

# Enabling Causal Discovery in Post-Nonlinear Models with Normalizing Flows

Nu Hoang\*, Bao Duong and Thin Nguyen

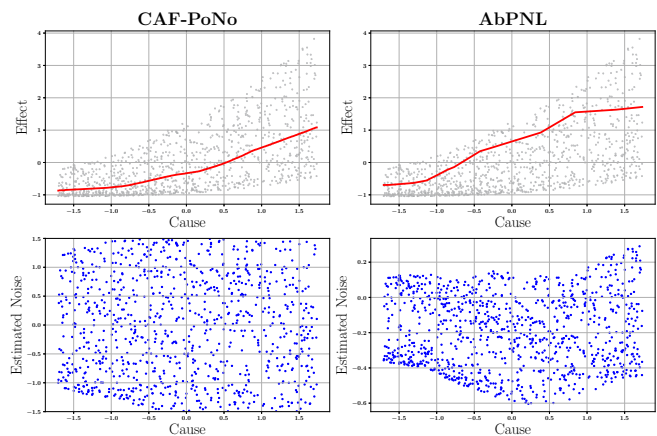
Applied Artificial Intelligence Institute (A<sup>2</sup>I<sup>2</sup>), Deakin University, Australia.

**Abstract.** Post-nonlinear (PNL) causal models stand out as a versatile and adaptable framework for modeling intricate causal relationships. However, accurately capturing the invertibility constraint required in PNL models remains challenging in existing studies. To address this problem, we introduce CAF-PoNo (Causal discovery via Normalizing Flows for Post-Nonlinear models), harnessing the power of the normalizing flows architecture to enforce the crucial invertibility constraint in PNL models. Through normalizing flows, our method precisely reconstructs the hidden noise, which plays a vital role in cause-effect identification through statistical independence testing. Furthermore, the proposed approach exhibits remarkable extensibility, as it can be seamlessly expanded to facilitate multivariate causal discovery via causal order identification, empowering us to efficiently unravel complex causal relationships. Extensive experimental evaluations on both simulated and real datasets consistently demonstrate that the proposed method outperforms several state-of-the-art approaches in both bivariate and multivariate causal discovery tasks.

## 1 Introduction

The need to uncover causal relationships from solely observational data where randomized controlled trials are impractical has led to the emergence of causal discovery methodologies as a crucial field of study. A key challenge in causal discovery stems from the non-uniqueness of causal models that can induce the same data distribution [32]. In other words, recovering the causal structure becomes impossible without making additional assumptions about the causal model. To overcome this, various functional causal models (FCMs) have been proposed with ensured identifiabilities, such as the linear non-Gaussian acyclic model (LiNGAM) [31], additive noise model (ANM) [12], and post-nonlinear (PNL) model [40]. While LiNGAM is limited to linear relationships, ANM adheres to the assumption of the additive noise with non-linear causal mechanisms.

Among them, the PNL model stands out for its generality in modeling complex non-linear causal systems. Specifically, under PNL models, the effect  $Y$  is generated from its cause  $X$  via the structural equation  $Y := g(h(X) + \epsilon_Y)$ , where  $\epsilon_Y$  is an exogenous noise independent of the cause. Here,  $h(\cdot)$  can be any function but  $g(\cdot)$  is constrained to be *invertible*, which has consequently introduced critical challenges in modeling and estimation. Several approaches have been investigated to effectively estimate PNL models [40, 35, 15], yet representing the exact invertibility required by the model still poses a major difficulty. For example, AbPNL [35] adopts auto-encoders



**Figure 1:** Comparison between CAF-PoNo and AbPNL [35] in cause-effect inference under PNL data. The red lines depict the estimated nonlinear functions based on the observed values of the effect and cause variables. In this example, AbPNL recovers the noises incorrectly due to the inability to model invertible functions of auto-encoders, leading to the noise visibly dependent on the cause. Meanwhile, with the adoption of normalizing flows, our CAF-PoNo method can accurately capture the invertibility constraint and recover the noises more correctly, which are independent of the cause.

to model a function and its “pseudo inverse” by minimizing the reconstruction error, so the full invertibility can only be obtained with zero reconstruction error everywhere, which is virtually impossible to achieve. Meanwhile, in [15] ranked-based methods are explored as a means to estimate strictly monotonic (and thus invertible) functions, however their analyses focus more on linear inner functions.

In this paper, to overcome the aforementioned challenges, we propose CAF-PoNo (**C**ausal discovery via **N**ormalizing **F**lows for **P**ost-**N**onlinear models)<sup>1</sup>, a novel estimation method for causal discovery under PNL models that capitalizes on the benefits of Normalizing Flows (NF), a powerful neural architecture known for its ability to model invertible functions [22]. This innovative adoption of NFs allows for parametrizing highly complex PNL models with the invertibility of  $g(\cdot)$  guaranteed without any ad-hoc enforcement, and thus drastically improves statistical efficiency in estimating them. The NF-enabled PNL models are then used to recover the underlying noises for the subsequent cause-effect identification task via statistical independence measures, in which our CAF-PoNo method exhibits remarkable capabilities compared with existing methods.

\* Corresponding Author. Email: nu.hoang@deakin.edu.au

<sup>1</sup> Source code is available at <https://github.com/htn274/CAFPoNo>.

As an example, Figure 1 demonstrates CAF-PoNo’s ability to correctly retrieve the underlying noise that is independent of the cause, whereas the autoencoder-based approach in AbPNL [35] suffers from the inaccurate invertible function estimation, resulting in an unfaithful noise that is still dependent on the cause. Furthermore, the proposed approach excels in addressing both bivariate and multivariate causal discovery, showcasing its effortless extensibility. Our main contributions are summarized as follows:

1. We address the long-standing challenge of accurately modeling the invertible functional mechanism in PNL causal models with the introduction of the CAF-PoNo method. By leveraging normalizing flows, our method can capture the exact invertibility of the PNL model, leading to improved statistical efficiency and estimability compared with existing PNL-based methods. To the best of our knowledge, this is the first application of normalizing flows to learn causal relationships under the PNL model.
2. We extend the CAF-PoNo framework from handling bivariate causal discovery (i.e., telling apart cause and effect) to handling multivariate causal discovery, by leveraging the causal ordering technique aided with independence scoring, exhibiting an attractive polynomial runtime with provable correctness.
3. We demonstrate the efficacy of the proposed CAF-PoNo method in causal discovery via an extensive set of numerical evaluations on both simulated and real datasets. The empirical results confirm the superiority of CAF-PoNo over competitive baselines in both bivariate and multivariate scenarios.

