Differentiable Cyclic Causal Discovery Under Unmeasured Confounders

Muralikrishnna G. Sethuraman School of Electrical & Computer Engineering Georgia Institute of Technology muralikgs@gatech.edu Faramarz Fekri School of Electrical & Computer Engineering Georgia Institute of Technology faramarz.fekri@ece.gatech.edu

Abstract

Understanding causal relationships between variables is fundamental across scientific disciplines. Most causal discovery algorithms rely on two key assumptions: (i) all variables are observed, and (ii) the underlying causal graph is acyclic. While these assumptions simplify theoretical analysis, they are often violated in real-world systems, such as biological networks. Existing methods that account for confounders either assume linearity or struggle with scalability. To address these limitations, we propose DCCD-CONF, a novel framework for differentiable learning of nonlinear cyclic causal graphs in the presence of unmeasured confounders using interventional data. Our approach alternates between optimizing the graph structure and estimating the confounder distribution by maximizing the loglikelihood of the data. Through experiments on synthetic data and real-world gene perturbation datasets, we show that DCCD-CONF outperforms state-of-the-art methods in both causal graph recovery and confounder identification. Additionally, we also provide consistency guarantees for our framework, reinforcing its theoretical soundness.

1 Introduction

Modeling cause-effect relationships between variables is a fundamental problem in science [1, 2, 3], as it enables the prediction of a system's behavior under previously unseen perturbations. These relationships are typically represented using *directed graphs* (DGs), where nodes correspond to variables, and directed edges capture causal dependencies. Consequently, causal discovery reduces to learning the structure of these graphs.

Existing causal discovery algorithms can be broadly classified into three categories: (i) constraintbased methods, (ii) score-based methods, and (iii) hybrid methods. Constraint-based methods, such as the PC algorithm [4, 5, 6], search for causal graphs that best satisfy the independence constraints observed in the data. However, since the number of conditional independence tests grows exponentially with the number of nodes, these methods often struggle with scalability. Score-based methods, such as the GES algorithm [7, 8], learn graph structures by maximizing a penalized score function, such as the Bayesian Information Criterion (BIC), over the space of graphs. Given the vast search space, these methods often employ greedy strategies to reduce computational complexity. A significant breakthrough came with Zheng et al. [9], who introduced a continuous constraint formulation to restrict the search space to acyclic graphs, inspiring several extensions [10, 11, 12, 13, 14, 15] that frame causal discovery as a continuous optimization problem under various model assumptions. Hybrid methods [16, 17, 18] integrate aspects of both constraint-based and score-based approaches, leveraging independence constraints while optimizing a score function.

Most causal structure learning methods assume (i) a *directed acyclic graph* (DAG) with no directed cycles and (ii) complete observability, meaning no unmeasured confounders.

While these assumptions simplify the search space, they are often unrealistic, as real-world systems—especially in biology—frequently exhibit feedback loops and hidden confounders [19]. Enforcing these constraints can also increase computational complexity, particularly in ensuring acyclicity, which often requires solving challenging combinatorial or constrained optimization problems. These limitations hinder the practical applicability of existing methods in settings where such violations are unavoidable.

Several approaches have been developed to address the challenge of feedback loops within causal graphs. Early work by Richardson [20] extended constraint-based approaches for Directed Acyclic Graphs (DAGs) to accommodate directed cycles. Another key contribution came from Lacerda et al. [21], who generalized Independent Component Analysis (ICA)-based causal discovery to handle linear non-Gaussian cyclic graphs. More recently, a



Figure 1: (a) Example of a directed mixed graph \mathcal{G} , where the bidirectional edges represent hidden confounders, with σ_{ij} indicating their corresponding strengths; (b) Mutilated graph, $do(I_k)(\mathcal{G})$, resulting from the interventional experiment $I_k = \{X_3\}$, where all incoming edges (including bidirectional edges) to X_3 are removed.

growing body of research has focused on score-based methods for learning cyclic causal graphs [22, 23, 24, 25]. Additionally, some approaches leverage interventional data to improve structure recovery in cyclic systems. For instance, Hyttinen et al. [26] and Huetter and Rigollet [22] introduced frameworks that explicitly incorporate interventions to refine cyclic graph estimation. Sethuraman et al. [27] further advanced this line of research by introducing a differentiable framework for learning nonlinear cyclic graphs. Unlike differentiable DAG learners that enforce acyclicity through augmented Lagrangian-based solvers, their approach sidesteps these constraints by directly modeling the data likelihood, enabling more efficient and flexible learning of cyclic causal structures. However, their method assumes the absence of unmeasured confounders, which limits its applicability in real-world settings where hidden confounders are often present.

Causal discovery in the presence of latent confounders has seen limited development, with most existing approaches grounded in constraint-based methodologies. Extensions of the PC algorithm, such as the Fast Causal Inference (FCI) algorithm [28], construct a *Partial Ancestral Graph* (PAG) to represent the equivalence class of DAGs in the presence of unmeasured confounders, and can accommodate nonlinear dependencies depending on the chosen conditional independence tests. However, standard FCI does not incorporate interventional data, prompting extensions such as JCI-FCI [29] and related approaches [30] that combine observational and interventional settings. Jaber et al. [31] further advanced this line of work by allowing for unknown interventional targets. Additionally, Forré and Mooij [32] introduced σ -separation, a generalization of *d*-separation, enabling constraint-based causal discovery in the presence of both cycles and latent confounders. A few recent approaches, such as Bhattacharya et al. [33], have explored continuous optimization frameworks using differentiable constraints, though these methods are currently limited to linear settings. Overall, a unified framework capable of handling nonlinearity, cycles, latent confounders, and interventions remains largely absent.

