_i}) = \mathcal{N}(\cdot)$, which contradicts Assumption 3 with $C = \text{pa}_i \setminus j$. \square

2.2 Causal Ordering

A valid causal order is an arrangement of the vertices in \mathcal{G} so that a cause is always positioned before all of its effects. Given a DAG \mathcal{G} , a valid causal ordering of \mathcal{G} is a permutation $\pi = [\pi_1, \dots, \pi_d]$ of the sequence $[1, \dots, d]$ such that $\pi_j \in \text{nd}_{\pi_i}^{\mathcal{G}}$ for all $j < i$, or equivalently, $\text{pa}_i \subseteq \pi_{<i} \subseteq \text{nd}_i \forall i$.

The causal order is of great interest because learning the causal structure from a causal order not only eases the necessary of the acyclicity constraint of the learned directed graph, but it can also be performed in polynomial runtime. Specifically, under our setting, the causal structure can be uniquely identified from any valid causal order by the following Lemma:

Algorithm 2 Causal Ordering Algorithm.

Input: Dataset $\mathcal{D} = \left\{ x^{(k)} \right\}_{k=1}^n \in \mathbb{R}^{n \times d}$, tolerance level $\epsilon \geq 0$, and additional hyperparameters.

Output: A causal order π .

- 1: $\pi \leftarrow []$.
 - 2: **while** $|\pi| < d$ **do**
 - 3: **for** $i \in \mathcal{V} \setminus \pi$ **do**
 - 4: Use Algorithm 1 to extract U_i from $\left\{ \left(x_{\pi}^{(k)}, x_i^{(k)} \right) \right\}_{k=1}^n$. \triangleright Sec. 3.2.1.
 - 5: $w_i \leftarrow W(U_i)$, where W is the Shapiro-Wilk statistics. \triangleright Sec. 3.2.2.
 - 6: **end for**
 - 7: $w^* \leftarrow \operatorname{argmax}_{i \in \mathcal{V} \setminus \pi} w_i$.
 - 8: $L \leftarrow \{i \in \mathcal{V} \setminus \pi \mid w^* - w_i \leq \epsilon\}$. \triangleright Sec. 3.2.3.
 - 9: Sort L decreasing by w .
 - 10: Concatenate L after π .
 - 11: **end while**
 - 12: **return** π .
-

Algorithm 3 HOST: Heteroscedastic Ordering-based causal Structure learning Algorithm.

Input: Dataset $\mathcal{D} = \left\{ x^{(k)} \right\}_{k=1}^n \in \mathbb{R}^{n \times d}$, a conditional independence test CI, significance level $\alpha \in [0, 1]$, and additional hyperparameters.

CI(X, Y, Z) should return the p -value for the null hypothesis $\mathcal{H}_0 : X \perp\!\!\!\perp Y \mid Z$.

Output: The causal DAG \mathcal{G} .

- 1: Use Algorithm 2 to find a causal order π . \triangleright Sec. 3.1.
 - 2: Initialize empty DAG \mathcal{G} with vertices set \mathcal{V} .
 - 3: **for** $j < i$ **do**
 - 4: **if** CI($X_{\pi_j}, X_{\pi_i}, X_{\pi_{<i} \setminus \pi_j}$) $< \alpha$ **then**
 - 5: Add edge $(\pi_j \rightarrow \pi_i)$ to \mathcal{G} . \triangleright Sec. 3.3.
 - 6: **end if**
 - 7: **end for**
 - 8: **return** \mathcal{G} .
-

Lemma 5. Given a causal order π and joint distribution P_X consistent with the true DAG $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ following Assumptions 2 and 3. Then, for all $j < i$,

$$(\pi_j \rightarrow \pi_i) \in \mathcal{E} \Leftrightarrow X_{\pi_j} \not\perp\!\!\!\perp_{P_X} X_{\pi_i} \mid X_{\pi_{<i} \setminus \pi_j} \quad (9)$$

where $< i$ represents the set of indices that come before i , specifically $1, 2, \dots, i-1$.

Proof. The \Rightarrow direction follows directly from Corollary 4. We now prove the \Leftarrow direction.

Suppose $\pi_j \in \pi_{<i} \setminus \text{pa}_{\pi_i}$. Since $\text{pa}_{\pi_i} \subseteq \pi_{<i}$, by the Causal Markov condition we have

$$P(X_{\pi_i} \mid X_{\text{pa}_{\pi_i}}) = P(X_{\pi_i} \mid X_{\text{pa}_{\pi_i}}, X_{\pi_{<i} \setminus \text{pa}_{\pi_i}} \setminus \pi_j) \quad (10)$$

$$= P(X_{\pi_i} \mid X_{\pi_{<i} \setminus \pi_j}), \text{ and} \quad (11)$$

$$P(X_{\pi_i} \mid X_{\text{pa}_{\pi_i}}) = P(X_{\pi_i} \mid X_{\text{pa}_{\pi_i}}, X_{\pi_{<i} \setminus \text{pa}_{\pi_i}} \setminus \pi_j, X_{\pi_j}) \quad (12)$$

$$= P(X_{\pi_i} \mid X_{\pi_{<i} \setminus \pi_j}, X_{\pi_j}) \quad (13)$$

Therefore, $X_{\pi_i} \perp\!\!\!\perp X_{\pi_j} \mid X_{\pi_{<i} \setminus \pi_j}$, which completes the proof. \square

Additionally, the DAG recovered this way is unique by construction.