## 2 Related works

### 2.1 Causal Discovery under PNL Models

While various identifiable Structural Equation Models (SEMs) and estimation methods have been proposed to address the identifiability issue in causal discovery [31, 12, 40, 16, 14], our study focuses on the PNL model, which is one of the most flexible model representing complicated causal relationships. Nevertheless, causal discovery under PNL models is very challenging and under-studied due to the invertibility constraint.

Several approaches have been proposed to address the challenge, yet they own certain limitations that hinder their ability to achieve higher performance. For instance, in [40], the PNL model is modeled by neural networks and estimated by minimizing the mutual information between the cause and noise. However, this approach ignores the invertibility constraint during the optimization process, which potentially results in invalid estimates that violate the invertibility requirement. To address this problem, AbPNL [35] leverages autoencoders to represent “approximately invertible” functions, which are estimated by minimizing both the reconstruction loss and mutual information of the noise and the presumed cause. However, AbPNL suffers from two key limitations. Firstly, it lacks a guarantee of learning a truly invertible function, leading to failures in reconstructing the underlying noise in the effect variable, inevitably causing incorrect causal direction identification. Secondly, the loss function necessitates computing mutual information during training, resulting in significant computational costs, especially when dealing with multiple variables. Meanwhile, instead of jointly learning both  $g(\cdot)$  and  $h(\cdot)$ , Keropyan et al. [15] employs ranked-based regression methods to separately learn the two transformations. Nevertheless, the study mainly focuses on linear settings for the inner function.

### 2.2 Multivariate Causal Discovery

Traditional causal discovery methods can be divided into three groups: constraint-based methods [32, 26, 3], score-based methods [5, 6, 10], and hybrid methods [34, 21]. Most of these methods search for the causal directed acyclic graph (DAG) within the combinatorial space of graph structure, which is known to be challenging due to the super-exponential explosion in complexity with respect to the number of variables [27]. This has changed in 2018 when Zheng et al. [42] introduced a novel approach known as NOTEARS that maps the discrete space to a continuous one, enabling the utilization of various gradient-based optimization techniques. This breakthrough opened up new possibilities for causal discovery by incorporating deep learning and continuous optimization methodologies [39, 19, 18].

Alternatively, instead of searching the vast space of all possible causal structures, another line of research focuses on finding a causal topological order [28, 30, 38], accomplished by searching over the space of permutations, which is orders of magnitude smaller than the space of DAGs. By fixing a topological order, the acyclicity constraint is naturally ensured, eliminating the need for additional checks. After obtaining a causal topological order, an additional pruning step is necessary to remove spurious edges. This pruning process helps refine the inferred causal relationships by discarding unnecessary connections between variables.

### 2.3 Normalizing Flows in Causal Discovery

Normalizing flows [22] have emerged as a class of expressive generative models within deep learning, capable of effectively modeling complex probability distributions by learning invertible transformations. In the context of causal discovery, many studies have applied normalizing flows to encapsulate intricate data distributions that may arise in practice [25, 4, 17]. For instance, a recent study [16] have employed autogressive flows in the context of location-scale (LS) models. More specifically, the causal mechanism for LS models is defined as  $y = g(x)\epsilon + h(x)$ , which deviates significantly from PNL models. Based on the formulations of PNL and LS models, they cannot generalize over each other, meaning that the techniques used in [16] are not applicable to PNL models, as empirically confirmed in Section 5. To the best of our knowledge, normalizing flows have yet been studied for causal discovery under PNL models, where the invertible causal mechanism is of central importance, which highlights the significance of our study.

## 3 Preliminary

Suppose we observe an empirical dataset of  $n$  i.i.d. samples of the random vector  $X = [x_1, x_2, \dots, x_d]^T$  of  $d$  variables. We assume that the data is generated under the PNL model as follows

$$x_i := g_i(h_i(X_{\text{Pa}(i)})) + \epsilon_i, i = 1, 2, \dots, d \quad (1)$$

where  $\text{Pa}(i)$  is a set of direct causes (or parents) of  $x_i$ . The noise variables  $\{\epsilon_i\}$  are mutually independent and therefore  $\epsilon_i \perp\!\!\!\perp X_{\text{Pa}(i)}$  for each  $i = 1, 2, \dots, d$ . The corresponding causal graph, denoted as  $\mathcal{G}$ , has directed edges  $(j \rightarrow i) : j \in \text{Pa}(i)$  representing the causal relationships between variables. Similar to previous studies [15, 19, 28], this work also relies on assumptions of acyclicity and sufficiency. That means the causal graph  $\mathcal{G}$  is a directed acyclic graph (DAG) and there is no hidden confounder. The goal of causal discovery is to infer the true DAG  $\mathcal{G}$  based on the empirical data of  $X$ .

For bivariate causal discovery, with a slight abuse of notation, we denote the cause and effect as  $X$  and  $Y$ , respectively, where:

$$y := g(h(x) + \epsilon_Y) \tag{2}$$

The PNL model is identifiable in most cases, except for specific instances where certain combinations of functions and noise distributions are involved as outlined in [40]. Moreover, the PNL model is one of the most generic SEMs as it generalizes both the prevalent LiNGAM [31] and ANM [12] models.

#### 4 CAF-PoNo: Causal Discovery via Normalizing Flows for Post-Nonlinear Models

In this section, we first elaborate how to incorporate normalizing flows to design an effective estimator of PNL causal models. Then, we explain how to leverage this estimator for distinguishing cause and effect. Lastly, the framework is generalized to unravel the causal structure among multiple variables under PNL data.

##### 4.1 Normalizing Flows for the PNL model

Normalizing flows is a powerful tool to express complex probability distributions from a simple base distribution through several invertible and differentiable transformations by exploiting the change of variables rule. The primary objective of normalizing flows research is to devise efficient transformations that possess the desirable properties of invertibility and differentiability with a tractable derivative/Jacobian for effectively applying the change of variables rule. For more details regarding normalizing flows, we refer to [22].

