

Scalable Variational Causal Discovery Unconstrained by Acyclicity

Nu Hoang*, Bao Duong and Thin Nguyen

Applied Artificial Intelligence Institute (A²I²), Deakin University, Australia.

Abstract. Bayesian causal discovery offers the power to quantify epistemic uncertainties among a broad range of structurally diverse causal theories potentially explaining the data, represented in forms of directed acyclic graphs (DAGs). However, existing methods struggle with efficient DAG sampling due to the complex acyclicity constraint. In this study, we propose a scalable Bayesian approach to effectively learn the posterior distribution over causal graphs given observational data thanks to the ability to generate DAGs without explicitly enforcing acyclicity. Specifically, we introduce a novel differentiable DAG sampling method that can generate a valid acyclic causal graph by mapping an unconstrained distribution of implicit topological orders to a distribution over DAGs. Given this efficient DAG sampling scheme, we are able to model the posterior distribution over causal graphs using a simple variational distribution over a continuous domain, which can be learned via the variational inference framework. Extensive empirical experiments on both simulated and real datasets demonstrate the superior performance of the proposed model compared to several state-of-the-art baselines.

1 Introduction

Causal inference [29] offers a powerful tool for tackling critical research questions in diverse fields, such as policy decision-making, experimental design, and enhancing AI trustworthiness. However, current causal inference algorithms typically require an input of a directed acyclic graph (DAG), encapsulating causal relationships among variables of interest. Unfortunately, identifying the true causal DAG often necessitates extensive experimentation, which can be time-consuming and ethically problematic in certain situations, hindering the application of causal inference to high-dimensional problems. Therefore, there is a pressing need to explore methods for discovering the causal DAG solely from observational data, which is typically more readily available [34, 17]. Nevertheless, a major hurdle in causal discovery using observational data is the non-identifiability issue of causal models when multiple DAGs may induce the same observed data mainly due to scarce data, model misspecification, and limited capability of optimizers. Bayesian inference is a promising approach to mitigate this problem by estimating the posterior distribution over causal DAGs, allowing for capturing epistemic uncertainties in causal structure learning. Moreover, this richer representation can then be leveraged for various tasks, including active causal discovery, where we strategically collect additional data to refine our understanding of causal relationships [1, 36].

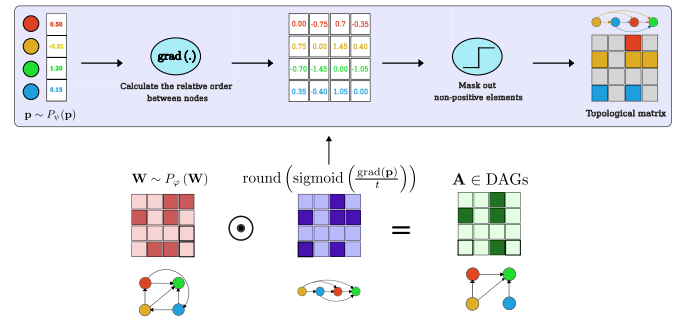


Figure 1: Overview of the proposed differentiable DAG sampling. Given two arbitrary probabilistic models over real vectors of d nodes and binary matrices of $d \times d$, we first sample a priority scores vector \mathbf{p} and a binary matrix \mathbf{W} . Then, we construct a binary adjacency matrix of a complete topological graph corresponding to \mathbf{p} , called a topological matrix, using the gradient operator followed by a tempered sigmoid function. The final DAG adjacency matrix \mathbf{A} is the element-wise multiplication of \mathbf{W} and the topological matrix derived from \mathbf{p} .

Due to the exponential explosion in the number of possible DAGs with increasing variables [7], inferring the posterior distribution becomes computationally impossible for large-scale problems. An efficient DAG sampler is therefore crucial to unlock the scalability of Bayesian causal discovery. In particular, several studies [38, 10, 3] have leveraged sequential models, e.g., Markov Chain Monte Carlo (MCMC) or GFlowNets, to sample DAGs in combinatorial spaces, leading to high computational demands. Recent advancements aim to enhance the inference efficiency through the development of gradient optimization methods for Bayesian structure learning [9, 23, 6, 2], which handle the acyclicity constraint by either integrating a smooth DAG regularization into the objective function [23] or relying on the permutation matrix-based DAG decomposition [9, 6, 2]. On the one hand, the inclusion of acyclicity regularization in the objective function may introduce additional computational costs, impeding scalability for high dimensional problems. For instance, the computational complexity for computing NOTEARS [43], a widely used DAG constraint function [39, 44, 21, 27], grows cubically with the number of nodes. In addition, DAG regularizer does not guarantee a complete DAG generation, potentially requiring additional post-processing steps [5]. On the other hand, the permutation matrix-based approach, factorizing the adjacency matrix into upper triangular and permutation matrices, may expose high complexities due to the difficulty of approximating the discrete distribution over the permutation matrix. For example, various Bayesian studies [9, 6, 2]

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

exploit the Gumbel-Sinkhorn operator, yet the process of turning a parameter matrix into a permutation matrix incurs significant computational overhead.

To address computational limitations of current approaches, we introduce a novel differentiable DAG sampling method, which maps a constraint-free distribution of implicit topological orders to a distribution over DAGs, eliminating the need for enforcing acyclicity explicitly. Inspired by [40], we sample a DAG via generating a binary adjacency matrix representing any directed graph and a vector of nodes' priority scores. These scores implicitly define an topological order when sorted, inducing a topological matrix that effortlessly ensures acyclicity. Specifically, the topological matrix can be computed promptly through the pairwise differences of a vector of nodes' priority scores with a tempered sigmoid function. Finally, the proposed approach obtains a DAG through the element-wise multiplication of the generated binary adjacency matrix and the topological matrix as shown in Figure 1, reducing the time complexity for DAG sampling to quadratic with respect to the number of nodes, enabling DAG generation with thousands of nodes in a matter of milliseconds. Based on this competent DAG sampling scheme, we are able to model the posterior distribution over DAGs using a simple variational distribution over a continuous domain, which can be learned via the variational inference framework. Our main contributions are outlined as follows:

- We present a novel differentiable probabilistic DAG model exhibiting scalable sampling for thousands of variables (Section 4.1). Furthermore, we provide theoretical justification substantiating the correctness of the proposed method.
- We introduce a fast and accurate Bayesian causal discovery method built upon on our efficient DAG sampling and the variational inference framework (Section 4.2). The proposed approach ensures the generation of valid acyclic causal structures at any time during training, accompanied by a competitive runtime compared to alternative methods.
- We demonstrate the efficiency of the proposed method through extensive numerical experiments on both synthetic and real datasets (Section 5)¹. The empirical results underscore the scalability of our DAG sampling model and showcase its superior performance in Bayesian causal discovery when incorporated with the variational inference framework compared with several baselines.

2 Related Work

Discrete Optimization. These methods encompass constraint-based methods, score-based methods, and hybrid methods, which typically search for the true causal graph within the original combinatorial space of DAGs. Constraint-based methods [34, 41, 8] depend on results from various conditional independence tests, while score-based methods [11, 12, 15] optimize a predefined score to identify the final DAG by adding, removing, or reversing edges. In the meantime, hybrid methods [37, 28] integrate both constraint-based and score-based techniques to trim the search space, thus accelerating the overall learning process.