Contributions. In this work, we tackle three key challenges in causal discovery: *directed cycles*, *nonlinearity*, and *unmeasured confounders*. Our main contributions are:

- We introduce DCCD-CONF, a novel differentiable causal discovery framework for learning nonlinear cyclic relationships under Gaussian exogenous noise, with confounders modeled as correlations in the noise term.
- We show that exact maximization of the proposed score function results in identification of the interventional equivalence class of the ground truth graph.
- We conduct extensive evaluations, comparing DCCD-CONF with state-of-the-art causal discovery methods on both synthetic and real-world datasets.

Organization. The paper is structured as follows: Section 2 introduces the problem setup. In Section 3, we present DCCD-CONF, our differentiable framework for nonlinear cyclic causal discovery with unmeasured confounders. We then evaluate its effectiveness on synthetic and real-world datasets in Section 4. Finally, Section 5 concludes the paper.

2 Problem Setup

2.1 Structural Equations for Cyclic Causal Graphs

Let $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathcal{B})$ represent a possibly cyclic *directed mixed graph* (DMG) that encodes the causal dependencies between the variables in the vertex set $\mathcal{V} = [d]$, where $[d] = \{1, \ldots, d\}$. \mathcal{E} denotes the set of directional edges of the form $i \to j$ in \mathcal{G} , and \mathcal{B} denotes the set of bidirectional edges of the form $i \to j$ in \mathcal{G} , and \mathcal{B} denotes the set of bidirectional edges of the form $i \to j$ in \mathcal{G} , and \mathcal{B} denotes the set of bidirectional edges of the form $i \to j \in \mathcal{E}$ representing a causal relation between X_i and X_j , and the bidirectional edge $i \to j \in \mathcal{B}$ indicates the presence of a hidden confounder between X_i and X_j . Following the framework proposed by Bollen [34] and Pearl [35], we use *structural equations model* (SEM) to algebraically describe the system:

$$X_i = F_i(\boldsymbol{X}_{\text{pa}_G(i)}, Z_i), \quad i = 1, \dots, d,$$
(1)

where $\operatorname{pa}_{\mathcal{G}}(i) := \{j \in [d] : j \to i \in \mathcal{E}\}$ represents the parent set of X_i in \mathcal{G} , and $\mathbf{X}_{\operatorname{pa}_{\mathcal{G}}(i)}$ denotes the components of $\mathbf{X} = (X_1, \ldots, X_d)$ indexed by the parent set $\operatorname{pa}_{\mathcal{G}}(i)$. We exclude *self-loops* (edges of the form $X_i \to X_i$) from \mathcal{G} , as their presence can lead to identifiability challenges [36]. The function F_i , referred to as *causal mechanism*, encodes the functional relationship between X_i and its parents $\mathbf{X}_{\operatorname{pa}_{\mathcal{G}}(i)}$, and the exogenous noise variable Z_i .

The collection of exogenous noise variables $\mathbf{Z} = (Z_1, \ldots, Z_d)$ account for the stochastic nature as well as the confounding observed in the system. We make the assumption that the exogenous noise vector follows a Gaussian distributions: $\mathbf{Z} \sim \mathcal{N}(\mathbf{0}, \mathbf{\Sigma}_Z)$. Notably, if $(\mathbf{\Sigma}_Z)_{ij} \neq 0$, then variables X_i and X_j are confounded, i.e., $i \leftrightarrow j \in \mathcal{B}$. In other words, confounding is modeled through correlations in the exogenous noise variables. Intuitively, if X_i and X_j share a hidden cause, their unexplained variation (the part not accounted for by their observed parents) will tend to move together. By allowing the noise terms Z_i and Z_j to be correlated, this shared influence can be effectively captured. This formulation generalizes prior work by allowing cycles, extending both nonlinear cyclic models without unmeasured confounders that assume independent noise terms [27], and acyclic models without confounders [37].

By collecting all the causal mechanisms into the joint function $\mathbf{F} = (F_1, \ldots, F_d)$, we can then combine (1) over $i = 1, \ldots, d$ to obtain the equation

$$\mathbf{X} = \mathbf{F}(\mathbf{X}, \mathbf{Z}). \tag{2}$$

We will use (2) to represent the causal system due to its simplicity for subsequent discussion. The observed data represents a snapshot of a dynamical process where the recursive equations in (2) define the system's state at its equilibrium. Thus, in our experiments we assume that the system has reached the equilibrium state. For a given random draw of Z, the value of X is defined as the solution to (2). To that end, we assume that (2) admits a unique fixed point for any given Z. We refer to the map $f_x : X \mapsto Z$ as the forward map, and $f_z : Z \mapsto X$ as the reverse map. In Section 3.1, we show that the chosen parametric family of functions indeed guarantees the existence of a unique fixed point. Under these restrictions, the probability density of X is well defined and is given by

$$p_{\mathcal{G}}(\boldsymbol{X}) = p_{Z}(\mathbf{f}_{x}(\boldsymbol{X})) |\det\left(\mathbf{J}_{\mathbf{f}_{x}}(\boldsymbol{X})\right)|, \qquad (3)$$

where $\mathbf{J}_{\mathbf{f}_r}(\mathbf{X})$ denotes the Jacobian matrix of the function \mathbf{f}_x at \mathbf{X} .