Here, consider an empirical dataset of two variables  $(X, Y)$  generated by Eq. (2), we are interested in estimating the functions  $g(\cdot)$  and  $h(\cdot)$ , so that we can recover the noise as

$$\epsilon_Y = g^{-1}(y) - h(x). \tag{3}$$

Normalizing flows are used in this study specifically to model the invertible function  $g(\cdot)$ . In particular, we employ the cumulative distribution function (CDF) flow [22], which is a simple, yet efficient invertible transformation for one dimensional (scalar) variables. A CDF flow is a positively weighted combination of a set of CDFs of any arbitrarily positive density function. For example, the following mixture is a CDF, and is thus invertible:

$$f(x) = \sum_{i=1}^k w_i \Phi_i(x, \mu_i, \sigma_i), \tag{4}$$

where  $k$  is the number of components,  $w_i \geq 0$ ,  $\sum_{i=1}^k w_i = 1$  and  $(\mu_i, \sigma_i)$  are the weights and parameters for each component, respectively, whereas  $\Phi_i$  is the cumulative distribution function of the  $i$ -th component. For simplicity, in our implementation, we consider a Gaussian distribution with mean  $\mu_i$  and variance  $\sigma_i^2$  for each component, yet it should be noted that *our method is not restricted to Gaussian data* but can be applied to more generic settings. Thanks to the universal approximation capability of Gaussian mixture models, the Gaussian mixture CDF flow can express any strictly monotonic  $\mathbb{R} \rightarrow (0, 1)$  map with arbitrary precision [13].

To apply the CDF flow for PNL estimation, let  $z = h(x) + \epsilon_Y$ , then we have  $y = g(z)$  and  $z = g^{-1}(y)$ . We use neural networks to parametrize the inner function  $h$ , while CDF flows are used to model  $g^{-1}$  instead of  $g$  since the noise directly relates to  $g^{-1}$  as shown in Eq. (3). To model  $g^{-1}$ , which is a  $\mathbb{R} \rightarrow \mathbb{R}$  map, we start with a CDF

---

#### Algorithm 1: CAF-PoNo algorithm.

---

**Input:** Empirical samples  $(X, Y) = \{(x_i, y_i)\}_{i=1}^n$ .

**Output:** The predicted causal direction.

Standardize  $X$  and  $Y$  to zero mean and unit variance. ;

Calculate the independence scores of two possible causal directions:

$$S_{X \rightarrow Y} = \text{CAF-PoNo}(X, Y)$$

$$S_{Y \rightarrow X} = \text{CAF-PoNo}(Y, X)$$

Determine the causal direction based on the independence scores:

$$\text{dir} = \begin{cases} 1, & \text{if } S_{X \rightarrow Y} > S_{Y \rightarrow X} \\ -1, & \text{if } S_{X \rightarrow Y} < S_{Y \rightarrow X} \end{cases}$$

**return** dir

**Function** CAF-PoNo( $X, Y$ ):

**Data:** Empirical samples for the cause  $X = \{x_i\}_{i=1}^n$  and the effect  $Y = \{y_i\}_{i=1}^n$ .

**Result:** The independence score for the causal direction  $X \rightarrow Y$ .

Perform train-test split to get  $(X_{\text{train}}, Y_{\text{train}})$  and  $(X_{\text{test}}, Y_{\text{test}})$ ;

Train the CDF Flow on  $(X_{\text{train}}, Y_{\text{train}})$  to obtain the learned parameters  $\theta$ ;

Estimate the noise:

$$\hat{\epsilon}_{Y_{\text{test}}} = g_{\theta}^{-1}(Y_{\text{test}}) - h_{\theta}(X_{\text{test}})$$

Compute the independence score:

$$S_{X \rightarrow Y} := -\text{HSIC}(\hat{\epsilon}_{Y_{\text{test}}}, X_{\text{test}})$$

**return**  $S_{X \rightarrow Y}$

**End Function**

---

flow, which has a  $(0, 1)$  co-domain, then wrap it with another invertible map  $(0, 1) \rightarrow \mathbb{R}$ , such as an inverse sigmoid function. In short, our proposed NF-based approximation of  $g^{-1}$  can be represented as follows:

$$z = g^{-1}(y) = \text{sigmoid}^{-1} \left( \sum_{i=1}^k w_i \Phi_i(y, \mu_i, \sigma_i) \right) \tag{5}$$

and the noise can be estimated as:  $\hat{\epsilon}_Y = z - h_{\theta}(x)$ .

We adopt the Maximum Likelihood Estimation (MLE) framework to learn the proposed PNL estimation in Eq. (5). Let  $\theta$  be the total set of all parameters, which includes the parameters for the neural networks  $h_{\theta}(x)$  and those for the CDF flow. The CDF flow consists of a fixed number of  $k$  Gaussian components, with each component having its own set of parameters  $w, \mu, \sigma$ . We chose Gaussian distribution as a base distribution for  $\epsilon_Y \sim \mathcal{N}(0, 1)$ . By the change of variables rule, the likelihood  $p_{\theta}(y | x)$  can be expressed as follows:

$$p_{\theta}(y | x) = p_{\theta}(\epsilon_Y | x) \left| \frac{\partial \epsilon_Y}{\partial y} \right| \tag{6}$$

$$= p_{\epsilon_Y}(\epsilon_Y) \left| \frac{\partial g_{\theta}^{-1}}{\partial y} \right| \tag{7}$$

$$= p_{\epsilon_Y}(g_{\theta}^{-1}(y) - h_{\theta}(x)) \left| \frac{\partial g_{\theta}^{-1}}{\partial y} \right| \tag{8}$$

**Algorithm 2:** Causal ordering identification.

---

**Input:** Data matrix  $X \in \mathbb{R}^{n \times d}$ .  
**Output:** The causal ordering  $\pi$ .  
 $\pi \leftarrow []$ ;  
nodes  $\leftarrow \{1, 2, \dots, d\}$ ;  
**repeat**  
  **for**  $i \in$  nodes **do**  
     $X_{(-i)} = \{x_j\}_{j \in \text{nodes} \setminus \{i\}}$ ;  
    Estimate the noise by CAF-PoNo model:  
       $\hat{\epsilon}_i = g_\theta^{-1}(x_i) - h_\theta(X_{(-i)})$   
    Compute the sink score:  
       $S(x_i) := - \left( \max_{x \in X_{(-i)}} \text{HSIC}(\hat{\epsilon}_i, x) \right)$   
  **end**  
  sink  $\leftarrow \arg \max_{x_i \in \text{nodes}} S(x_i)$ ;  
   $\pi \leftarrow [\pi, \text{sink}]$ ;  
  nodes  $\leftarrow \text{nodes} - \{\text{sink}\}$ ;  
**until** nodes =  $\emptyset$ ;  
**return**  $\pi$