Continuous Optimization. Optimizing in the discrete space of DAGs is known to be challenging due to the super exponential increase in complexity with the number of variables. To address this issue, several studies map the discrete space to a continuous one, thereby unlocking the application of various continuous optimization techniques. A pioneering study is NO TEARS [43], which introduced a smooth function to evaluate the DAG-ness of a weighted adjacency

matrix. The causal structure learning problem is then tackled using an augmented Lagrangian optimization method. Subsequent studies, inspired by NO TEARS, have enhanced its efficiency by introducing low-complexity DAG constraints [22, 42], or extending it for non-linear functional models [39, 44, 21, 27]. In contrast to NO TEARS, recent studies [40, 24] introduce various DAG mapping functions, which facilitate direct optimization within the DAG space. Therefore, these mapping functions provide more scalable and direct approaches without the need to evaluate the DAG constraint.

Bayesian Causal Structure Learning. The above studies usually output the Markov equivalence class (MEC) of the true DAG or a single DAG, which may not adequately represent the uncertainty in certain practical scenarios. To address this challenge, Bayesian causal discovery methods produce a posterior distribution over causal DAGs. Several studies demonstrate DAG sampling in the discrete space using either Markov Chain Monte Carlo (MCMC) [35, 38, 20] or GFlowNets [10, 3]. However, these approaches expose slow mixing and convergence. Recent advancements aim for more efficient inference through the development of gradient optimization methods for Bayesian structure learning. However, existing studies still struggle in representing the acyclicity constraint. For instance, DiBS [23] exploits NO TEARS as a DAG regularizer and utilizes Stein variational approach to learn the joint distribution over DAGs and causal model parameters. However, its scalability is limited for large graphs due to the computational complexity associated with NO TEARS. In contrast, BCDnets [9] and DDS [6] exploit the ordering-based DAG decomposition to parameterize the DAGs distribution through the multiplication of upper triangular and permutation matrices. For approximating a discrete distribution over the permutation matrix, they utilized Gumbel-Sinkhorn operator, which poses a high time complexity, i.e., cubic with respect to the number of node. To reduce the complexity of the Gumbel-Sinkhorn operator, BayesDAG [2] exploits No-Curl constraint [40] which can decrease the number of iterations required for the Gumbel-Sinkhorn operator, yet the scalability of this approach remains a challenge, when dealing with large graphs.

3 Preliminary

3.1 Problem formulation

Let $\mathbf{X} \in \mathbb{R}^{n \times d}$ be an observational dataset consisting of n i.i.d. samples of d random variables from a joint distribution $P(X)$. The marginal joint distribution $P(X)$ factorizes according to a DAG $\mathcal{G} = \langle \mathbf{V}, \mathbf{E} \rangle$, where $\mathbf{V} = \{1, 2, \dots, d\}$ is a set of nodes corresponding to d random variables and \mathbf{E} is a set of edges representing the dependency between nodes. In other words, $P(X) = \prod_{i=1}^d P(X_i | X_{\text{pa}(i)})$, where $X_{\text{pa}(i)}$ denotes a set of parents of nodes X_i . We can model the data of a node X_i with a structural equation model as follows:

$$X_i = g_i(f_i(X_{\text{pa}(i)}), \epsilon_i), \quad (1)$$

where g_i and f_i are deterministic functions and ϵ_i is an arbitrary noise. In this work, we consider a Gaussian additive noise model (ANM) [33, 14] as follows:

$$X_i = f_i(X_{\text{pa}(i)}) + \epsilon_i, \epsilon_i \sim \mathcal{N}(0, \sigma^2), \quad (2)$$

The DAG \mathcal{G} can be represented by a binary adjacency matrix $\mathbf{A} \in \{0, 1\}^{d \times d}$, where $A_{ij} = 1$ indicates an edge from X_i to X_j . As a result, Eq. (1) can be rewritten as:

$$X_i = f_i(A_i \circ X) + \epsilon_i, \quad (3)$$

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

where A_i is the i^{th} column in the adjacency matrix \mathbf{A} . Intuitively, the matrix \mathbf{A} plays as a mask to extract parental nodes of X_i . Using a binary adjacency matrix \mathbf{A} to represent a DAG offers a more flexible approach to model the functional relationship between each node X_i and its parents.

Given \mathbf{X} , we aim to learn the adjacency matrix \mathbf{A} of the DAG \mathcal{G} . Let \mathbb{D} be the DAG space of d nodes. Score-based methods usually solve an optimization problem by maximizing a fitness score of a candidate graph and the data, i.e., $F(\mathbf{A}, \mathbf{X})$:

$$\max_{\mathbf{A}} F(\mathbf{A}, \mathbf{X}) \text{ s.t. } \mathbf{A} \in \mathbb{D}. \quad (4)$$

3.2 Bayesian causal structure learning

Solving Eq. (4) yields a single point DAG solution that comes with practical limitations, particularly in addressing the non-identifiability problem in DAG learning. This limitation stems from the fact that the true causal DAG is only identifiable under specific conditions. For instance, identifiability in the linear Gaussian SEM holds in the equal variance noise setting [30]. In real-world scenarios, the non-identifiability issue may surface due to the limited number of observations, leading the point estimation approach to converge toward an incorrect solution. Considering these challenges, it proves beneficial to model the uncertainty in DAG learning, where Bayesian learning emerges as a standard approach.

The ultimate goal of Bayesian causal structure learning methods is to approximate the posterior distribution over the causal graph given the observational data, denoted as $P(\mathcal{G} | \mathbf{X})$. Using Bayes' rule, the posterior $P(\mathcal{G} | \mathbf{X})$ can be expressed through the prior distribution $P(\mathcal{G})$ and the marginal likelihood $P(\mathbf{X} | \mathcal{G})$ as follows:

$$P(\mathcal{G} | \mathbf{X}) = \frac{P(\mathbf{X} | \mathcal{G})P(\mathcal{G})}{\sum_{\mathcal{G}} P(\mathbf{X} | \mathcal{G})P(\mathcal{G})}, \quad (5)$$

where the marginal likelihood $P(\mathbf{X} | \mathcal{G})$ is defined as a marginalization of the likelihood function over all possible parameters for \mathcal{G} :

$$P(\mathbf{X} | \mathcal{G}) = \int P(\mathbf{X} | \mathcal{G}, \theta_{\mathcal{G}})P(\theta_{\mathcal{G}} | \mathcal{G})d\theta_{\mathcal{G}}. \quad (6)$$

The main challenge in Bayesian causal discovery is the intractability of the denominator in Eq. (5) due to the expansive space of DAGs.