2.2 Interventions

In our work, we consider *surgical interventions* [35], also known as *hard interventions*, where all the incoming edges to the intervened nodes are removed from \mathcal{G} . Given a set of intervened upon nodes (also known as *interventional targets*), denoted as $I \subseteq \mathcal{V}$, the structural equations in (1) are modified as follows

$$X_{i} = \begin{cases} C_{i}, & \text{if } X_{i} \in I, \\ F_{i}(\boldsymbol{X}_{\text{pa}_{\mathcal{G}}(i)}, Z_{i}), & \text{if } X_{i} \notin I, \end{cases}$$
(4)

where C_i is a random variable sampled from a known distribution, i.e., $C_i \sim p_I(C_i)$. We denote $do(I)(\mathcal{G})$ to be the mutilated graph under the intervention I (see Figure 1). Note that X_i is no longer confounded if it is intervened on.

We consider a family of K interventional experiments $\mathcal{I} = \{I_k\}_{k \in [K]}$, where I_k represents the interventional targets for the k-th experiment. Let $\mathbf{U}_k \in \{0,1\}^{d \times d}$ denote a diagonal matrix with

 $(\mathbf{U}_k)_{ii} = 1$ if $i \notin I_k$, and $(\mathbf{U}_k)_{ii} = 0$ if $i \in I_k$. Similar to the observational setting, (4) can be vectorized to obtain the following form

$$\boldsymbol{X} = \boldsymbol{\mathrm{U}}_k \boldsymbol{\mathrm{F}}(\boldsymbol{X}, \boldsymbol{Z}) + \boldsymbol{C},\tag{5}$$

where $C = (C_1, \ldots, C_d)$ is a vector with $C_i \sim p_I(C_i)$ if $i \in I_k$, and $C_i = 0$ otherwise. For the interventional targets $I_k \in \mathcal{I}$, let $\mathbf{f}_x^{(I_k)}$ denote the forward map. Similar to the observational setting, we make the following assumption on the set of interventions.

Assumption 1 (Interventional stability). Let $\mathcal{I} = {I_k}_{k \in [K]}$ be a family of interventional targets. For each $I_k \in \mathcal{I}$, the structural equations in (5) admits a unique fixed point given the exogenous noise vector \mathbf{Z} .

Thus, the probability distribution of X for the interventional targets I_k is given by

$$p_{\mathrm{do}(I_k)(\mathcal{G})}(\boldsymbol{X}) = p_I(\boldsymbol{C})p_Z\Big(\big[\mathbf{f}_x^{(I_k)}(\boldsymbol{X})\big]_{\mathcal{U}_k}\Big)\Big|\det\big(\mathbf{J}_{\mathbf{f}_x^{(I_k)}}(\boldsymbol{X})\big)\Big|,\tag{6}$$

where $\mathcal{U}_k = \{i : i \in \mathcal{V} \setminus I\}$ denotes the index of purely observed nodes, and $p_Z \left(\left[\mathbf{f}_x^{(I_k)}(\mathbf{X}) \right]_{\mathcal{U}_k} \right)$ is the marginal distribution of the combined vector Z, restricted to the components indexed by \mathcal{U}_k .

Given a family of interventions \mathcal{I} , our goal is to learn the structure of the DMG by maximizing the log-likelihood of the data, in addition to identifying the variables that are being confounded by the unmeasured confounders Z. The next section presents our approach to addressing this problem.

3 DCCD-CONF: Differentiable Cyclic Causal Discovery with Confounders

In this section, we present our framework for differentiable learning of cyclic causal structures in the presence of unmeasured confounders. We start by modeling the causal mechanisms, then define the score function used for learning, followed by a theorem that validates its correctness. Finally, we outline the algorithm for estimating the model parameters.

3.1 Modeling Causal Mechanism

We model the structural equations in (2) using *implicit flows* [38], which define an invertible mapping between x and z by solving the root of a function $\mathbf{G}(x, z) = \mathbf{0}$, where $\mathbf{G} : \mathbb{R}^{2d} \to \mathbb{R}^d$. Specifically, we take $\mathbf{G}(x, z) = x - \mathbf{F}(x, z)$. General implicit mappings, however, do not guarantee invertibility or permit efficient computation of the log-determinant required for evaluating (6). To balance expressiveness with tractability, we adopt the structured form proposed by Lu et al. [38] for the causal mechanism:

$$\mathbf{F}(\boldsymbol{x}, \boldsymbol{z}) = -\mathbf{g}_{\boldsymbol{x}}(\boldsymbol{x}) + \mathbf{g}_{\boldsymbol{z}}(\boldsymbol{z}) + \boldsymbol{z}, \tag{7}$$

where \mathbf{g}_x and \mathbf{g}_z are restricted to be contractive functions. A function $\mathbf{g} : \mathbb{R}^d \to \mathbb{R}^d$ is contractive if there exists a constant L < 1 such that $\|\mathbf{g}(\mathbf{x}) - \mathbf{g}(\mathbf{y})\| \le L \|\mathbf{x} - \mathbf{y}\|$ for all $\mathbf{x}, \mathbf{y} \in \mathbb{R}^d$. This contractiveness ensures that the associated implicit map is uniquely solvable and invertible (see Theorem 1 in [38]).