This property suggests performing a series of conditional independence (CI) tests to recover the DAG with only $\mathcal{O}(d^2)$ tests, compared with an exponential number of tests in the classical PC algorithm [30].

3 HOST: Heteroscedastic Ordering-based Causal Structure Learning

Lemma 5 has established that a causal ordering is sufficient to recover the underlying DAG. Therefore, what remains is how to retrieve such orderings. In this section, we explain in details the causal ordering procedure and related technical considerations.

3.1 Causal Order Identification

Identifying the precise causal order is usually a challenging computational task. Nonetheless, certain assumptions can make the problem more manageable. Typically, the primary source of computational advantage in polynomial algorithms is the ability to identify a source or sink node in polynomial time. This is usually accomplished by assuming the “equal variance” condition in additive noise models ([4], [11]), which allows a source node to be effectively identified among the remaining variables since it has the smallest conditional variance. However, in our study, this condition is not applicable due to the heteroscedastic nature of the model under consideration. Instead, we rely on the normality of the residuals to detect a source node, based on the assumption that only source nodes will have Gaussian residuals (Assumption 3). Therefore, at each step, we can efficiently extract the residuals and test their normality in polynomial time. By repeating this process until all variables have been examined, we obtain a polynomial time algorithm for causal ordering.

Indeed, here we show that a valid causal order is fully identifiable under Assumption 3. To begin with, the following Lemma says that by conditioning on an ancestral set, one can identify the subsequent source nodes in the reduced graph where the ancestral set is removed, wherein the case of empty ancestral set allows detecting source nodes in the original graph.

Lemma 6. *Under assumptions 2 and 3, let C be an ancestral set in \mathcal{G} (i.e., C satisfies: $C \subseteq \bigcap_{i \notin C} \text{nd}_i^{\mathcal{G}}$). Define $\mathcal{G} \setminus C$ as a reduced DAG with vertices set $\mathcal{V} \setminus C$ and edges set $\{(j \rightarrow i) \in \mathcal{E} \mid j, i \in \mathcal{V} \setminus C\}$. For $i \notin C$, if $X_i \mid X_C \sim \mathcal{N}$ then i is a source node in $\mathcal{G} \setminus C$.*

Proof. Since $X_i \mid X_C \sim \mathcal{N}$, we have $\text{pa}_i^{\mathcal{G}} \subseteq C$ by Assumption 3. Thus, $\text{pa}_i^{\mathcal{G} \setminus C} = \emptyset$, i.e., i is a source node in $\mathcal{G} \setminus C$. \square

Therefore, by leveraging Lemma 6 we can derive a causal ordering algorithm specialized for HCM models. Algorithm 2 demonstrates the key steps of our causal ordering algorithm, in which we iteratively employ normality testing subroutines to detect new source nodes every step. The latents extraction and normality testing components are discussed more thoroughly in the following subsection.

3.2 Identifying Source Nodes via Normality Statistics

We have shown that testing for $X_i \mid X_C \sim \mathcal{N}$ is essential to the causal ordering procedure. More precisely, given an ancestral set C and a variable X_i , we need to test if $\frac{X_i - \mathbb{E}[X_i \mid X_C]}{\sqrt{\mathbb{V}[X_i \mid X_C]}} \sim \mathcal{N}(0, 1)$. Again, let us denote $\frac{X_i - \mathbb{E}[X_i \mid X_C]}{\sqrt{\mathbb{V}[X_i \mid X_C]}}$ by U_i , which we term as the “latent variable” as argued. Then, the problem is reduced to the conventional normality testing problem with the null hypothesis $\mathcal{H}_0 : U_i \sim \mathcal{N}(0, 1)$ and alternative hypothesis $\mathcal{H}_1 : U_i \not\sim \mathcal{N}(0, 1)$. We next show how to extract the latents and test for their normality in our framework.

3.2.1 Extracting the Latents

To excerpt U_i from X_i and X_C , it is natural to estimate $\mathbb{E}[X_i \mid X_C]$ and $\mathbb{V}[X_i \mid X_C]$ and plug the estimates into the expression of U_i . These estimations, which are nonlinear regression problems, can be done separately, for example, as in [33]. However, that means several regression stages must be performed sequentially, which possibly become computationally involved.