3.3 DAG representation

To deal with the computational challenge in DAG learning, several studies [43, 22, 42] have shifted the combinatorial search to a continuous optimization problem, proving to be more scalable and adaptable to different SEMs. Specifically, researchers have introduced various smooth functions to evaluate whether a directed adjacency matrix represents a DAG, i.e., $\mathbf{A} \in \mathbb{D} \Leftrightarrow h(\mathbf{A}) = 0$. Consequently, solving Eq. (4) can be accomplished through a constrained optimization method, such as the augmented Lagrangian method. Despite these advancements, existing methods still fall short in ensuring a valid acyclic output, possibly requiring additional post-processing steps [5]. Therefore, we aim to find a convenient representation for the DAG space that allows direct optimization within it. This approach guarantees the output of a DAG at any stage during the learning process.

A common approach for representing a DAG involves utilizing a permutation matrix $\mathbf{\Pi} \in \{0, 1\}^{d \times d}$ and an upper triangular matrix $\mathbf{U} \in \{0, 1\}^{d \times d}$, i.e., $\mathbf{A} = \mathbf{\Pi}^T \mathbf{U} \mathbf{\Pi}$. This formulation arises from

the inherent property of a DAG, where there exists at least one valid permutation (or causal order) $\pi \in \mathbb{R}^d$, implying that there is no direct edge from node $X_{\pi(j)}$ and $X_{\pi(i)}$ if $\pi(i) < \pi(j)$. To seamlessly integrate this formula into a continuous DAG learning framework, we can parameterize the permutation matrix and the upper triangular matrix using Gumbel-Sinkhorn [25] and Gumbel-Softmax [16] distributions, respectively. Yet, this approach introduces a high complexity (i.e., $\mathcal{O}(d^3)$) due to the Hungarian algorithm [19] used in the forward pass of Gumbel-Sinkhorn.

Motivated by the same objective, [40] introduce No-Curl, a simple DAG mapping function that facilitates a projection of any weighted adjacency matrix from an arbitrary directed graph to the DAG space through a straightforward process: applying element-wise multiplication with the gradient of a priority scores vector associated with the graph vertices. We will provide a detailed explanation of the No-Curl characterization, as it constitutes a crucial component of our proposed method.

Let $\mathbf{p} \in \mathbb{R}^d$ be a vector representing priority scores of d nodes in a directed graph \mathcal{G} . The gradient of \mathbf{p} , denoted by $\text{grad}(\mathbf{p}) : \mathbb{R}^d \rightarrow \mathbb{R}^{d \times d}$, is defined as follows:

$$\text{grad}(\mathbf{p})_{ij} = \mathbf{p}_j - \mathbf{p}_i. \quad (7)$$

Then, No-Curl proposes mapping \mathbf{p} and a weighted directed adjacency matrix \mathbf{W} , denoted by $\gamma(\mathbf{W}, \mathbf{p}) : \mathbb{R}^{d \times d} \times \mathbb{R}^d \rightarrow \mathbb{R}^{d \times d}$, to the DAG space by:

$$\gamma(\mathbf{W}, \mathbf{p}) = \mathbf{W} \circ \text{ReLU}(\text{grad}(\mathbf{p})), \quad (8)$$

where ReLU is the rectified linear unit activation function. For a detailed proof of No-Curl, we refer to [40]. Here, we provide an intuitive explanation of the mapping $\gamma(\mathbf{W}, \mathbf{p})$: the priority scores vector \mathbf{p} implicitly corresponds to a causal order of d nodes when we sort \mathbf{p} in increasing order. To avoid cycles, we exclusively permit edges from nodes with lower scores to nodes with higher scores, i.e., $p_i < p_j \Rightarrow (i \rightarrow j)$, which is equivalent to $\text{grad}(\mathbf{p})_{ij} > 0$. Hence, $\text{ReLU}(\text{grad}(\mathbf{p}))$ eliminates cycle-inducing edges by zeroing out all $\text{grad}(\mathbf{p})_{ij} \leq 0$. Finally, we remove spurious edges through a Hadamard product with a weighted adjacency matrix \mathbf{W} .

Unfortunately, the No-Curl characterization is specifically crafted for a weighted DAG representation. Indeed, we favor a binary representation, since it offers greater flexibility to model both linear and non-linear functional relationships. To address this preference, we propose a novel adaptation of the No-Curl approach for representing binary adjacency matrices of DAGs, as detailed in the next section.

4 Proposed method

The main goal of this study is to propose a scalable and fully differentiable framework to approximate the posterior distribution over DAGs given observational data. To this end, we first introduce a novel measure to parameterize the probabilistic model over DAGs by extending the No-Curl characterization. Leveraging this probabilistic model for DAG sampling and integrating it with the variational inference framework, we then introduce a new Bayesian causal discovery method, named VCUDA (Variational Causal Discovery Unconstrained by Acylicity), offering precise capture and effective generation of samples from the complex posterior distribution of DAGs.

4.1 Differentiable DAG sampling

As discussed in the previous section, a weighted adjacency matrix to represent a DAG does not align with our purpose. Therefore, we extend the No-Curl characterization by substituting the $\text{ReLU}(\cdot)$ by a tempered sigmoid(\cdot) function, allowing us to represent a DAG using a binary adjacency matrix.

Theorem 1. Let $\mathbf{A} \in \{0, 1\}^{d \times d}$ be an adjacency matrix of a graph of d nodes. Then, \mathbf{A} is DAG if and only if there exists a vector of priority scores $\mathbf{p} \in \mathbb{R}^d$ and a corresponding binary matrix $\mathbf{W} \in \{0, 1\}^{d \times d}$ such that:

$$\mathbf{A} = \nu(\mathbf{W}, \mathbf{p}) = \mathbf{W} \circ \lim_{t \rightarrow 0} \text{sigmoid}\left(\frac{\text{grad}(\mathbf{p})}{t}\right)$$

where $t > 0$ is a strictly positive temperature and \mathbf{p} contains no duplicate elements, i.e., $p_i \neq p_j \forall i, j$.

Proof. See the Appendix [13] for more details. \square

Closely related to our method, COSMO [24] also introduces a smooth orientation matrix for unconstrained DAG learning. It is crucial to highlight that our study is motivated by distinct objectives, specifically, addressing challenges related to scalability and generalization in Bayesian causal discovery. This convergence in ideas reaffirms the significance of our approach in independently tackling common challenges, underscoring its broader applicability and relevance.

Based on Theorem 1, we further introduce a new probabilistic model over DAGs space as follows:

$$P(\mathbf{A}) = \sum_{\mathbf{W}} \int_{\mathbf{p}} P(\mathbf{W}, \mathbf{p}) d\mathbf{p} \quad \text{s.t. } \mathbf{A} = \mathbf{W} \circ \lim_{t \rightarrow 0} \text{sigmoid}\left(\frac{\text{grad}(\mathbf{p})}{t}\right) \quad (9)$$

where $P(\mathbf{W})$ and $P(\mathbf{p})$ are distributions over edges and priority scores, respectively. As a result, Eq. (9) provides a fast and assured sampling approach of DAGs without evaluating any explicit acyclicity constraints. We follow [6], utilizing the Gumbel-Softmax [16] to model the discrete distribution over edges. Let $\varphi \in [0, 1]$ be the probability of the existence of an edge from node X_i to X_j . The Gumbel-Softmax, which is a continuous distribution, enables a differentiable approximation of samples from a discrete distribution, e.g., Bernoulli distribution: $\hat{W}_{ij} \in [0, 1] \sim \text{Gumbel-Softmax}_{\tau}(\varphi_{ij})$, where $\tau > 0$ is the temperature parameter controlling the smoothness of the categorical during sampling. For example, a low value of τ generates more likely one-hot encoding samples, making the Gumbel-Softmax distribution resemble the original categorical distribution. Consequently, we can directly generate a DAG by sampling an edge matrix from a Gumbel-Softmax distribution $\mathbf{W} \sim \text{Gumbel-Softmax}(\varphi)$, and a priority score vector from a multivariate distribution $\mathbf{p} \sim P_{\psi}(\mathbf{p})$: $\mathbf{A} \sim P_{\varphi, \psi}(\mathbf{A})$ where φ and ψ are parameters defined distributions of \mathbf{W} and \mathbf{p} , respectively.