Under this formulation, the forward map takes the form $\mathbf{f}_x(\mathbf{x}) = (\mathbf{id} + \mathbf{g}_z)^{-1} \circ (\mathbf{id} + \mathbf{g}_x)(\mathbf{x})$, where \mathbf{id} denotes the identity map. Given \mathbf{x} (or \mathbf{z}), the corresponding value of \mathbf{z} (or \mathbf{x}) can be computed via a root-finding procedure, i.e., $\mathbf{z} = \text{RootFind}(\mathbf{x} - \mathbf{F}(\mathbf{x}, \cdot))$, specifically, we employ a quasi-Newton method (i.e., Broyden's method [39]) to find the roots. To capture more complex nonlinear interactions between the observed variables \mathbf{X} and latent confounders \mathbf{Z} , multiple such implicit blocks can be stacked. For simplicity, we focus on a single implicit flow block for subsequent discussion.

We parameterize the functions \mathbf{g}_x and \mathbf{g}_z using neural networks. The adjacency matrix of the causal graph \mathcal{G} is encoded as a binary matrix $\mathbf{M}^{\mathcal{G}} \in \{0,1\}^{d \times d}$, representing the presence of directed edges and serving as a mask on the inputs to \mathbf{g}_x . The diagonal entries of $\mathbf{M}^{\mathcal{G}}$ are explicitly enforced to be zero to prevent self-loops. Similarly, the identity matrix is used to mask the inputs to \mathbf{g}_z . Consequently, the causal mechanism is defined as:

$$[\mathbf{F}_{\boldsymbol{\theta}}(\boldsymbol{X}, \boldsymbol{Z})]_{i} = \left[-\operatorname{NN}(\mathbf{M}_{*,i}^{\mathcal{G}} \odot \boldsymbol{X} \mid \boldsymbol{\theta}_{x}) + \operatorname{NN}(Z_{i} \mid \boldsymbol{\theta}_{z}) + Z_{i}\right]_{i},$$
(8)

where NN($\cdot \mid \theta$) denotes a fully connected neural network parameterized by θ , \odot denotes the Hadamard product, and $\mathbf{M}_{*,i}^{\mathcal{G}}$ is the *i*-th column of $\mathbf{M}^{\mathcal{G}}$. The contractivity of \mathbf{g}_x and \mathbf{g}_z can be enforced by rescaling their weights using spectral normalization [40]. Moreover, the contractive nature of the causal mechanism facilitates efficient computation of the score function used for learning causal graphs, as discussed in Section 3.2.

While the contractivity assumption may seem restrictive, it ensures stability and well-posedness in the presence of directed cycles. If the causal graph is known to be acyclic, this assumption can be relaxed (see Appendix C.1).

3.2 Score function

Given a family of interventions $\mathcal{I} = \{I_k\}_{k \in [K]}$, we would like to learn the parameters of the structural equation model, i.e., causal graph structure, causal mechanism, and confounder distribution. To that end, similar to prior work [27, 41, 15] in this domain we employ regularized log-likelihood of the observed nodes as the score function to be maximized. That is,

$$S_{\mathcal{I}}(\mathcal{G}) := \sup_{\boldsymbol{\theta}, \boldsymbol{\Sigma}_{Z}} \sum_{k=1}^{K} \mathbb{E}_{\boldsymbol{X} \sim p^{(k)}} \log p_{\operatorname{do}(I_{k})(\mathcal{G})}(\boldsymbol{X}) - \lambda |\mathcal{G}|$$
(9)

where $p^{(k)}$ is the data generating distribution for the k-th interventional experiment I_k , Σ_Z is the parameter (covariance matrix) governing the confounder distribution p_Z , $\boldsymbol{\theta} = (\boldsymbol{\theta}_x, \boldsymbol{\theta}_z)$ is the combined causal mechanism parameters, and $|\mathcal{G}|$ denotes a sparsity enforcing regularizer on the edges of \mathcal{G} , and $p_{\text{do}(I_k)(\mathcal{G})}(\boldsymbol{X})$ is given by (6).

We now present the main theoretical result of this paper. The following theorem establishes that, under appropriate assumptions, the graph $\hat{\mathcal{G}}$ estimated by maximizing (9) belongs to the same general directed Markov equivalence class (introduced by [36]) as the ground truth graph \mathcal{G}^* for each interventional setting $I_k \in \mathcal{I}$, denoted as $\hat{\mathcal{G}} \equiv_{\mathcal{I}} \mathcal{G}^*$, see Appendix A.1. Due to space constraints we provide the proof sketch below, see Appendix A.3 for complete proof of Theorem 2.

Theorem 2. Let $\mathcal{I} = \{I_k\}_{k=1}^K$ be a family of interventional targets, let \mathcal{G}^* denote the ground truth directed mixed graph, let $p^{(k)}$ denote the data generating distribution for I_k , and $\hat{\mathcal{G}} := \arg \max_{\mathcal{G}} \mathcal{S}(\mathcal{G})$. Then, under the Assumptions 1, A.13, A.14, and A.15, and for a suitably chosen $\lambda > 0$, we have that $\hat{\mathcal{G}} \equiv_{\mathcal{I}} \mathcal{G}^*$. That is, $\hat{\mathcal{G}}$ is \mathcal{I} -Markov equivalent to \mathcal{G}^* .

Proof (Sketch). Building on the characterization of general directed Markov equivalence class by Bongers et al. [36], extended to the interventional setting, we show that any graph outside this equivalence class has a strictly lower score than the ground truth graph \mathcal{G}^* . This follows from the fact that certain independencies present in the data are not captured by graphs outside the equivalence class. Combined with the expressiveness of the model class, this prevents such graphs from fitting the data properly.