---

We aim to maximize the log-likelihood of the observed  $Y$  conditioned on  $X$  over the space of  $\theta$ :

$$\begin{aligned} \mathcal{L}(\theta) &= \frac{1}{n} \sum_{i=1}^n \ln p_\theta(y_i | x_i) \\ &= \frac{1}{n} \sum_{i=1}^n \left( \ln p_{\epsilon_Y}(g_\theta^{-1}(y) - h_\theta(x)) + \ln \left| \frac{\partial g_\theta^{-1}}{\partial y} \right| \right) \end{aligned} \quad (9)$$

where  $\{(x_i, y_i)\}_{i=1}^n$  is the set of  $n$  observed samples of  $(X, Y)$ .

An important advantage of the employed CDF flow is the ability to tractably calculate the log-derivative  $\ln \left| \frac{\partial g_\theta^{-1}}{\partial y} \right|$  in a closed-form fashion. Particularly, let  $t = \sum_{i=1}^k w_i \Phi(y, \mu_i, \sigma_i)$ , we then have  $z = -\ln\left(\frac{1}{t} - 1\right)$ , which yields:

$$\begin{aligned} \ln \left| \frac{\partial g_\theta^{-1}}{\partial y} \right| &= \ln \left| \frac{\partial g_\theta^{-1}}{\partial t} \right| + \ln \left| \frac{\partial t}{\partial y} \right| \\ &= \ln \frac{1}{t(1-t)} + \ln \left| \sum_{i=1}^k w_i \mathcal{N}(y, \mu_i, \sigma_i) \right| \end{aligned} \quad (10)$$

By substituting Eq. (5) and Eq. (10) into Eq. (9), we get a fully differentiable loss function that can be optimized using off-the-shelf gradient-based methods. Moreover, the adoption of the MLE framework in cause-effect estimation is justified by established theory [41] in the following Lemma.

**Lemma 1.** (Maximum likelihood for causal discovery under PNL models, Theorem 2 of [41]). *The parameter set  $\theta^*$  that maximizes the likelihood  $\mathbb{E}[\ln p_\theta(Y | X)]$  also minimizes the mutual information of the cause  $X$  and the noise  $\epsilon_Y$ .*

The proof can be found in [41]. Simply put, Lemma 1 ensures that with a sufficiently capable model, maximizing MLE is equivalent to reaching the ground truth PNL in the true direction, given access to the ground truth joint distribution.

## 4.2 Cause-Effect Inference

When the PNL model is identifiable, by definition, there exists no reverse model  $x = g_X(h_X(y) + \epsilon_X)$  such that  $g_X$  is invertible and  $Y \perp\!\!\!\perp \epsilon_X$  [40], which implies that the best model found by MLE in the anti-causal direction will have  $\epsilon_X \not\perp\!\!\!\perp Y$ , since invertibility is already guaranteed. This insight enables us to design an independence-based cause-effect identification method.

More specifically, to decide the causal direction, we fit a model corresponding to each direction  $X \rightarrow Y$  and  $Y \rightarrow X$  to the data, after which we obtain the estimated noises  $\hat{\epsilon}_Y$  and  $\hat{\epsilon}_X$ , respectively. Subsequently, the model that exhibits independence between the estimated noise and the putative cause is chosen as the more likely causal explanation. To quantify independence, we employ the Hilbert-Schmidt independence criterion (HSIC) [9] which is a popular kernel-based dependence measure that requires no parameter estimation. A higher HSIC value indicates a stronger dependence, while a lower value indicates a lower level of dependence. Consequently, the independence score is the negative value of the HSIC. The details of the proposed bivariate causal discovery algorithm is illustrated in Algorithm 1.

## 4.3 Extension to Multivariate Causal Discovery

An important aspect of our approach is that it is readily extendable to multiple variables, thus effectively addressing the challenges of the multivariate causal discovery problem.

Towards this end, we decompose the causal structure learning task into two stages, where we first identify the causal ordering among the variables, then the edges are recovered with respect to said ordering. This approach ensures the acyclicity of the resultant graph, while attaining an efficient polynomial runtime.

### 4.3.1 Causal Ordering Identification

**Definition 1.** *A causal ordering of a graph  $\mathcal{G}$  is a non-unique permutation  $\pi$  of  $d$  nodes such that a given node always precedes its descendants in  $\pi$ , i.e.,  $i \prec_\pi j \Leftrightarrow j \in \text{Des}_{\mathcal{G}}(i)$ .*

Each permutation  $\pi$  corresponds to a unique, fully connected DAG  $\mathcal{G}^\pi$  where every pair of nodes  $i \prec_\pi j$  defines directed edges  $(i \rightarrow j)$  in  $\mathcal{G}^\pi$ . Therefore, the graph  $\mathcal{G}$  that we seek for is a subgraph of  $\mathcal{G}^\pi$  if  $\pi$  is among the topological sorts of  $\mathcal{G}$ . As a consequence, a pruning procedure needs to be performed to eliminate redundant edges from  $\mathcal{G}^\pi$ .

Following [36], we utilize the two following propositions to find one of causal orders of the underlying causal structure, by recursively detecting and excluding the sink node.