Therefore, we prefer to jointly estimate the conditional expectation and standard deviation instead. A naive and common approach for this would be directly parametrizing the conditional mean and standard deviation using neural networks, such as in [16], e.g.,

$$t_i(X_C; \theta) \approx \mathbb{E}[X_i \mid X_C] \quad (14)$$

$$s_i(X_C; \theta) \approx \sqrt{\mathbb{V}[X_i \mid X_C]} \quad (15)$$

under the model $X_i \sim \mathcal{N}(t_i(X_C; \theta), s_i^2(X_C; \theta))$. The optimal parameters θ^* can be found by maximizing the Gaussian log likelihood:

$$\ln p_{\theta}(X_i \mid X_C) = \ln \mathcal{N}(X_i; t_i(X_C; \theta), s_i^2(X_C; \theta)) \quad (16)$$

$$= -\frac{(X_i - t_i(X_C; \theta))^2}{2s_i^2(X_C; \theta)} - \ln s_i(X_C; \theta) + \text{const} \quad (17)$$

$$\mathcal{L}(\theta) = \mathbb{E}[\ln p_{\theta}(X_i \mid X_C)] \quad (18)$$

$$\theta^* = \underset{\theta}{\text{argmax}} \mathcal{L}(\theta) \quad (19)$$

However, Eqn. (17) is not a jointly concave objective function w.r.t. t_i and s_i . This is because its Hessian matrix is not always negative-definite, since

$$\frac{\partial^2 \ln p_{\theta}(X_i \mid X_C)}{\partial s_i^2} = \frac{-3(X_i - t_i(X_C; \theta))^2}{s_i^4(X_C; \theta)} + \frac{1}{s_i^2(X_C; \theta)} \quad (20)$$

can be positive when $s_i^2(X_C; \theta) > 3(X_i - t_i(X_C; \theta))^2$. Hence, jointly optimizing for them both via—e.g., common gradient-based solvers—does not guarantee the global maxima, even if the neural networks have infinite capacities.

To mitigate this issue, following [18, 14], we adopt the natural parametrization of the Gaussian distribution. Particularly, we parametrize $\mathcal{N}(\mu, \sigma^2)$ using two natural parameters η_1, η_2 such that $\mu = -\frac{\eta_1}{2\eta_2}$ and $\sigma^2 = -\frac{1}{2\eta_2}$. The log likelihood now becomes

$$\ln p_{\theta}(X_i \mid X_C) = \eta_1 X_i + \eta_2 X_i^2 + \frac{\eta_1^2}{4\eta_2} + \frac{1}{2} \ln(-2\eta_2) + \text{const} \quad (21)$$

where $\eta_1(X_C; \theta)$ and $\eta_2(X_C; \theta)$ are functions of X_C that are parametrized by θ , and $\eta_2(X_C; \theta) < 0$. One can then show that the objective function (21) is now jointly concave in both η_1 and η_2 [18], which makes gradient-based solutions to the maximum likelihood objective consistent.

To adapt to arbitrarily nonlinear relationships, η_1 and η_2 can be parametrized with neural networks. More specifically, we can parametrize $\eta_1(X_C; \theta) : \mathbb{R}^{|C| \times |\Theta|} \rightarrow \mathbb{R}$ as a simple Multiple Layer Perceptron (MLP) with parameters space Θ . Regarding η_2 , since it must be negative, we should instead parametrize $\ln(-2\eta_2(X_C; \theta)) : \mathbb{R}^{|C| \times |\Theta|} \rightarrow \mathbb{R}$ as another MLP. However, for merely linear maps η_1 and $\ln(-2\eta_2)$, the conditional distribution $P_{\theta}(X_i \mid X_C)$ can still have nonlinear expectations and standard deviations, which already represent a wide class of distributions.

To summarize, the latents extraction subroutine is described in Algorithm 1.

3.2.2 Testing for the Latents' Normality

Having the latents retrieved, we are now in a position to test for their normality to detect source nodes.

Normality testing is a well-studied problem where a broad variety of methods is available. For an overview, see, e.g., [34, 7]. In this study, we particularly employ the well-regarded Shapiro-Wilk test [29] since it has been shown to have a better power against other common alternative approaches [23]. That being said, the considering component of our method is modular and any valid normality testing method can be employed in place of the Shapiro-Wilk test.

Simply put, the test's statistics of the Shapiro-Wilk test, denoted as W , has the range of $[0, 1]$ where the value of one indicates perfect normality and the value of zero suggests strong non-normality.

We can now detect new source nodes using the Shapiro-Wilk test. To proceed, it is intuitive to adopt the whole hypothesis testing procedure for each remaining variable. More specifically, with a significance level α chosen prior to seeing the data, we compute the test statistics W for each variable and their associated p -values, then select those whose p -values less than α .

However, the null distribution of W is complicated and unknown [29], thus computing its p -value requires Monte Carlo simulation, which can be computationally intense. Moreover, in practice there may be no variable with p -value $< \alpha$, making the iterative process unhalting. Hence, we take a slight detour by selecting the variable with the highest W statistics, which does not require p -value calculation and always exists.

3.2.3 Layer Decomposition

For sparse graphs, each iteration may find multiple valid source nodes. If only one is chosen then the others must have their calculations re-done multiple times, which is wasteful of resource, and the chance of miscalculation may arise in the subsequent steps as the ancestral set increases in cardinality. Therefore, we should select as many source nodes as possible in each step to avoid these pitfalls.