If the intervention set consists of all single-node interventions, $\mathcal{I} = \{I_k\}_{k=1}^d$ with $I_k = \{k\}$, Hyttinen et al. [26] showed that the ground truth DMG can be uniquely recovered in the linear setting. Moreover, in the absence of cycles and confounders, this result extends to the nonlinear case, as demonstrated by Brouillard et al. [15]. However, determining the necessary conditions on interventional targets for perfect recovery in general DMGs with cycles and confounders remains an open problem. Nonetheless, in practice, we find that observational distribution in combination with single-node interventions across all nodes lead to perfect recovery of the ground truth, even in the nonlinear case, as shown in Section 4.

3.3 Updating model parameters

In practice, we use gradient based stochastic optimization to maximize (9). For this purpose, following Sethuraman et al. [27] and Brouillard et al. [15], the entries of adjacency matrix M_{ij} are modeled as Bernoulli random variable with parameters b_{ij} , grouped into the matrix $\sigma(\mathbf{B})$. We denote $\mathbf{M} \sim \sigma(\mathbf{B})$ to indicate that $M_{ij} \sim \text{Bern}(b_{ij})$ for all $i, j \in [d]$. In this formulation, the sparsity regularizer is $\|\mathbf{M}\|_0$, which is computationally intractable and thus we use the ℓ_1 -norm, $\|\mathbf{M}\|_1$ as a proxy. Consequently, the score function in (9) is replaced by the following relaxation:

$$\hat{\mathcal{S}}_{\mathcal{I}}(\mathbf{B}) := \sup_{\boldsymbol{\theta}, \boldsymbol{\Sigma}_{Z}} \mathbb{E}_{\mathbf{M} \sim \sigma(\mathbf{B})} \left[\sum_{k=1}^{K} \sum_{i=1}^{N_{k}} \log p_{\mathrm{do}(I_{k})(\mathcal{G})}(\boldsymbol{x}^{(i,k)}) - \lambda \|\mathbf{M}\|_{1} \right],$$
(10)

where we replace the expectation with respect to data distribution in (9) with sum over the finite samples, $\boldsymbol{x}^{(i,k)}$ represents the *i*-th data sample in the *k*-the interventional setting. We note that, since $p_Z = \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}_Z)$, the covariance of the exogenous confounder vector, $\boldsymbol{\Sigma}_Z$, is implicitly embedded within $p_{do(I_k)\mathcal{G}}(\boldsymbol{x}^{(i,k)})$ in the score function.

The optimization of the score function is carried in two steps. First, we optimize $\hat{S}(\mathbf{B})$ with respect to the neural network parameters $\boldsymbol{\theta}$ and the graph structure parameters \mathbf{B} . Next, we optimize $\hat{S}(\mathbf{B})$ with respect to the parameters of the exogenous noise distribution, Σ_Z . However, maximizing $\hat{S}_{\mathcal{I}}(\mathbf{B})$ presents two main challenges: (i) computing log $p_X(\mathbf{X})$ is computationally expensive due to the presence of $|\det(\mathbf{J}_{\mathbf{f}_x^{(I_k)}}(\mathbf{X}))|$, which requires $\mathcal{O}(d^2)$ gradient calls, and (ii) updating Σ_Z via stochastic gradients could lead to stability issues as Σ_Z may loose its positive definiteness.

We now describe how these challenges are addressed, along with the specific procedures for updating the individual model parameters.

3.3.1 Computing log determinant of the Jacobian

As discussed earlier, computing $\log |\mathbf{J}_{\mathbf{f}_x^{(I_k)}}(\mathbf{X})|$ is a significant challenge in maximizing the score function $\hat{\mathcal{S}}(\mathcal{B})$. To address this, we utilize the unbiased estimator of the log-determinant of the Jacobian introduced by Behrmann et al. [40], which is based on the power series expansion of $\log(1 + x)$. Since $\mathbf{f}_x^{(I_k)}(\mathbf{x}) = (\mathbf{id} + \mathbf{U}_k \mathbf{g}_z)^{-1} \circ (\mathbf{id} + \mathbf{U}_k \mathbf{g}_z)(\mathbf{x})$

$$\log \left| \det \left(\mathbf{J}_{\mathbf{f}_{x}^{(I_{k})}}(\boldsymbol{X}) \right) \right| = \log \left| \det \left(\mathbf{I} + \mathbf{J}_{\mathbf{U}_{k}\mathbf{g}_{x}}(\boldsymbol{X}) \right) \right| - \log \left| \det \left(\mathbf{I} + \mathbf{J}_{\mathbf{U}_{k}\mathbf{g}_{z}}(\boldsymbol{Z}) \right) \right|$$
$$= \sum_{m=1}^{\infty} \frac{(-1)^{m+1}}{m} \left[\operatorname{Tr} \left\{ \mathbf{J}_{\mathbf{U}_{k}\mathbf{g}_{x}}^{m}(\boldsymbol{X}) \right\} - \operatorname{Tr} \left\{ \mathbf{J}_{\mathbf{U}_{k}\mathbf{g}_{z}}^{m}(\boldsymbol{Z}) \right\} \right], \qquad (11)$$

where $\mathbf{I} \in \mathbb{R}^{d \times d}$ denotes the identity matrix, $\mathbf{J}_{\mathbf{U}_k \mathbf{g}_x}^m$ represents the Jacobian matrix raised to the *m*-th power, and Tr denotes the trace of matrix. The series in (11) is guaranteed to converge if the causal functions \mathbf{g}_x and \mathbf{g}_z are contractive [42].