Computational complexity: Our proposed probabilistic model significantly speeds up the DAG sampling time compared with related studies using the Gumbel-Sinkhorn approach, such as BCD-nets [9]. To elaborate, our proposed approach requires $\mathcal{O}(d^2)$ for sampling the edge matrix \mathbf{W} and $\mathcal{O}(d)$ for sampling the priority scores vector \mathbf{p} . This leads to an overall computational complexity of $\mathcal{O}(d^2)$. A closely related study to our approach is BayesDAG [2], which suggests replacing $\text{ReLU}(\cdot)$ with $\text{Step}(\cdot)$. However, their

Algorithm 1: VCUDA (Variational Causal Discovery Unconstrained by Acyclicity)

Input: Observational dataset \mathbf{X} ; prior distributions $P_{\text{prior}}(\mathbf{W})$, $P_{\text{prior}}(\mathbf{p})$; temperature t , regularizers λ_1 , λ_2 , training iterations T

Initialize parameters φ, ψ, θ ;

for $i = 1 \dots T$ **do**

for $\mathbf{X}_{\text{batch}} \in \mathbf{X}$ **do**

 Sample $\mathbf{W} \sim P_{\varphi}(\mathbf{W})$;

 Sample $\mathbf{p} \sim P_{\psi}(\mathbf{p})$;

 Compute $\mathbf{A} = \mathbf{W} \circ \text{sigmoid}\left(\frac{\text{grad}(\mathbf{p})}{t}\right)$;

 Compute $\hat{\mathbf{X}} = f_{\theta}(\mathbf{X}_{\text{batch}}, \mathbf{A})$;

 Maximize ELBO loss (Eq. (21)) w.r.t φ, ψ, θ ;

end

end

return φ, ψ, θ

approach encounters the problem of uninformative gradients due to the intrinsic property of $\text{Step}(\cdot)$. Consequently, [2] still utilizes the Gumbel-Sinkhorn operator to approximate the distribution over permutation matrices, incurring high complexity at $\mathcal{O}(d^3)$.

4.2 Variational Inference DAG Learning

As mentioned earlier, Bayesian causal structure learning aims to determine the posterior distribution over DAGs. However, directly computing the posterior becomes infeasible due to the intractability of the marginal data distribution. To address this challenge, we turn to variational inference. Here, we leverage the probabilistic DAG model $P_{\varphi, \psi}(\mathbf{A})$ introduced in Section 4.1 to approximate the true posterior distribution $P(\mathbf{A} | \mathbf{X})$. In essence, we aim to optimize the variational parameters to minimize the KL divergence between the approximate and true posterior distributions that are equivalent to maximizing the evidence lower bound (ELBO) objective as follows:

$$\max_{\theta, \varphi, \psi} \mathcal{L} = \underbrace{\mathbb{E}_{\mathbf{W}, \mathbf{p} \sim P_{\varphi, \psi}(\mathbf{W}, \mathbf{p})} [\log P_{\theta}(\mathbf{X} | \mathbf{W}, \mathbf{p})]}_{(i)} - \underbrace{D_{\text{KL}}(P_{\varphi, \psi}(\mathbf{W}, \mathbf{p}) \| P_{\text{prior}}(\mathbf{W}, \mathbf{p}))}_{(ii)}. \quad (10)$$

The objective in Eq. (10) consists of two terms: i) the first is the log-likelihood of the data given the causal structure model and ii) the second is the KL divergence between the approximate posterior distribution and the prior distribution. With appropriate choices of variational families and prior models, the optimized parameters θ, φ, ψ from Eq. (10) minimize the divergence between the approximate distribution and the true distribution, i.e., $P_{\varphi, \psi}(\mathbf{A}) \approx P(\mathbf{A} | \mathbf{X})$.

To compute (i), we begin by sampling a DAG adjacency matrix \mathbf{A} from the approximate distribution in each iteration. For every node X_i , we reconstruct its values by applying masking on the observed data \mathbf{X} with the sampled \mathbf{A} . This is then followed by a transformation $f_{i, \theta}$, parameterized using neural networks:

$$\hat{X}_i = f_{i, \theta}(A_i \circ \mathbf{X}), \quad (11)$$

where A_i is the i^{th} column in the adjacency matrix \mathbf{A} . By assuming that the data has a Gaussian distribution with unit variance, we can approximate the first term by the least square loss, i.e., $\|\mathbf{X} - \hat{\mathbf{X}}\|^2$.

To compute (ii), we initially employ a mean-field factorization for the variational model, i.e., $P_{\varphi, \psi}(\mathbf{W}, \mathbf{p}) = P_{\varphi}(\mathbf{W}) P_{\psi}(\mathbf{p})$. This

mean-field factorization provides us a convenient way to calculate the KL divergence, represented as:

$$D_{KL}(P_{\varphi,\psi}(\mathbf{W}, \mathbf{p}) \parallel P_{\text{prior}}(\mathbf{W}, \mathbf{p})) \quad (12)$$

$$= \mathbb{E}_{\mathbf{p}, \mathbf{W} \sim P_{\varphi,\psi}(\mathbf{p}, \mathbf{W})} \left[\log \frac{P_{\varphi}(\mathbf{W}) P_{\psi}(\mathbf{p})}{P_{\text{prior}}(\mathbf{W}) P_{\text{prior}}(\mathbf{p})} \right] \quad (13)$$

$$= \mathbb{E}_{\mathbf{p}, \mathbf{W} \sim P_{\varphi,\psi}(\mathbf{p}, \mathbf{W})} \left[\log \frac{P_{\varphi}(\mathbf{W})}{P_{\text{prior}}(\mathbf{W})} + \log \frac{P_{\psi}(\mathbf{p})}{P_{\text{prior}}(\mathbf{p})} \right] \quad (14)$$

$$= \int \int P_{\varphi}(\mathbf{W}) P_{\psi}(\mathbf{p}) \left[\log \frac{P_{\varphi}(\mathbf{W})}{P_{\text{prior}}(\mathbf{W})} + \log \frac{P_{\psi}(\mathbf{p})}{P_{\text{prior}}(\mathbf{p})} \right] d\mathbf{W} d\mathbf{p} \quad (15)$$

$$= \int \int P_{\varphi}(\mathbf{W}) P_{\psi}(\mathbf{p}) \left[\log \frac{P_{\varphi}(\mathbf{W})}{P_{\text{prior}}(\mathbf{W})} \right] d\mathbf{W} d\mathbf{p} \quad (16)$$