**Proposition 2.** *The noise  $\epsilon_i$  of every node  $x_i$  is independent of all its non-descendants.*

**Proposition 3.** *An arbitrary node  $x_i$  is considered as a sink node if and only if its corresponding noise  $\epsilon_i$  is independent of all other variables, which are denoted as  $X_{(-i)} := X \setminus \{x_i\}$ .*

We refer to [36] for the detailed proof of these propositions. To identify the sink node, we extend the proposed bivariate model to accommodate multiple causes. Specifically, the exogenous noise of an arbitrary node  $x_i$ , denoted as  $\epsilon_i$ , is estimated from the other variables excluding  $x_i$ , i.e.,  $\hat{\epsilon}_i := g_\theta^{-1}(x_i) - h_\theta(X_{(-i)})$ . To check whether  $x_i$  is a sink node, we introduce the sink score, which quantifies the independence degree of the estimated noise  $\hat{\epsilon}_i$  and  $X_{(-i)}$ . The sink

score of a node  $x_i$ , denoted as  $S(x_i)$ , is defined as the negative of the maximum value of HSIC between  $\hat{\epsilon}_i$  and every  $x \in X_{(-i)}$ :

$$S(x_i) := - \left( \max_{x \in X_{(-i)}} \text{HSIC}(\hat{\epsilon}_i, x) \right) \quad (11)$$

By evaluating the sink score for each variable, we can identify the sink node as the variable with the highest sink score. Once the sink node is identified, it is removed from the list and the process is repeated with the remaining variables until the full causal ordering is identified. The details of the causal ordering process is outlined in Algorithm 2.

### 4.3.2 Pruning Method

To eliminate spurious edges and refine the causal structure, we employ an array of conditional independence tests based on a given permutation. In particular, we remove edges  $(i \rightarrow j)$  if  $X_i \perp\!\!\!\perp X_j | X_{\text{pre}_\pi(j) \setminus \{i\}}$  where  $\text{pre}_\pi(j)$  is a set of preceding variables in the causal ordering  $\pi$  [37]. This pruning method does not depend on the SEM assumption, and yet is more flexible than the CAM pruning, a popular pruning method in other ordering-based studies [28, 30, 38], which is however only designed for generalized additive models. That being said, we also investigate both approaches comparatively in the Appendix [11].

To perform the conditional independence tests, we employ a recent state-of-the-art latent representation learning based conditional independence testing (LCIT) method [7], which does not make any parametric assumption and is shown to scale linearly with the sample size. It is worth mentioning that alternative conditional independence testing methods can also be employed during this pruning process. We summarize the pruning method in Algorithm 3. Together, Algorithm 2 and Algorithm 3 highlights the key steps for the CAF-PoNo method in multivariate causal discovery.

### 4.3.3 Complexity Analysis

The causal ordering identification in Algorithm 2 requires  $\mathcal{O}(d^2)$  iterations, each spending  $\mathcal{O}(nd)$  for the noise estimation using CAF-PoNo and  $\mathcal{O}(n^2d)$  for the sink score calculation using HSIC, leading to a total complexity of  $\mathcal{O}(n^2d^3)$ . Similarly, the pruning step in Algorithm 3 performs  $\mathcal{O}(d^2)$  conditional independence tests, each costing  $\mathcal{O}(nd)$ , resulting to the final complexity of  $\mathcal{O}(nd^3)$ .

## 5 Numerical Evaluations

To assess the quality of the proposed method, we compare it with state-of-the-art causal discovery models on two scenarios: bivariate and multivariate cause-effect inference, using both synthetic and real data, where the ground truth causal directions/structures are available.

Regarding parameter selection in our method, for each dataset, the hyperparameters are chosen to maximize the AUC score, with values among the following sets:

- The number of Gaussian components  $k$ : {4, 8, 10, 12}.
- The hidden size of a single layer neural network: {16, 32, 64, 128}.
- The number of epochs: {200, 300, 500, 700, 1000}.

All parameters are trained using the Adam optimizer with early stopping, learning rate of 0.001, and batch size of 128. We randomly

---

### Algorithm 3: Pruning algorithm.

---

**Input:** Data matrix  $X \in \mathbb{R}^{n \times d}$ , a causal ordering  $\pi$ , a conditional independence testing method CIT returning the  $p$ -value for the hypothesis  $X_i \perp\!\!\!\perp X_j | X_Z$ , and a significance level  $\alpha$  for the test.

**Output:** The adjacency matrix  $A$  of the final DAG.  
Initialize  $A$  as a zero value matrix of size  $d \times d$ ;

```

for  $j \in \pi$  do
     $\text{Pre}_j = \pi[: j]$ ;
    for  $i \in \text{Pre}_j$  do
         $Z = \text{Pre}_j \setminus \{i\}$ ;
         $A[i, j] = \begin{cases} 1 & \text{if } \text{CIT}(X_i, X_j, X_Z) < \alpha \\ 0 & \text{otherwise} \end{cases}$ ;
    end
end
return  $A$ 

```

---

**Table 1:** Bivariate causal discovery performance on various synthetic datasets. The evaluation metric is AUC (higher is better). The proposed CAF-PoNo method is compared against ANM [12], PNL [40], CAREFL [16], AbPNL [35], RECI [1], and LOCI [14]. We highlight in bold the **best** result and in underline the **second best** result. CAF-PoNo achieves the highest AUC on most benchmark datasets.

Method	Simple PNL (different noises)			GP PNL	LS	AN	Overall
	Gaussian	Laplace	Uniform				
ANM	56.10	60.26	54.97	82.71	48.78	53.84	59.44
PNL	<u>91.41</u>	68.42	88.64	85.71	67.83	66.59	78.10
CAREFL	85.43	79.24	<u>89.61</u>	<u>90.10</u>	75.53	82.46	<u>83.73</u>
AbPNL	45.43	51.81	47.94	13.49	40.86	65.66	44.19
RECI	81.38	<u>89.72</u>	39.70	87.72	84.91	91.88	79.22
LOCI	00.00	00.00	00.00	00.00	<b>97.04</b>	<b>100.00</b>	32.84
Ours	<b>95.44</b>	<b>93.51</b>	<b>98.73</b>	<b>95.88</b>	<u>89.16</u>	<u>97.17</u>	<b>94.98</b>

split the datasets into training and test sets with a ratio of 6:4. The validation set comes from 20% of the training set. The training set is used to learn model parameters, while the validation set is utilized for early stopping. The testing set is for computing the independence score for causal direction identification. Following previous studies [40, 35, 36], we use the same setting for HSIC, including Gaussian kernel and bandwidth  $\sigma$  of 1.

## 5.1 Bivariate Causal Discovery

### 5.1.1 Baselines and the Evaluation Metric

We compare the bivariate causal discovery performance of the proposed method against several popular baselines as described below:

- Additive noise model (ANM) [12] estimates the causal mechanism via Gaussian Process regressions and utilizes HSIC to assess the independence between the cause and residuals.
- Post nonlinear (PNL) [40] reconstructs the noise using two neural networks trained to minimize the mutual information between the putative cause and the noise, and utilizes HSIC test for independence test between the cause and the estimated noise.
- Causal Autoregressive Flow (CAREFL) [16] utilizes affine normalizing flows to estimate the affine causal model introduced in the paper.
- Autoencoder-based post nonlinear (AbPNL) [35] models PNL using an auto-encoder trained with a combination of independence and reconstruction losses.