In practice, the power series is truncated to a finite number of terms, which may introduce bias into the estimator. To mitigate this issue, we follow the stochastic approach of Chen et al. [43]. Specifically, we sample a random cut-off point $n \sim p_{\mathbb{N}}(n)$ for truncating the power series and weight the *i*-term in the finite series by the inverse probability of the series not ending at *i*. This yields the following unbiased estimator

$$\log\left|\det\left(\mathbf{J}_{\mathbf{f}_{x}^{(I_{k})}}(\boldsymbol{X})\right)\right| = \mathbb{E}_{n \sim p_{\mathbb{N}}(N)}\left[\sum_{m=1}^{n} \frac{(-1)^{m+1}}{m} \cdot \frac{\operatorname{Tr}\left\{\mathbf{J}_{\mathbf{U}_{k}\mathbf{g}_{x}}^{m}(\boldsymbol{X})\right\} - \operatorname{Tr}\left\{\mathbf{J}_{\mathbf{U}_{k}\mathbf{g}_{z}}^{m}(\boldsymbol{Z})\right\}}{p_{\mathbb{N}}(\ell \geq m)}\right].$$
(12)

The gradient calls can be reduced even further using the *Hutchinson trace estimator* [44], see Appendix B for more details.

3.3.2 Updating neural network and graph parameters.

In the first step of the parameter update, keeping Σ_Z fixed, the parameters of the neural network θ and the graph structure **B** are updated using the backpropagation algorithm with stochastic gradients. The gradient of the score function $\hat{S}_{\mathcal{I}}(\mathbf{B})$ with respect to **B** is computed using the Straight-Through Gumbel estimator. This involves using Bernoulli samples in the forward pass while computing score, and using samples from Gumbel-Softmax distribution in the backward pass to compute the gradient, which can be differentiated using the reparameterization trick [45].

3.3.3 Updating the confounder-noise distribution parameters

In second parameter update step, we fix the value of θ and B and focus on the confounder-noise distribution parameter Σ_Z . First, consider the case where no interventions are applied, i.e, $I_k =$

 \emptyset . Note that the dependence of $\hat{S}(\mathbf{B})$ on Σ_Z arises solely from p_Z , which is embedded within $p_{\mathrm{do}(I_k)(\mathcal{G})}(\mathbf{X})$. Therefore, we can thus ignore the remaining terms in $\hat{S}(\mathbf{B})$ and focus exclusively on p_Z . Let $\{\mathbf{x}^{(i)}\}_{i=1}^N$ denote the observational data. From the forward map, we have $\mathbf{z}^{(i)} = \mathbf{f}_x(\mathbf{x}^{(i)})$. Given that $p_Z = \mathcal{N}(\mathbf{0}, \Sigma_Z)$, the relevant parts of $\hat{S}(\mathbf{B})$ with respect to Σ_Z , denoted as $\tilde{\mathcal{L}}(I_k)$, are expressed as:

$$\tilde{\mathcal{L}}(I_k) = \sup_{\boldsymbol{\Sigma}_Z} \sum_{i=1}^N -\frac{1}{2} \log |\boldsymbol{\Sigma}_Z| - \frac{(\boldsymbol{z}^{(i)})^\top \boldsymbol{\Sigma}_Z^{-1} \boldsymbol{z}^{(i)}}{2}.$$
(13)

Simplifying (13) yields a more convenient form:

$$\tilde{\mathcal{L}}(I_k) = \sup_{\Sigma_Z} -\text{Tr}(\mathbf{S}\Sigma_Z^{-1}) - \log |\Sigma_Z|,$$
(14)

where $\mathbf{S} = \frac{1}{n} \sum_{i=1}^{N} \boldsymbol{z}^{(i)} (\boldsymbol{z}^{(i)})^{\top}$ is the sample covariance of \boldsymbol{Z} .

Maximizing (14) directly using backpropagation and stochastic gradients results in stability issues as Σ_Z may loose its positive definiteness. However, Friedman et al. [46] demonstrated that the sparsity-regularized version of (14) is a concave optimization problem in Σ_Z^{-1} that can be efficiently solved by optimizing the columns of Σ_Z individually. This is achieved by formulating the column recovery as a lasso regression problem. We adopt this strategy while updating the Σ_Z during the maximization of $\hat{S}(\mathbf{B})$.

Let $\mathbf{W} = \Sigma_Z$ be the estimate of the covariance matrix. We reorder \mathbf{W} such that the column and row being updated can be placed at the end, resulting in the following partition

$$\mathbf{W} = \begin{pmatrix} \mathbf{W}_{11} & \mathbf{w}_{12} \\ \mathbf{w}_{12}^{\top} & \mathbf{w}_{22} \end{pmatrix}, \quad \mathbf{S} = \begin{pmatrix} \mathbf{S}_{11} & \mathbf{s}_{12} \\ \mathbf{s}_{12}^{\top} & \mathbf{s}_{22} \end{pmatrix}.$$
(15)

Then, as shown by Friedman et al. [46], $w_{12} = W_{11}\beta$, where β is the solution to the following lasso regression problem, denoted as $lasso(W_{11}, s_{12}, \rho)$:

$$\min_{\boldsymbol{\beta}} \frac{1}{2} \| \mathbf{W}_{11}^{1/2} \boldsymbol{\beta} - \boldsymbol{y} \|^2 + \rho \| \boldsymbol{\beta} \|_1,$$
(16)

where $y = \mathbf{W}_{11}^{-1/2} s_{12}$, and ρ is the regularization constant that promotes sparsity in $\boldsymbol{\Sigma}_Z^{-1}$.