$$+ \int \int P_{\varphi}(\mathbf{W}) P_{\psi}(\mathbf{p}) \left[\log \frac{P_{\psi}(\mathbf{p})}{P_{\text{prior}}(\mathbf{p})} \right] d\mathbf{W} d\mathbf{p} \quad (17)$$

$$= \int P_{\psi}(\mathbf{p}) D_{KL}(P_{\varphi}(\mathbf{W}) \parallel P_{\text{prior}}(\mathbf{W})) d\mathbf{p} \quad (18)$$

$$+ \int P_{\varphi}(\mathbf{W}) D_{KL}(P_{\psi}(\mathbf{p}) \parallel P_{\text{prior}}(\mathbf{p})) d\mathbf{W} \quad (19)$$

$$= D_{KL}(P_{\varphi}(\mathbf{W}) \parallel P_{\text{prior}}(\mathbf{W})) + D_{KL}(P_{\psi}(\mathbf{p}) \parallel P_{\text{prior}}(\mathbf{p})) \quad (20)$$

Consequently, we can compute the KL divergence between the approximate posterior distribution and the prior distribution over DAGs via the sum of the KL divergence between the variational model and the prior distribution over the edge matrix \mathbf{W} and the priority scores vector \mathbf{p} .

Variational Families: For the distribution over the priority scores vector \mathbf{p} , we opt for the isotropic Gaussian, i.e., $\mathbf{p} \sim \mathcal{N}(\boldsymbol{\mu}, \sigma^2 \mathbf{I})$. As discussed earlier, we choose the Gumbel-Softmax distribution [16] for the variational model of the edge matrix \mathbf{W} . These choices enable us to utilize the pathwise gradient, offering a lower variance approach compared to the score-function method [26]. To compute the gradient, we leverage the straight-through estimator [4]: for \mathbf{p} , we use the rounded value of sigmoid $\left(\frac{\text{grad}(\mathbf{p})}{t}\right)$ in the forward pass, and its continuous value in the backward pass. For \mathbf{W} , we use the discrete value $W_{ij} = \arg \max [1 - \hat{W}_{ij}, \hat{W}_{ij}]$ in the forward pass, and the continuous approximation \hat{W}_{ij} in the backward pass.

Prior Distribution: A well-chosen prior encapsulates existing knowledge about the model parameters, thereby guiding the inference process. In line with the belief in the sparsity of causal DAGs, we set a small prior $P_{\text{prior}}(W_{ij})$ on the edge probability. For an effective gradient estimation, we define the prior distribution of the priority scores vector as a normal distribution with a mean of zero and a small variance.

Incorporating all the above design choices, the final loss can be expressed as follows:

$$\begin{aligned} \max_{\theta, \varphi, \psi} \mathcal{L} = & - \sum (X_{ij} - \hat{X}_{ij})^2 \\ & - D_{KL}(P_{\varphi}(\mathbf{W}) \parallel P_{\text{prior}}(\mathbf{W})) \\ & - D_{KL}(P_{\psi}(\mathbf{p}) \parallel P_{\text{prior}}(\mathbf{p})), \end{aligned} \quad (21)$$

where $\hat{X}_i = f_{i,\theta}(\mathbf{A}_i \circ \mathbf{X})$ is the reconstructed data from the sampled DAG. In the implementation, we divide the total loss by the number of nodes d for stable numerical optimization. The training process of the proposed approach are summarized in Algorithm 1.

5 Experiment

In this section, we demonstrate extensive experiments showing the empirical performance of our proposed method on DAG sampling

and DAG structure learning tasks.

Baselines. Regarding DAG sampling, we compare the proposed method with the Gumbel-Sinkhorn and Gumbel-Top-k approaches [6] in terms of sampling time for thousands of variables. Unlike our method, which samples a priority score vector, these approaches focus on sampling a permutation matrix. The Gumbel-Sinkhorn method [25] leverages the Sinkhorn operator to approximate the distribution over the permutation matrix. Meanwhile, Gumbel-Top-k combines the Gumbel-Top-k distribution [18] with the Soft-Sort operator [31] to achieve faster sampling compared to Gumbel-Sinkhorn. Regarding DAG structure learning, we focus our comparison on differentiable methods and hence select five state-of-the-art causal discovery methods that belong to both point-estimations and Bayesian based baselines:

- **GraN-DAG** [21] utilizes the product of neural network computation path as a proxy for the adjacency matrix and NOTEARS [43] for the DAG constraint.²
- **Masked-DAG** [27] leverages Gumbel-Sigmoid to parameterize the binary adjacency matrix and NOTEARS [43] to impose the DAG constraint.³
- **DiBS** [23] models the Bayesian causal structure problem from the latent space of a probabilistic graph representation and employs Stein variational gradient descent to solve the problem. The study also exploits NOTEARS [43] to impose the DAG constraint on the latent space.⁴
- **DDS** [6] introduces a differentiable DAG sampling via sampling an edge matrix and a permutation matrix, which is integrated with a variational inference model to solve Bayesian causal structure learning.⁵
- **BaDAG** [2] combines both stochastic gradient Markov Chain Monte Carlo and variational inference for Bayesian causal discovery. The study leverages No-Curl constraint and Sinkhorn algorithm to sample DAGs.⁶

We obtain the original implementations and the recommended hyperparameters for these baselines.

Datasets. We benchmark these methods on both synthetic and real datasets. For synthetic datasets, we closely follow [21, 27, 6]. For generating causal DAGs, we consider Erdős-Rényi (ER) and scale-free (SF) network models with average degree equal to 1. We vary the graph size in terms of number of nodes $d = \{10, 50, 100\}$ and consider both linear and nonlinear Gaussian SEMs. For linear model, we generate a weighted adjacency matrix $\mathbf{W} \in \mathbb{R}^{d \times d}$ with edges' weights randomly sampled from $\mathcal{U}([-2, -0.5] \cup [0.5, 2])$. We then generate the data $\mathbf{X} \in \mathbb{R}^{n \times d}$ following the linear SEM: $\mathbf{X} = \mathbf{W}^T \mathbf{X} + \epsilon$, where $\epsilon \sim \mathcal{N}(0, 1)$. For nonlinear model, we generate the data following $X_i = f_i(X_{\text{pa}(i)}) + \epsilon_i$, where the functional model f_i is generated from Gaussian Process with RBF kernel of bandwidth one and $\epsilon_i \sim \mathcal{N}(0, 1)$. In our experiments, we sample 10 datasets per setting where each dataset includes a ground truth of the causal DAG's adjacency matrix, a training dataset of 1,000 samples and a held out testing dataset of 100 samples. For real datasets, we closely follow [40, 23, 27]. We use Sachs dataset [32] which measures the expression level of different proteins and phospholipids in human cells. The data contains 853 observational samples generated from a protein interaction network of 11 nodes and 17 edges.