To do this, we employ a tolerance threshold ϵ onto the selection of the source nodes based on their W statistics. Particularly, let w^* be the largest W statistics value of the remaining variables in each step, we will include all variables with W within the ϵ -radius of w^* as new source nodes, which form a "layer" similarly to [11].

Further, there is a trade-off between accuracy and runtime when choosing ϵ . If ϵ is too small then the effect of computation reduction is negligible, whereas larger ϵ will allow non-source nodes into the layer. In our implementation, we choose a value of 10^{-4} , which is relatively small in comparison with the range $[0, 1]$ of W , since accuracy is preferred in our experiments.

Resultantly, the causal order will be the concatenation of the layers collected after each step. Algorithm 2 summarizes the main steps of the Causal Ordering algorithm with all the technical considerations discussed.

Nevertheless, it is important to mention that our algorithm can still perform effectively even without utilizing the layer decomposition procedure. Furthermore, our layer decomposition only necessitates a single hyperparameter, the tolerance ϵ . If this hyperparameter is set to zero, the feature is disabled, and this could be set as the default behavior for less experienced users of our algorithm.

3.3 DAG Recovery From the Causal Order

As stated in Lemma 5, one can obtain the full causal DAG from any valid causal order π . This is done by first starting with an empty DAG,

then examining every ordered pair of vertices (π_j, π_i) to see if there is a directed edge connecting them with the help of a series of CI tests.

Alternatively, instead of CI tests, one can employ feature selection techniques with lower computational demands if runtime is preferred. For example, GAM (Generalized Additive Model) feature selection is widely used in several ordering-based methods that follow nonlinear additive noise models, e.g., [2, 26, 28]. More specifically, a GAM model is fitted to $(X_{\pi_{<i}}, X_{\pi_i})$, then any significant feature X_{π_j} with $j < i$, evidenced by a small p -value, will add an edge $(\pi_j \rightarrow \pi_i)$ to an initially empty DAG. Refer to [2] for more details on this procedure.

Finally, we put together all the pieces of our **HOST** method in Algorithm 3. The computational complexity of **HOST** is given in Appendix B [9], which is polynomial in both sample size and dimensionality.

4 Theoretical Properties

4.1 Identifiability of HCMs

The identifiability of HCM is given by the following Theorem.

Theorem 4.1. *Under assumptions 2 and 3, HCMs are fully identifiable, i.e., there exists a unique graph \mathcal{G} consistent with data P_X generated according to Eqn. 1.*

Proof sketch. By Lemma 6 a valid source node can be identified at every step, thus by induction a valid causal ordering can be identified. Then by Lemma 5 every true edge can be identified from the causal ordering, rendering the whole DAG identifiable. \square

4.2 Invariance to Scaling and Translation

Another important property of our method is the invariance to scaling and invariance, which is given by the following Theorem.

Theorem 4.2. *The **HOST** algorithm is invariant to scaling and translation.*

Proof sketch. If we scale and translate X_i by a scale $c > 0$ and location d , i.e., $X'_i := cX_i + d$, the new causal model will have all the structural assignments and noise variables unchanged, except for X_i . Instead, we can replace it with $X'_i := \mu'_i(X_{\text{pa}_i}) + \sigma'_i E_i$, where $\mu'_i := \mu_i \times c + d$ and $\sigma'_i := \sigma_i \times c$. In case $c < 0$, the negation can be absorbed to $E'_i := -E_i$, which is still a standard Gaussian noise mutually independent of other noises.

Note that both μ'_i and σ'_i maintain the properties of μ_i and σ_i as described in Assumption 2, namely being deterministic and differentiable, $\sigma'_i > 0$, and are non-constant with respect to any parent variable $j \in \text{pa}_i$.

Additionally, since scaling nor translation does not alter normality, Assumption 3 also applies to the new HCM. Consequently, the newly obtained HCM is also recoverable using our method. \square

This property allows us to standardizing data before applying the **HOST** algorithm without changing the result. This is useful since neural networks, which are employed in our framework, are sensitive to data scale, meaning if the original data scale is too large or too small, the learning process may struggle to converge.

We note that this step is not possible with methods based on the "equal variance" assumption (e.g., EqVar [4] or NPVar [11]), since scaling breaks the equality between the variables' variances, making their algorithms misbehave.