In an interventional setting I_k , the dependence of $\hat{S}(\mathbf{B})$ on Σ_Z arrises from the marginal distribution of Z restricted to components indexed by \mathcal{U}_k , i.e., purely observed nodes. Since Z follows a Gaussian distribution, $Z_{\mathcal{U}_k}$ also follows a Gaussian distribution with $Z_{\mathcal{U}_k} \sim \mathcal{N}(\mathbf{0}, \tilde{\Sigma}_{I_k})$. From the properties of Gaussian distribution [47], we have $\tilde{\Sigma}_{I_k} = (\Sigma_Z)_{\mathcal{U}_k, \mathcal{U}_k}$. Consequently, for the interventional setting I_k , (14) becomes

$$\mathcal{L}(I_k) = \sup_{\Sigma_Z} -\text{Tr}\left(\mathbf{S}_{I_k}(\Sigma_Z)^{-1}_{\mathcal{U}_k,\mathcal{U}_k}\right) - \log\left|(\Sigma_Z)_{\mathcal{U}_k,\mathcal{U}_k}\right|,\tag{17}$$

where $\mathbf{S}_{I_k} = \frac{1}{n} \sum_{i=1}^{N} \mathbf{z}_{\mathcal{U}_k}^{(i)} (\mathbf{z}_{\mathcal{U}_k}^{(i)})^{\top}$ is the sample covariance of \mathbf{Z} corresponding to the purely observed nodes. In this case, we set $\mathbf{W} = (\mathbf{\Sigma}_Z)_{\mathcal{U}_k, \mathcal{U}_k}$ and the rest of the update procedure remains the same. The overall parameter update procedure is summarized in Algorithm 1 in Appendix B.

4 **Experiments**

The code is included in the supplementary materials and will be made public upon publication.

We evaluated DCCD-CONF on both synthetic and real-world datasets, comparing its performance against several state-of-the-art baselines: NODAGS-Flow [27], LLC [26], and DCDI [15]. NODAGS-Flow learns nonlinear cyclic causal graphs but does not model unmeasured confounders. LLC accounts for confounders but is limited to linear cyclic SEMs. DCDI supports nonlinear SEMs but does not handle confounders or cycles. We also include a comparison between DCCD-CONF and JCI-FCI [29], a constraint-based method that models confounders, interventions, and nonlinearity, but outputs a partial ancestral graph (PAG), in the Appendix (see Appendix C).



Figure 2: Performance of causal graph and confounder recovery under varying problem dimensions. (Top row) Number of observed variables is fixed at d = 10, while the number of latent confounders ranges from 2 to 8. (Bottom row) The number of confounders is fixed to 0.4d, while d is varied from 10 to 80. Left column shows results for linear SEMs, and right column for nonlinear SEMs.

4.1 Synthetic data

In all synthetic experiments, the cyclic graphs were generated using Erdős-Rényi (ER) random graph model with the outgoing edge density set to 2. We evaluated DCCD-CONF and the baselines on both linear as well as nonlinear SEMs described in Section 2. Our training data set consists of observational data and single-node interventional over all the nodes in the graph, i.e, $\mathcal{I} = \{\emptyset, \{1\}, \dots, \{d\}\}$, with $N_k = 1000$ samples per intervention. Furthermore, in all the experiments presented here, the SEM was constrained to be contractive. However, we also compare the performance of DCCD-CONF to the baselines on non-contractive SEMs in the appendix. For causal graph recovery (directed edges), we use the *normalized structural Hamming distance* (SHD) as the error metric. SHD counts the number of operations (addition, deletion, and reversal) needed to match the estimated causal graph to the ground truth, and normalization is done with respect to the number of nodes in the graph (lower the better). For confounder identification (bidirectional edges), we compare the non-diagonal entries of the estimated confounder-noise covariance matrix to those of the ground truth. We use F1 score as the error metric (higher the better). More details regarding the experimental setup is provided in Appendix B.

Impact of confounder count. We evaluate the performance of DCCD-CONF and the baselines using the previously defined error metrics, varying the confounder ratio (number of confounders divided by the number of nodes) from 0.2 to 0.8. In this case, the number of nodes in the graph is set to d = 10. The results, summarized in Figures 2a and 2b, show that DCCD-CONF consistently achieves lower SHD across all confounder ratios in both linear and nonlinear SEMs. Notably, in nonlinear SEMs, DCCD-CONF outperforms all baselines in causal graph recovery. Additionally, it demonstrates competitive results in confounder identification, highlighting its robustness in both tasks.

Scaling with number of nodes. We compare the performance of DCCD-CONF and the baselines as the number of nodes (d) varies from 10 to 80, with results summarized in Figures 2c and 2d. The number of confounders is set to 0.4d. As the number of nodes increases, SHD rises across all methods, reflecting the increased difficulty of causal graph recovery in larger graphs. However, DCCD-CONF consistently outperforms the baselines in many cases, achieving lower SHD and higher F1 score, suggesting superior scalability with increasing graph size.

Scaling with interventions. In this section, we evaluate graph recovery performance as the number of training interventions K varies from 0 to d, with d = 10 fixed. The case K = 0 corresponds to the observational dataset. Results for the nonlinear SEM setting are presented in Figure 3. As illustrated, DCCD-CONF consistently outperforms LLC across all values of K, and achieves near-perfect graph recovery when all the nodes in the graph are intervened. It is also worth noting that LLC cannot operate in the purely observational setting (K = 0).



Figure 3: Performance comparison between DCCD-CONF and LLC for d = 10 and a confounder ratio of 0.4, as the number of training interventions varies from 0 to 10.