² <https://github.com/kurowasan/GraN-DAG>

³ <https://github.com/huawei-noah/trustworthyAI>

⁴ <https://github.com/larslorch/dibs>

⁵ <https://github.com/sharpenb/Differentiable-DAG-Sampling>

⁶ <https://github.com/microsoft/Project-BayesDAG>

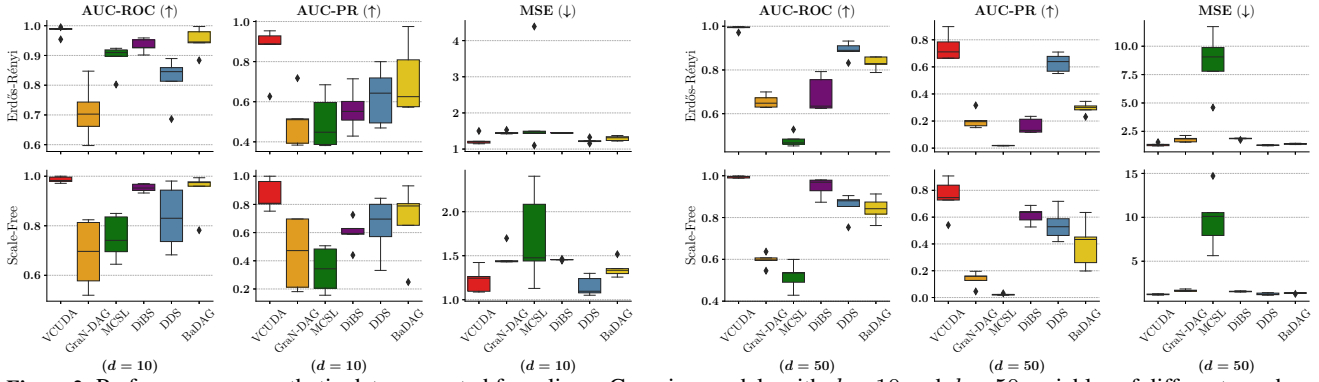


Figure 2: Performance on synthetic data generated from linear Gaussian models with $d = 10$ and $d = 50$ variables of different graph models. The reported values are aggregated from 10 independent runs. VCUDA achieves the best results across most metrics and outperforms other Bayesian approaches (DiBS and DDS). \downarrow denotes lower is better and \uparrow denotes higher is better.

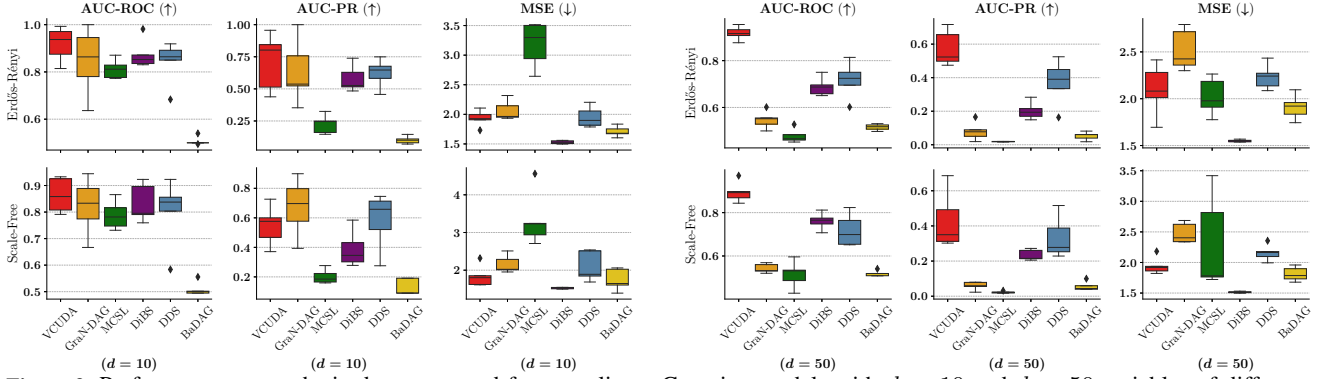


Figure 3: Performance on synthetic data generated from nonlinear Gaussian models with $d = 10$ and $d = 50$ variables of different graph models. The reported values are aggregated from 10 independent runs. Our proposed approach VCUDA achieves the best results across most metrics and outperforms other Bayesian based approaches (DiBS and DDS). \downarrow denotes lower is better and \uparrow denotes higher is better.

Evaluation metrics. We use the area under the curve of precision-recall (AUC-PR) and the area under the receiver operating characteristic curve (AUC-ROC) between the ground-truth binary adjacency matrix and the output scores matrix, denoted as \mathbf{S} , where S_{ij} represents the possibility of the presence of an edge from X_i to X_j . For point estimation methods, we get the scores matrix from the output before thresholding. For Bayesian-based methods, we get the scores matrix by averaging 100 sampled binary adjacency matrix from the learned probabilistic DAG model. We also evaluate the learned functional model $f_\theta(\cdot)$ by computing the mean squared error (MSE) between the ground-truth node value X_i and the estimated node values $\hat{X}_i = f_{i,\theta}(\mathbf{A}_i \circ \mathbf{X})$ on a held-out dataset.

Hyperparameters. We use a neural network with one hidden layer and ReLU activation to parameterize the functional models f_i in the nonlinear setting and real-world dataset. We perform a random search over the learning rate $lr \in [10^{-3}, 10^{-1}]$. The prior edge probability and scale of the priority scores vector are set to 0.01 and 0.1, respectively. Based on our ablation study (refer to Section 5.2), a temperature of 0.3 is chosen for its benefits. To prevent overfitting, VCUDA is trained by the Adam optimizer with the l2-regularization of $1e - 4$. Furthermore, early stopping is employed with validation loss checks every 10 epochs. We find that the model is convergent within 500 epochs.

5.1 DAG sampling

We compare our DAG sampling model to two well-known models, including Gumbel-Sinkhorn and Gumbel-Top-k [6] on large-scale DAGs sampling. The results are shown in the Appendix [13], indi-

cating superiority in running time of our proposed model compared to the others, paving a road for scalable Bayesian causal discovery. More importantly, our sampling method achieves a consistently high performance when integrated into the variational framework to infer the DAGs' posterior distribution, which will be shown in the following section.

5.2 DAG structure learning

Synthetic datasets. We present the results of DAG structure learning on ER/SF graphs with $d = \{10, 50\}$ for both linear and nonlinear models in Figure 2 and Figure 3, respectively. The results exhibit the superior performance of VCUDA across all settings, particularly in terms of AUC-ROC. Specifically, the AUC-ROC values of VCUDA remain consistently high, always surpassing 0.9 for linear models and 0.8 for nonlinear models. In contrast, the AUC-ROC values of the other baselines show volatility depending on the settings. Additionally, VCUDA achieves a low MSE measures, outperforming most baselines including GranDAG, MCSL, and DDS across all settings. In comparison to DiBS, the MSE measures of VCUDA are comparable for both linear and nonlinear settings. We observe a notably superior performance of Bayesian baselines compared to point-estimation baselines, especially in scenarios with higher dimensions (e.g., $d = 50$). This disparity in performance can be attributed to the Bayesian methods' capacity to capture the uncertainty effectively. For more results on denser graphs, we refer to the Appendix [13].