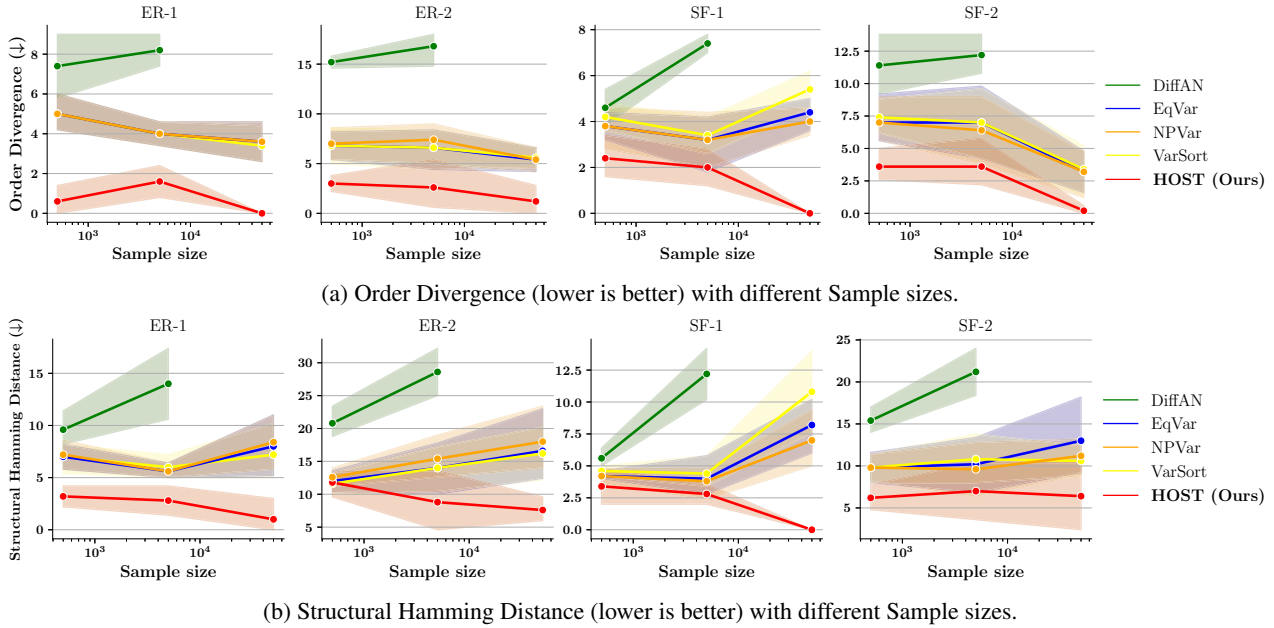


Figure 1. Causal structure learning performance on synthetic data as function of Sample size under Linear parameterization. We fix $d = 10$ and vary the sample size. The proposed method is compared against DiffAN [28], EqVar [4], NPVar [11], and VarSort [24]. Column: Graph type. Shaded areas are 95% confidence intervals over five independent runs. LCIT [8] is used to recover the DAGs from the causal orders. Missing data of DiffAN is due to excessive runtime.

5 Numerical Evaluations

5.1 Experiment Setup

Baselines. We evaluate both the causal ordering and causal structure learning performance of the proposed **HOST** method with competitive baselines in ordering-based causal structure learning, including:

VarSort [24] simply sorts the variables according to their marginal variances based on the observation that an effect usually has higher variance than its causes in the common evaluation practice of structure learning methods on simulated data. This acts as a sanity check for the complexity of our problem.

EqVar [4] models the data with linear relationships and homoscedastic additive noises of equal variances. At each step, the variables with minimum conditional variances given the current ancestral set are chosen as next source nodes.

NPVar [11] extends EqVar by considering nonlinear causal mechanisms using nonparametric regression techniques. It is claimed to be able to handle heteroscedasticity, but only with constant expected noise variances, i.e., $\mathbb{E}[\mathbb{V}[X_i | X_{\text{pa}_i}]]$.

DiffAN [28] considers nonlinear additive noise models with homoscedastic Gaussian noises. Variables with constant partial derivatives of the score function, learned with diffusion models, are selected as sink nodes every step.

Additionally, for completeness, in Appendix C [9] we also compare **HOST** with other popular baselines that are not polynomial-time, including CAM [2], GOLEM [19], and GraN-DAG [17].

Metrics. For assessing the causal ordering accuracy, we employ the Order Divergence measure proposed by [26], which counts the number of edges that are incorrectly ordered:

$$\text{OrderDivergence}(\pi, \mathcal{E}) = \sum_{j < i} \mathbb{1}[(\pi_i \rightarrow \pi_j) \in \mathcal{E}]$$

Regarding causal structure learning, we adopt the common metric of Structural Hamming Distance, which measures the number of edge additions, removals, and reversals to transform the predicted DAG to the ground truth DAG.

In addition, to tackle scenarios where there is an imbalance between missing and extra edges, we also employ the F_1 score and AUC (Area Under the Receiver Operating Characteristic Curve) measures. The performance of our proposed method under these metrics is provided in Appendix C [9], demonstrating the superiority of our proposed **HOST** method over the baselines in terms of both F_1 score and AUC. This finding is consistent with the results obtained using the SHD measure given in the next subsection.

DAG recovery methods. To remove the influence of the choice of the DAG recovery method from the orderings onto the structure learning performance, we use the same algorithm for all methods. Particularly, we consider two options, first is the CI testing approach using the recent state-of-the-art test LCIT [8] since it is generic and linearly scalable in both dimensionality and sample size, along with the GAM feature selection approach which is usually employed in methods based on the equal variances assumption [2, 26, 28]. Both of them use $\alpha = 0.001$.