**Table 2:** Bivariate causal discovery performance on real datasets (Tübingen dataset [20]). The evaluation metrics are AUC and Accuracy (higher is better). The proposed CAF-PoNo method is compared against ANM [12], PNL [40], CAREFL [16], AbPNL [35], RECI [1], and LOCI [14]. We highlight in bold the **best** result and in underline the second best result. CAF-PoNo attains the highest scores on both AUC and accuracy.

Method	ANM	PNL	AbPNL	CAREFL	RECI	LOCI	CAF-PoNo
AUC	58.86%	63.22%	68.51%	65.10%	<u>69.67%</u>	62.04%	<b>73.84%</b>
Accuracy	52.52%	59.59%	62.62%	52.52%	62.62%	<u>63.63%</u>	<b>71.71%</b>

- RECI [1] employs a monomial regressor to predict the target variable using normalized versions of potential causal variables and then compares the mean squared errors (MSE) of these regressions for each possible causal direction. The direction with the lower MSE is selected as the inferred causal relationship.
- LOCI [14] harnesses the estimation of location scale noise models (LSNM) via heteroscedastic regression to reconstruct the underlying noise in the observed data. The causal direction is then determined by comparing the level of independence between the estimated noise and the presumed cause under two possible causal directions.

We utilize code from Causal Discovery Toolbox (CDT)<sup>2</sup> and Causal-learn framework<sup>3</sup> for ANM, RECI and PNL methods. The implementations of the other methods are derived from the original code published by the authors.

As for the evaluation metric, we utilize the bi-directional area under the receiver operating characteristic curve (AUC), where higher values indicate better performance. The AUC was computed by averaging two scores: one measures classification performance for determining if a pair  $(X, Y)$  belongs to the  $X \rightarrow Y$  class, and the other accounts for the  $Y \rightarrow X$  class.

### 5.1.2 Synthetic data

We generate datasets that follow the PNL model in Eq. (2), considering both simple and complex functional relationships. Following [40, 35], simple functional relationships are captured in the Simple-PNL dataset, which is constructed using various function choices as outlined in the Appendix [11]. We consider three different noise distributions: Gaussian, Uniform, and Laplace noises. For each distribution, we generate 500 pairs of  $(X, Y)$ , each pair consisting of 1,000 i.i.d. samples. Furthermore, we introduce PNL-GP dataset with more intricate functional relationships within the PNL model where  $h(\cdot)$  and  $g(\cdot)$  are the weighted sum of Gaussian Processes and sigmoid functions, respectively. To assess CAF-PoNo’s effectiveness beyond the PNL assumption, we also utilize synthetic datasets generated from non-linear additive noise (AN) and location-scale (LS) noise models, as provided by [33].

The results are presented in Table 1, demonstrating that CAF-PoNo outperforms other methods on all PNL datasets, achieving the highest AUC score overall. Notably, the results highlight the robustness of CAF-PoNo across various types of noise even though the base noise type is solely set as Gaussian, implying CAF-PoNo’s effectiveness in addressing noise mis-specification. Interestingly, the results also reveal CAF-PoNo’s capability in handling model mis-specification. In more details, CAF-PoNo attains an AUC of approximately 90% under LS data, demonstrating a comparable performance to LOCI, a method specifically designed for the LS scenario.

<sup>2</sup> <https://github.com/FenTechSolutions/CausalDiscoveryToolbox>

<sup>3</sup> <https://github.com/py-why/causal-learn>

### 5.1.3 Real world data

To demonstrate the effectiveness of CAF-PoNo in real world scenarios, we utilize the Tübingen dataset [20], a common benchmark in bivariate causal discovery comprising of 99 cause-effect pairs from a variety of domains, e.g., abalone measurements, census income, and liver disorders. The number of samples for each dataset varies from 94 to 16,382.

Table 2 illustrates the comparative performance in terms of AUC and Accuracy scores. The results clearly demonstrate that CAF-PoNo outperforms other baselines on both metrics when applied to real-world scenarios. Specifically, CAF-PoNo achieves the highest AUC, surpassing the runner-up RECI model by a significant margin of around 5%. Following prior studies [16, 14], we also provide the accuracy which is a common metric in bivariate causal discovery. As evidenced by the table, CAF-PoNo again accomplishes the highest accuracy of 71.71%, leaving the second-best method LOCI by a large margin of over 8%.

## 5.2 Multivariate Causal Discovery

### 5.2.1 Baselines and Evaluation Metrics

We consider the multivariate version of AbPNL [36] which is the first work demonstrating multivariate causal discovery for PNL model. In addition, we also include three ordering-based methods designed for ANM models including RESIT [24], NPVar [8] and SCORE [28]. The causal orders obtained from these methods are applied with the same pruning procedure as described in Section 4.3.2 with a commonly adopted significance level of 0.001.


We consider two common structural metrics for multivariate causal discovery evaluation: the structural Hamming distance (SHD) and structural intervention distance (SID) [23]. These metrics enable us to quantify the error between the predicted DAG with the true DAG. In particular, the SHD measures the total number of missing, extra, and reverse edges between the estimate and the true DAG, while the SID measures the minimum number of interventions required to transform the output to the true DAG. Hence, lower values for both SHD and SID indicate a better fit of the output DAG to the given data. In addition, we also consider the order divergence ( $D_{\text{order}}$ ) [28] as a measure of the quality of the causal ordering. More specifically, the order divergence reflects the number of directed edges in the true DAG that disagree with the causal ordering.

### 5.2.2 Synthetic data


We generate Erdős-Rényi causal graphs [2] with  $d$  nodes with an expected in-degree of 2. The data is generated following the PNL model (1), where  $h_i$  represents the weighted sums of Gaussian processes and  $g_i$  represents sigmoid functions. For each causal structure, we generate 1,000 samples from uniform noises  $\mathcal{U}(0, 1)$ .