Furthermore, we study the performance of VCUDA and baseline methods on the high dimensional causal graph with $d = 100$ for both linear and nonlinear models. We find that DiBS is computation-

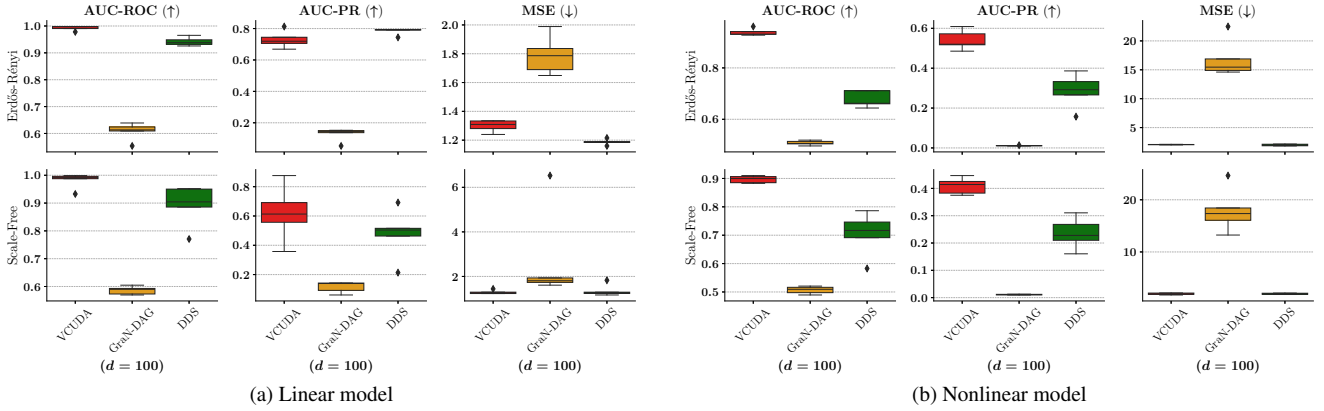


Figure 4: Performance on high dimensional data with $d = 100$ for different graphs and causal functional models. The reported values are aggregated from 10 independent runs. Our proposed approach VCUDA achieves the best results across most metrics. \downarrow denotes lower is better and \uparrow denotes higher is better.




Figure 5: The running time for causal discovery on synthetic datasets generated from a nonlinear model and ER graphs with $d = [10, 30, 50]$. VCUDA runs faster than 3 of 4 baselines, especially in high dimensions.

ally excessive producing an error in the benchmarking device, while the running time results of MCSL and BaDAG exceeded our defined time limit of 1 hour for each dataset. Therefore, Figure 4 visualizes the results of VCUDA, GraNDAG, and DDS. Among these models, VCUDA shows a consistent outperformance compared to other baselines despite the challenge of high dimensional problems.

Table 1: Performance on a real dataset of the protein signaling network. We report the mean \pm std of AUC-ROC and AUC-PR metrics. Results are averaged over 10 different restarts. \uparrow denotes higher is better. We highlight in bold the **best** result and in underline the second best result. VCUDA achieves the best AUC-ROC and the second best AUC-PR.

Method	AUC-ROC (\uparrow)	AUC-PR (\uparrow)
GraNDAG	0.57 \pm 0.02	0.26 \pm 0.03
MCSL	0.58 \pm 0.04	0.21 \pm 0.03
DiBS	0.66 \pm 0.03	0.34 \pm 0.07
DDS	0.43 \pm 0.02	0.16 \pm 0.04
BaDAG	0.48 \pm 0.01	0.17 \pm 0.02
VCUDA (Ours)	0.71 \pm 0.04	0.32 \pm 0.05

Real datasets. Table 1 displays the AUC-ROC and AUC-PR metrics on the real dataset of the protein signaling network. It is crucial to acknowledge a notable model misspecification in real-world data, given that the data might not adhere to the additive noise model assumption. Nevertheless, VCUDA achieves the best AUC-ROC and the second-best AUC-PR, which underscores the adaptability of VCUDA in navigating the intricate of real-world scenarios.

Running time. We assess the running times of all methods in the ER-nonlinear setting with varying the number of nodes. The results are presented in Figure 5. The results reveal that VCUDA demon-




Figure 6: The performance of VCUDA on different values of temperature t . The numerical results are obtained from 10 random datasets generated from a linear SEM with both ER and SF graph models.

strates significantly faster running times compared to GraN-DAG, MCSL, and DiBS, particularly as the number of nodes increases. This observation can be explained by the additional computational burden imposed on these methods due to the need to assess the DAG-ness constraint (e.g., No-TEARS) throughout the process. While VCUDA’s runtime is competitive with DDS using Gumbel-Top-k for DAG sampling, it achieves a substantial improvement in AUC-ROC, highlighting its efficacy in capturing the underlying causal structure. This translates to VCUDA being a superior solution with negligible runtime increase. Notably, even with more training iterations, DDS fails to significantly improve its overall performance.

Ablation Study. We investigate the impact of temperature t on the performance of VCUDA by evaluating different values within the range $\{0.1, 0.3, 0.5, 1.0\}$. As shown in Figure 6, lower temperatures can improve VCUDA’s performance metrics, including AUC-ROC, AUC-PR, and MSE. Based on these results, we chose a temperature of 0.3 for all the experiments due to its consistent performance.

6 Conclusion

We introduce VCUDA, a scalable approach for Bayesian causal discovery from observational data. By eliminating explicit acyclicity constraints, we propose a differentiable approach for DAGs sampling, enabling fast generation of large DAGs within seconds. In addition, the efficient sampling model enhances Bayesian inference for causal structure models when integrated into the variational inference framework. Extensive experiments on synthetic and real datasets showcase VCUDA’s superior performance, outpacing other baselines in terms of multiple metrics, all achieved with remarkable efficiency.