Synthetic data. We generate data under two typical DAG settings. First is the Erdős-Rényi (ER) graph [10], where edges are independently added with an equal probability, and we control the expected in-degree to be one (ER-1 graphs) or two (ER-2 graphs). Second is the Scale Free (SF) graph [1], with the SF-1 (SF-2) variant being the DAG initially started with one (two) nodes and every subsequent node is added with one (two) random edges from the previously added nodes. We also consider the scenarios with denser graphs (ER-4 and SF-4) in Appendix C [9].

About the functional mechanisms, we generate $\mu_i(\cdot)$ and $\sigma_i(\cdot)$ from the natural parametrizations and consider both linear and nonlinear scenarios. For the linear case, η_1 and $\ln(-2\eta_2)$ are homogeneous

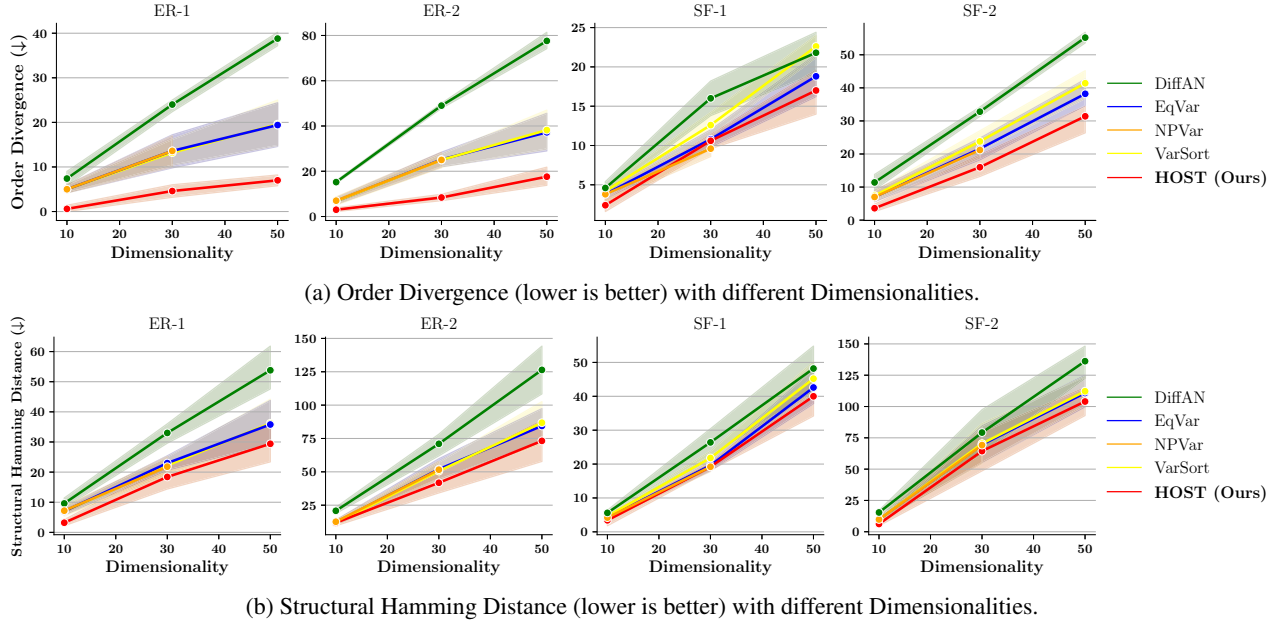


Figure 2. Causal structure learning performance on synthetic data as function of Dimensionality under Linear parameterization. We fix $n = 500$ and vary the dimensionality. The proposed method is compared against DiffAN [28], EqVar [4], NPVar [11], and VarSort [24]. Column: Graph type. Shaded areas are 95% confidence intervals over five independent runs. LCIT [8] is used to recover the DAGs from the causal orders. Missing data of NPVar is due to excessive runtime.

linear maps $a^\top x_{pa_i}$ for each node. Meanwhile, in the nonlinear case, these parameters associated with each parent are chosen from the set $\left\{a^\top x, x^2, \sin(2\pi x), \ln(x - \min(x) + 1), \frac{1}{1+e^{-x}}\right\}$, and then summed afterward, e.g., $\eta_1(x_{pa_i}) = \sum_{j \in pa_i} \eta_1(x_j)$.

We emphasize here that Assumption 3 is not imposed in our data generating process, which means that it is possible for it to be violated in the generated data. However, even without enforcing this assumption, our method still shows robustness compared to the baselines in situations where the model is not well-specified, as evidenced by the empirical results reported in the next subsection.

5.2 Results on Synthetic Data

Here we present the empirical results for the linear parametrization setting with LCIT being the DAG recovery method. The additional experimental results, including the consideration of other evaluation metrics, nonlinear settings, GAM feature selection, denser graphs, as well as the recorded runtimes can be found in Appendix C [9].

Effect of sample size. We first study the influence of sample size to the learning performance of the considering methods by fixing the number of nodes at $d = 10$ and vary the sample size n from 500 to 50,000 (Figure 1). The results suggest that our method is far more effective than the baseline counterparts across all scenarios, especially in the task of causal ordering. Additionally, it can be observed that our method converges to zero error on both metrics as the sample size increases, empirically suggesting its consistency.