**Effect of dimensionality.** To study the performance of the proposed method across different numbers of nodes, we vary  $d$  from 4 to 11 variables. The results are shown in Figure 2, showing that our proposed method consistently outperforms other baselines in most cases. Remarkably, the proposed method shows a significantly low  $D_{\text{order}}$  compared with competitors when the graph size grows. Surprisingly, AbPNLMulti exhibits the poorest performance although the method is specifically designed for the PNL model.


**Effect of sample size.** In order to assess the impact of sample sizes on the performance of different models, we vary the sample size from 100 to 1,000 while fixing  $d = 4$ . The results in Figure 3 reveal that



**Figure 2:** Multivariate causal discovery performance on synthetic data as a function of the number of variables. We fix  $n = 1,000$  and vary the number of variables. The evaluation metrics are  $D_{order}$ , SHD, and SID (lower is better). The reported values are aggregated over 10 independent runs. We compare the proposed CAF-PoNo method with RESIT [24], NPVar [8], SCORE [28], and AbPNLMulti [36].



**Figure 3:** Multivariate causal discovery performance on synthetic data as a function of sample size. We fix  $d = 4$  and vary the sample size. The evaluation metrics are  $D_{order}$ , SHD, and SID (lower is better). The reported values are aggregated over 10 independent runs. We compare the proposed CAF-PoNo method with RESIT [24], NPVar [8], SCORE [28], and AbPNLMulti [36].



**Figure 4:** The running time in seconds as a function of the number of variables. The proposed method exhibits a significant reduction in running time compared to AbPNL, demonstrating its potential scalability to high dimensional data.

CAF-PoNo constantly surpasses other methods across various settings. In these cases, CAF-PoNo achieves extraordinarily low  $D_{order}$  values, hovering around 1. Notably, CAF-PoNo exhibits exceptional efficiency when the sample size exceeds 500. In particular, both SHD and SID decrease to below 2 and 4, respectively. These results underscore the robustness and effectiveness of CAF-PoNo, particularly its sample efficiency.

**Running time.** Figure 4 illustrates the running time of both AbPNL and CAF-PoNo algorithms as the number of variables increases. The results clearly demonstrate that CAF-PoNo exhibits a substantial reduction in running time with up to 5x speed-up compared to AbPNL. This considerable gain stems from the lightweight of the CDF flow and our proposed MLE framework, which eliminates the need for mutual information calculations in AbPNL, leading to a more computationally efficient model.

### 5.2.3 Real data

We test these methods on the Sachs dataset [29], a popular benchmark dataset in multivariate causal discovery. The dataset contains

**Table 3:** Multivariate causal discovery performance on a real data of the protein signaling network. The evaluation metrics are SHD, and SID (lower is better). We compare the proposed CAF-PoNo method with four baselines RESIT [24], NPVar [8], SCORE [28], and AbPNLMulti [36].

Method	SHD ( $\downarrow$ )	SID ( $\downarrow$ )
RESIT	13	47
NPVar	14	57
SCORE	12	45
AbPNLMulti	13	47
<b>CAF-PoNo (Ours)</b>	<b>11</b>	<b>41</b>

853 observational samples, representing expression data of the protein signaling network of 11 nodes and 17 edges. The results are shown in Table 3, which demonstrates the superior performance of our method in comparison with state-of-the-art baselines, evidenced by lowest SHD and SID among all methods. These results highlight the remarkable capability of our proposed method in successfully uncovering intricate causal structures in real-world datasets.

## 6 Conclusion

In this paper, we introduce CAF-PoNo, an innovative and efficient flow-based approach for cause-effect identification under the PNL model. Moreover, CAF-PoNo surpasses existing methods in parameter estimation for the PNL model by adhering to the invertibility constraint, resulting in the highest AUC on synthetic and real datasets. We also extend CAF-PoNo for multivariate causal discovery and show its superiority over state-of-the-art methods.