References

- [1] R. Agrawal, C. Squires, K. Yang, K. Shanmugam, and C. Uhler. ABCD-strategy: Budgeted experimental design for targeted causal structure discovery. In *Proceedings of the International Conference on Artificial Intelligence and Statistics*, pages 3400–3409, 2019.
- [2] Y. Annadani, N. Pawlowski, J. Jennings, S. Bauer, C. Zhang, and W. Gong. BayesDAG: Gradient-based posterior inference for causal discovery. In *Advances in Neural Information Processing Systems*, 2023.
- [3] L. Atanackovic, A. Tong, B. Wang, L. Lee, Y. Bengio, and J. Hartford. DynGFN: Towards Bayesian inference of gene regulatory networks with GFlowNets. In *Advances in Neural Information Processing Systems*, 2023.
- [4] Y. Bengio, N. Léonard, and A. Courville. Estimating or propagating gradients through stochastic neurons for conditional computation. *arXiv preprint arXiv:1308.3432*, 2013.
- [5] P. Bühlmann, J. Peters, and J. Ernest. CAM: Causal additive models, high-dimensional order search and penalized regression. *The Annals of Statistics*, 42:2526–2556, 2014.
- [6] B. Charpentier, S. Kibler, and S. Günnemann. Differentiable DAG sampling. In *Proceedings of the International Conference on Learning Representations*, 2022.
- [7] D. M. Chickering, D. Heckerman, and C. Meek. Large-sample learning of Bayesian networks is NP-hard. *Journal of Machine Learning Research*, 5:1287–1330, 2004.
- [8] D. Colombo and M. Maathuis. Order-independent constraint-based causal structure learning. *Journal of Machine Learning Research*, 15: 3741–3782, 2014.
- [9] C. Cundy, A. Grover, and S. Ermon. BCD nets: Scalable variational approaches for Bayesian causal discovery. In *Advances in Neural Information Processing Systems*, volume 34, pages 7095–7110, 2021.
- [10] T. Deleu, A. Göis, C. Emezue, M. Rankawat, S. Lacoste-Julien, S. Bauer, and Y. Bengio. Bayesian structure learning with generative flow networks. In *Proceedings of the Conference on Uncertainty in Artificial Intelligence*, pages 518–528, 2022.
- [11] D. Haughton. On the choice of a model to fit data 613 from an exponential family. *The Annals of Statistics*, 16:342–355, 1988.
- [12] D. Heckerman, D. Geiger, and D. M. Chickering. Learning Bayesian networks: The combination of knowledge and statistical data. *Machine Learning*, 20:197–243, 1995.
- [13] N. Hoang, B. Duong, and T. Nguyen. Scalable variational causal discovery unconstrained by acyclicity. *arXiv preprint arXiv:2407.04992*, Full version of this paper.
- [14] P. Hoyer, D. Janzing, J. Mooij, J. Peters, and B. Schölkopf. Nonlinear causal discovery with additive noise models. In *Advances in Neural Information Processing Systems*, volume 21, 2008.
- [15] B. Huang, K. Zhang, Y. Lin, B. Schölkopf, and C. Glymour. Generalized score functions for 625 causal discovery. In *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 1551–1560, 2018.
- [16] E. Jang, S. Gu, and B. Poole. Categorical reparameterization with Gumbel-softmax. In *Proceedings of the International Conference on Learning Representations*, 2017.
- [17] P. Jonas, J. Dominik, and S. Bernhard. *Elements of causal inference: foundations and learning algorithms*. The MIT Press, 2017.
- [18] W. Kool, H. V. Hoof, and M. Welling. Stochastic beams and where to find them: The Gumbel-Top-k trick for sampling sequences without replacement. In *Proceedings of the International Conference on Machine Learning*, pages 3499–3508, 2019.
- [19] H. W. Kuhn. The Hungarian method for the assignment problem. *Naval Research Logistics Quarterly*, 2:83–97, 1955.
- [20] J. Kuipers, P. Suter, and G. Moffa. Efficient sampling and structure learning of bayesian networks. *Journal of Computational and Graphical Statistics*, 31:639–650, 2022.
- [21] S. Lachapelle, P. Brouillard, T. Deleu, and S. Lacoste-Julien. Gradient-based neural dag learning. In *Proceedings of the International Conference on Learning Representations*, 2020.
- [22] H.-C. Lee, M. Danieleto, R. Miotto, S. Cherng, and J. Dudley. Scaling structural learning with NO-BEARS to infer causal transcriptome networks. In *Pacific Symposium On Biocomputing*, pages 391–402, 2019.
- [23] L. Lorch, J. Rothfuss, B. Schölkopf, and A. Krause. DiBS: Differentiable Bayesian structure learning. In *Advances in Neural Information Processing Systems*, volume 34, pages 24111–24123, 2021.
- [24] R. Massidda, F. Landolfi, M. Cinquini, and D. Bacciu. Constraint-free structure learning with smooth acyclic orientations. In *Proceedings of the International Conference on Learning Representations*, 2024.
- [25] G. Mena, D. Belanger, S. Linderman, and J. Snoek. Learning latent permutations with Gumbel-Sinkhorn networks. In *Proceedings of the International Conference on Learning Representations*, 2018.
- [26] S. Mohamed, M. Rosca, M. Figurnov, and A. Mnih. Monte Carlo gradient estimation in machine learning. *Journal of Machine Learning Research*, 21:132:5183–132:5244, 2020.
- [27] I. Ng, S. Zhu, Z. Fang, H. Li, Z. Chen, and J. Wang. Masked gradient-based causal structure learning. In *Proceedings of the SIAM International Conference on Data Mining*, pages 424–432, 2022.
- [28] 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.
- [29] J. Pearl. *Causality: Models, Reasoning, and Inference*. Cambridge University Press, 2009.
- [30] J. Peters and P. Bühlmann. Identifiability of Gaussian structural equation models with equal error variances. *Biometrika*, 101:219–228, 2014.
- [31] S. Prillo and J. Eisenschlos. SoftSort: A continuous relaxation for the argsort operator. In *Proceedings of the International Conference on Machine Learning*, pages 7793–7802, 2020.
- [32] 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.
- [33] S. Shimizu, P. O. Hoyer, A. Hyvärinen, and A. Kerminen. A linear non-Gaussian acyclic model for causal discovery. *Journal of Machine Learning Research*, 7:2003–2030, 2006.
- [34] P. Spirtes, C. Glymour, and R. Scheines. *Causation, Prediction, and Search*. The MIT Press, 2001.
- [35] C. Su and M. E. Borsuk. Improving structure MCMC for Bayesian networks through Markov blanket resampling. *Journal of Machine Learning Research*, 17:1–20, 2016.
- [36] P. Tigas, Y. Annadani, D. R. Ivanova, A. Jesson, Y. Gal, A. Foster, and S. Bauer. Differentiable multi-target causal bayesian experimental design. In *Proceedings of the International Conference on Machine Learning*, pages 34263–34279, 2023.
- [37] 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.
- [38] J. Viinikka, A. Hyttinen, J. Pensar, and M. Koivisto. Towards scalable bayesian learning of causal DAGs. In *Advances in Neural Information Processing Systems*, volume 33, pages 6584–6594, 2020.
- [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*, pages 7154–7163, May 2019.
- [40] Y. Yu, T. Gao, N. Yin, and Q. Ji. DAGs with No Curl: An efficient DAG structure learning approach. In *Proceedings of the International Conference on Machine Learning*, pages 12156–12166, 2021.
- [41] K. Zhang, J. Peters, D. Janzing, and B. Schölkopf. Kernel-based conditional independence test and application in causal discovery. In *Proceedings of the Conference on Uncertainty in Artificial Intelligence*, pages 804–813, 2011.
- [42] Z. Zhang, I. Ng, D. Gong, Y. Liu, E. M. Abbasnejad, M. Gong, K. Zhang, and J. Q. Shi. Truncated matrix power iteration for differentiable DAG learning. In *Advances in Neural Information Processing Systems*, 2022.
- [43] X. Zheng, B. Aragam, P. K. Ravikumar, and E. P. Xing. DAGs with NO TEARS: Continuous optimization for structure learning. In *Advances in Neural Information Processing Systems*, volume 31, 2018.
- [44] X. Zheng, C. Dan, B. Aragam, P. Ravikumar, and E. Xing. Learning sparse nonparametric DAGs. In *Proceedings of the International Conference on Artificial Intelligence and Statistics*, pages 3414–3425, 2020.