Effect of dimensionality. Next, we study the performance variation of all methods when the dimensionality changes. To this end, we fix the sample size at $n = 500$ and vary the number of nodes from 10 to 50 (Figure 2). In this setting, while all methods show the same degradation in performance, our proposed **HOST** method is still the leading performer over all aspects.

Table 1. Causal structure learning performance on real data. We compare the proposed **HOST** with DiffAN [28], EqVar [4], NPVar [11], and VarSort [24] on the Sachs data set [27]. The values are *mean* \pm *standard error* over the same ten independent subsamples of size 700.

	Order Divergence (\downarrow)	Structural Hamming Distance (\downarrow) (CI testing)	(GAM feature selection)
DiffAN	9.4 ± 0.70	15.0 ± 1.05	15.6 ± 0.84
EqVar	7.2 ± 0.42	13.5 ± 0.71	12.0 ± 0.47
NPVar	7.2 ± 0.42	13.5 ± 0.71	12.0 ± 0.47
VarSort	7.2 ± 0.42	13.7 ± 0.95	12.0 ± 0.47
HOST (Ours)	4.4 ± 0.70	13.5 ± 0.97	13.3 ± 0.95

5.3 Results on Real Data

To demonstrate the effectiveness of our method in the real-life setting, we conduct experiments on the well-known benchmark data set Sachs [27], where the ground truth causal network is available. We employ the observational portion of the data set with 853 observations, 11 vertices, and 17 edges in the causal graph.

Table 1 displays the empirical results. Our method achieves the best accuracy in recovering the causal order with an error at only half of that for other methods, while it is competitive with the state-of-the-arts in recovering the causal DAG.

6 Conclusions

This study presents the **HOST** algorithm for identifying the causal structures under heteroscedastic causal models. By exploiting the conditional normalities with the help of normality tests, we devise a simple procedure for recovering causal orderings, which are used to uniquely recover the causal structures. The empirical results on a wide range of synthetic and real data show that **HOST** is able to consistently outperform existing state-of-the-art ordering-based methods in both causal ordering and structure learning.