## References

- [1] P. Blöbaum, D. Janzing, T. Washio, S. Shimizu, and B. Schölkopf. Cause-effect inference by comparing regression errors. In *Proceedings of the International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 900–909. PMLR, 2018.
- [2] B. Bollobás. The Evolution of Random Graphs. *Transactions of the American Mathematical Society*, 286:257–274, 1984.
- [3] D. Colombo, M. H. Maathuis, M. Kalisch, and T. S. Richardson. Learning High-Dimensional Directed Acyclic Graphs with Latent and Selection Variables. *The Annals of Statistics*, 40:294–321, 2012.
- [4] C. Cundy, A. Grover, and S. Ermon. BCD Nets: Scalable Variational Approaches for Bayesian Causal Discovery. In *Advances in Neural Information Processing Systems*, pages 7095–7110, 2021.
- [5] J. Cussens. Bayesian network learning with cutting planes. In *Proceedings of the Conference on Uncertainty in Artificial Intelligence (UAI)*, pages 153–160, 2011.
- [6] J. Cussens, M. Järvisalo, J. H. Korhonen, and M. Bartlett. Bayesian Network Structure Learning with Integer Programming: Polytopes, Facets and Complexity. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*, pages 4990–4994, 2017.
- [7] B. Duong and T. Nguyen. Conditional Independence Testing via Latent Representation Learning. In *Proceedings of The IEEE International Conference on Data Mining (ICDM)*, pages 121–130, 2022.
- [8] M. Gao, Y. Ding, and B. Aragam. A polynomial-time algorithm for learning nonparametric causal graphs. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 33, pages 11599–11611, 2020.
- [9] A. Gretton, O. Bousquet, A. Smola, and B. Schölkopf. Measuring statistical dependence with hilbert-schmidt norms. In *Proceedings of the International Conference on Algorithmic Learning Theory (ALT)*, pages 63–77, 2005.
- [10] R. Hemmecke, S. Lindner, and M. Studený. Characteristic imsets for learning Bayesian network structure. *International Journal of Approximate Reasoning*, 53:1336–1349, 2012.
- [11] N. Hoang, B. Duong, and T. Nguyen. Enabling causal discovery in post-nonlinear models with normalizing flows. *arXiv preprint arXiv:2407.04980*, Full version of this paper.
- [12] P. O. Hoyer, D. Janzing, J. Mooij, J. Peters, and B. Schölkopf. Nonlinear causal discovery with additive noise models. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 689–696, 2008.
- [13] G. Ian, B. Yoshua, and C. Aaron. *Deep Learning*. MIT Press, 2016.
- [14] A. Immer, C. Schultheiss, J. E. Vogt, B. Schölkopf, P. Bühlmann, and A. Marx. On the Identifiability and Estimation of Causal Location-Scale Noise Models. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2023.
- [15] G. Keropyan, D. Strieder, and M. Drton. Rank-Based Causal Discovery for Post-Nonlinear Models. In *Proceedings of The International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 7849–7870, 2023.
- [16] I. Khemakhem, R. Monti, R. Leech, and A. Hyvarinen. Causal Autoregressive Flows. In *Proceedings of The International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 3520–3528, 2021.
- [17] P. Lippe, T. Cohen, and E. Gavves. Efficient Neural Causal Discovery without Acyclicity Constraints. In *Proceedings of The International Conference on Learning Representations (ICLR)*, 2022.
- [18] L. Lorch, J. Rothfuss, B. Schölkopf, and A. Krause. DiBS: Differentiable Bayesian Structure Learning. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 24111–24123, 2021.
- [19] L. Lorch, S. Sussex, J. Rothfuss, A. Krause, and B. Schölkopf. Amortized Inference for Causal Structure Learning. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 35, pages 13104–13118, 2022.
- [20] J. M. Mooij, J. Peters, D. Janzing, J. Zscheischler, and B. Schölkopf. Distinguishing cause from effect using observational data: methods and benchmarks. *Journal of Machine Learning Research*, 17:1103–1204, 2016.
- [21] J. M. Ogarrio, P. Spirtes, and J. Ramsey. A Hybrid Causal Search Algorithm for Latent Variable Models. In *Proceedings of the International Conference on Probabilistic Graphical Models*, pages 368–379, 2016.
- [22] G. Papamakarios, E. Nalisnick, D. J. Rezende, S. Mohamed, and B. Lakshminarayanan. Normalizing flows for probabilistic modeling and inference. *Journal of Machine Learning Research*, 22:2617–2680, 2021.
- [23] J. Peters and P. Bühlmann. Structural intervention distance for evaluating causal graphs. *Neural Computation*, 27:771–799, 2015.
- [24] J. Peters, J. M. Mooij, D. Janzing, and B. Schölkopf. Causal Discovery with Continuous Additive Noise Models. *Journal of Machine Learning Research*, 15:2009–2053, 2014.
- [25] S. Ren, H. Yin, M. Sun, and P. Li. Causal Discovery with Flow-based Conditional Density Estimation. In *Proceedings of The IEEE International Conference on Data Mining (ICDM)*, pages 1300–1305, 2021.
- [26] T. S. Richardson. *Models of Feedback: Interpretation and Discovery*. Carnegie-Mellon University, 1996.
- [27] R. W. Robinson. Counting unlabeled acyclic digraphs. In *Combinatorial Mathematics V: Proceedings of the Fifth Australian Conference, Held at the Royal Melbourne Institute of Technology, August 24–26, 1976*, pages 28–43. Springer, 1977.
- [28] P. Rolland, V. Cevher, M. Kleindessner, C. Russell, D. Janzing, B. Schölkopf, and F. Locatello. Score Matching Enables Causal Discovery of Nonlinear Additive Noise Models. In *Proceedings of the International Conference on Machine Learning (ICML)*, pages 18741–18753, 2022.
- [29] K. Sachs, O. Perez, D. Pe’er, D. A. Lauffenburger, and G. P. Nolan. Causal Protein-Signaling Networks Derived from Multiparameter Single-Cell Data. *Science*, 308:523–529, 2005.
- [30] P. Sanchez, X. Liu, A. Q. O’Neil, and S. A. Tsafaris. Diffusion Models for Causal Discovery via Topological Ordering. In *Proceedings of The International Conference on Learning Representations (ICLR)*, 2023.
- [31] S. Shimizu, P. O. Hoyer, A. Hyvärinen, and A. Kerminen. A Linear Non-Gaussian Acyclic Model for Causal Discovery. *The Journal of Machine Learning Research*, 7:2003–2030, 2006.
- [32] P. Spirtes, C. Glymour, and R. Scheines. *Causation, Prediction, and Search*. The MIT Press, 2001.
- [33] N. Tagasovska, V. Chavez-Demoulin, and T. Vatter. Distinguishing cause from effect using quantiles: Bivariate quantile causal discovery. In *Proceedings of the International Conference on Machine Learning (ICML)*, pages 9311–9323. PMLR, 2020.
- [34] I. Tsamardinos, L. E. Brown, and C. F. Aliferis. The max-min hill-climbing Bayesian network structure learning algorithm. *Machine Learning*, 65:31–78, 2006.
- [35] K. Uemura and S. Shimizu. Estimation of Post-Nonlinear Causal Models Using Autoencoding Structure. In *Proceedings of The IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 3312–3316, 2020.
- [36] K. Uemura, T. Takagi, K. Takayuki, H. Yoshida, and S. Shimizu. A Multivariate Causal Discovery based on Post-Nonlinear Model. In *Proceedings of the Conference on Causal Learning and Reasoning*, pages 826–839, 2022.
- [37] T. Verma and J. Pearl. Causal networks: semantics and expressiveness. In *Proceedings of the Conference on Uncertainty in Artificial Intelligence (UAI)*, pages 69–78, 1990.
- [38] D. Yang, G. Yu, J. Wang, Z. Wu, and M. Guo. Reinforcement Causal Structure Learning on Order Graph. In *Proceedings of The AAAI Conference on Artificial Intelligence*, 2023.
- [39] Y. Yu, J. Chen, T. Gao, and M. Yu. DAG-GNN: DAG Structure Learning with Graph Neural Networks. In *Proceedings of the International Conference on Machine Learning (ICML)*, pages 7154–7163, 2019.
- [40] K. Zhang and A. Hyvärinen. On the identifiability of the post-nonlinear causal model. In *Proceedings of the Conference on Uncertainty in Artificial Intelligence (UAI)*, pages 647–655, 2009.
- [41] K. Zhang, Z. Wang, J. Zhang, and B. Schölkopf. On estimation of functional causal models: general results and application to the post-nonlinear causal model. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 7(2):1–22, 2015.
- [42] X. Zheng, B. Aragam, P. Ravikumar, and E. P. Xing. DAGs with NO TEARS: continuous optimization for structure learning. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 9492–9503, 2018.