References

- [1] Albert-László Barabási and Réka Albert, 'Emergence of scaling in random networks', *Science*, **286**, 509–512, (1999).
- [2] Peter Bühlmann, Jonas Peters, and Jan Ernest, 'CAM: Causal additive models, high-dimensional order search and penalized regression', *The Annals of Statistics*, **42**, 2526–2556, (2014).
- [3] Yinan Cao, Christopher Summerfield, Hame Park, Bruno Lucio Gordan, and Christoph Kayser, 'Causal inference in the multisensory brain', *Neuron*, **102**, 1076–1087, (2019).
- [4] Wenyu Chen, Mathias Drton, and Y Samuel Wang, 'On causal discovery with an equal-variance assumption', *Biometrika*, **106**, 973–980, (2019).
- [5] David Maxwell Chickering, 'Optimal structure identification with greedy search', *Journal of Machine Learning Research*, **3**, 507–554, (2002).
- [6] Diego Colombo, Marloes H Maathuis, Markus Kalisch, and Thomas S Richardson, 'Learning high-dimensional directed acyclic graphs with latent and selection variables', *The Annals of Statistics*, 294–321, (2012).
- [7] Keya Rani Das and Rahmatullah Imon, 'A brief review of tests for normality', *American Journal of Theoretical and Applied Statistics*, **5**, 5–12, (2016).
- [8] Bao Duong and Thin Nguyen, 'Conditional independence testing via latent representation learning', in *Proceedings of the IEEE International Conference on Data Mining (ICDM)*, pp. 121–130, (2022).
- [9] Bao Duong and Thin Nguyen, 'Heteroscedastic causal structure learning', *arXiv preprint arXiv:2307.07973*, (2023).
- [10] Paul Erdős and Alfréd Rényi, 'On the evolution of random graphs', *Publications of the Mathematical Institute of the Hungarian Academy of Sciences*, (1960).
- [11] Ming Gao, Yi Ding, and Bryon Aragam, 'A polynomial-time algorithm for learning nonparametric causal graphs', in *Advances in Neural Information Processing Systems*, volume 33, pp. 11599–11611, (2020).
- [12] Asish Ghoshal and Jean Honorio, 'Learning linear structural equation models in polynomial time and sample complexity', in *Proceedings of the International Conference on Artificial Intelligence and Statistics*, pp. 1466–1475, (2018).
- [13] Paul Hünermund and Elias Bareinboim, 'Causal inference and data fusion in econometrics', *The Econometrics Journal*, utad008, (2023).
- [14] Alexander Immer, Christoph Schultheiss, Julia E Vogt, Bernhard Schölkopf, Peter Bühlmann, and Alexander Marx, 'On the identifiability and estimation of causal location-scale noise models', in *Proceedings of the International Conference on Machine Learning*, pp. 14316–14332, PMLR, (2023).
- [15] Dominik Janzing and Bernhard Schölkopf, 'Causal inference using the algorithmic Markov condition', *IEEE Transactions on Information Theory*, **56**, 5168–5194, (2010).
- [16] Ilyes Khemakhem, Ricardo Monti, Robert Leech, and Aapo Hyvarinen, 'Causal autoregressive flows', in *Proceedings of the International Conference on Artificial Intelligence and Statistics*, pp. 3520–3528, (2021).
- [17] Sébastien Lachapelle, Philippe Brouillard, Tristan Deleu, and Simon Lacoste-Julien, 'Gradient-based neural DAG learning', in *Proceedings of the International Conference on Learning Representations*, (2020).
- [18] Quoc V Le, Alex J Smola, and Stéphane Canu, 'Heteroscedastic gaussian process regression', in *Proceedings of the International Conference on Machine Learning*, pp. 489–496, (2005).
- [19] Ignavier Ng, AmirEmad Ghassami, and Kun Zhang, 'On the role of sparsity and dag constraints for learning linear DAGs', *Advances in Neural Information Processing Systems*, 17943–17954, (2020).
- [20] Goutham Rajendran, Bohdan Kivva, Ming Gao, and Bryon Aragam, 'Structure learning in polynomial time: Greedy algorithms, Bregman information, and exponential families', in *Advances in Neural Information Processing Systems*, volume 34, pp. 18660–18672, (2021).
- [21] Joseph Ramsey, Madelyn Glymour, Ruben Sanchez-Romero, and Clark Glymour, 'A million variables and more: the Fast Greedy Equivalence Search algorithm for learning high-dimensional graphical causal models, with an application to functional magnetic resonance images', *Proceedings of the International Journal of Data Science and Analytics*, **3**, 121–129, (2017).
- [22] Joseph D Ramsey, 'Scaling up greedy causal search for continuous variables', *arXiv preprint arXiv:1507.07749*, (2015).
- [23] Normadiah Mohd Razali, Yap Bee Wah, et al., 'Power comparisons of Shapiro-Wilk, Kolmogorov-Smirnov, Lilliefors and Anderson-Darling tests', *Journal of statistical modeling and analytics*, **2**, 21–33, (2011).
- [24] Alexander Reisch, 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).
- [25] Robert W Robinson, 'Counting unlabeled acyclic digraphs', in *Combinatorial Mathematics V*, 28–43, Springer, (1977).
- [26] Paul Rolland, Volkan Cevher, Matthäus Kleindessner, Chris Russell, Dominik Janzing, Bernhard Schölkopf, and Francesco Locatello, 'Score matching enables causal discovery of nonlinear additive noise models', in *Proceedings of the International Conference on Machine Learning*, pp. 18741–18753, (2022).
- [27] Karen Sachs, Omar Perez, Dana Pe'er, Douglas A Lauffenburger, and Garry P Nolan, 'Causal protein-signaling networks derived from multi-parameter single-cell data', *Science*, **308**, 523–529, (2005).
- [28] Pedro Sanchez, Xiao Liu, Alison Q O'Neil, and Sotirios A. Tsaftaris, 'Diffusion models for causal discovery via topological ordering', in *Proceedings of the International Conference on Learning Representations*, (2023).
- [29] Samuel Sanford Shapiro and Martin B Wilk, 'An analysis of variance test for normality (complete samples)', *Biometrika*, **52**, 591–611, (1965).
- [30] Peter Spirtes and Clark Glymour, 'An algorithm for fast recovery of sparse causal graphs', *Social Science Computer Review*, **9**, 62–72, (1991).
- [31] Peter Spirtes, Clark N Glymour, Richard Scheines, and David Heckerman, *Causation, prediction, and search*, MIT Press, 2000.
- [32] Chandler Squires and Caroline Uhler, 'Causal structure learning: A combinatorial perspective', *Foundations of Computational Mathematics*, 1–35, (2022).
- [33] Eric V Strobl and Thomas A Lasko, 'Identifying patient-specific root causes with the heteroscedastic noise model', *Journal of Computational Science*, 102099, (2023).
- [34] Henry C Thode, *Testing for normality*, CRC press, 2002.
- [35] Thomas Verma and Judea Pearl, 'Causal networks: Semantics and expressiveness', in *Machine Intelligence and Pattern Recognition*, volume 9, 69–76, Elsevier, (1990).
- [36] Sascha Xu, Osman A Mian, Alexander Marx, and Jilles Vreeken, 'Inferring cause and effect in the presence of heteroscedastic noise', in *Proceedings of the International Conference on Machine Learning*, pp. 24615–24630, (2022).
- [37] Xun Zheng, Bryon Aragam, Pradeep Ravikumar, and Eric P. Xing, 'DAGs with NO TEARS: Continuous optimization for structure learning', in *Advances in Neural Information Processing Systems*, pp. 9472–9483, (2018).
- [38] Xun Zheng, Chen Dan, Bryon Aragam, Pradeep Ravikumar, and Eric P. Xing, 'Learning sparse nonparametric DAGs', in *Proceedings of the International Conference on Artificial Intelligence and Statistics*, (2020).