4.2 Real World data

We evaluate DCCD-CONF on learning the causal graph structure of a gene regulatory network from real-world gene expression data with genetic interventions. Specifically, we use the Perturb-CITE-seq dataset [48], which contains gene expression data from 218,331 melanoma cells across three conditions: (i) control, (ii) co-culture, and (iii) IFN- γ . Due to computational constraints, we restrict our analysis to a subset of 61 genes from the 20,000 genes in the genome, following the experimental setup of Sethuraman et al. [27] (see Appendix B for details). Each cell condition is treated as a separate dataset consisting of single-node interventions on the selected 61 genes.

Since the dataset does not provide a ground truth causal graph, SHD cannot be used for direct performance comparison. Instead, we assess DCCD-CONF and the baselines based on predictive performance over unseen interventions. To evaluate performance, we split each dataset 90-10, using the smaller por-

Since the dataset does not provide a ground truth causal graph, SHD cannot be used for direct performance com-

Method	Control	Co-Culture	IFN- γ	
DCCD-CONF	1.375 (0.103)	1.245 (0.039)	1.235 (0.338)	
NODAGS	1.465 (0.015)	1.406 (0.012)	1.504 (0.009)	
LLC	1.385 (0.039)	1.325 (0.029)	1.430 (0.048)	
DCDI	1.523 (0.036)	1.367 (0.018)	1.517 (0.041)	

tion as the test set, and measure performance using negative log-likelihood (NLL) on the test data after model training (lower the better). The results are presented in Table 1. From Table 1, we can see that DCCD-CONF outperforms all the baselines across all the three cell conditions, showcasing the efficacy of the model and prevalence of confounders in real-world systems.

Additional experiments. Additionally, we also provide results in Appendix C for the following settings: (i) performance comparison on non-contractive SEMs when the underlying graph is restricted to DAGs, (ii) performance comparison as a function of training data size, (iii) performance comparison as a function of training data size, (iii) performance comparison as a function of training data size, (iii) performance comparison as a function of training data size, (iii) performance comparison as a function of training data size, (iii) performance comparison as a function of training data size, (iii) performance comparison as a function of training data size, (iii) performance comparison as a function of training data size, (iii) performance comparison as a function of training data size, (iii) performance comparison as a function of training data size, (iii) performance comparison as a function of training data size, (iii) performance comparison as a function of training data size, (iii) performance comparison as a function of training data size, (iii) performance comparison as a function of training data size, (iii) performance comparison as a function of training data size, (iii) performance comparison as a function outgoing edge density, and (v) performance comparison between DCCD-CONF and JCI-FCI.

5 Discussion

In this work, we introduced DCCD-CONF, a novel differentiable causal discovery framework that handles directed cycles and unmeasured confounders, assuming Gaussian exogenous noise. It models causal mechanisms via neural networks and learns the causal graph structure by maximizing penalized data likelihood. We provide consistency guarantees in the large-sample regime and demonstrate, through extensive synthetic and real-world experiments, that DCCD-CONF outperforms state-of-the-art methods, maintaining robustness with increasing confounders and graph size. On the Perturb-CITE-seq dataset, our model achieves superior predictive accuracy.

While the focus of this work is limited to Gaussian exogenous noise, we plan to investigate other noise distributions for future research. Other future directions include supporting missing data, and relaxing interventional assumptions by incorporating soft interventions and unknown interventional targets.

References

- Karen Sachs, Omar Perez, Dana Pe'er, Douglas A. Lauffenburger, and Garry P. Nolan. Causal protein-signaling networks derived from multiparameter single-cell data. *Science*, 308(5721): 523–529, 2005.
- [2] Eran Segal, Dana Pe'er, Aviv Regev, Daphne Koller, Nir Friedman, and Tommi Jaakkola. Learning module networks. *Journal of Machine Learning Research*, 6(4), 2005.
- [3] Bin Zhang, Chris Gaiteri, Liviu-Gabriel Bodea, Zhi Wang, Joshua McElwee, Alexei A. Podtelezhnikov, Chunsheng Zhang, Tao Xie, Linh Tran, and Radu Dobrin. Integrated systems approach identifies genetic nodes and networks in late-onset Alzheimer's disease. *Cell*, 153(3):707–720, 2013.
- [4] Peter Spirtes, Clark N Glymour, Richard Scheines, and David Heckerman. *Causation, prediction, and search.* MIT press, 2000.
- [5] Sofia Triantafillou and Ioannis Tsamardinos. Constraint-based causal discovery from multiple interventions over overlapping variable sets. *The Journal of Machine Learning Research*, 16(1): 2147–2205, 2015.
- [6] Christina Heinze-Deml, Jonas Peters, and Nicolai Meinshausen. Invariant causal prediction for nonlinear models. *Journal of Causal Inference*, 6(2), 2018.
- [7] Christopher Meek. *Graphical Models: Selecting causal and statistical models*. PhD thesis, Carnegie Mellon University, 1997.
- [8] Alain Hauser and Peter Bühlmann. Characterization and greedy learning of interventional markov equivalence classes of directed acyclic graphs. *The Journal of Machine Learning Research*, 13(1):2409–2464, 2012.
- [9] Xun Zheng, Bryon Aragam, Pradeep K Ravikumar, and Eric P Xing. DAGs with NO TEARS: Continuous optimization for structure learning. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, Advances in Neural Information Processing Systems, volume 31, 2018. URL https://proceedings.neurips.cc/paper/2018/ file/e347c51419ffb23ca3fd5050202f9c3d-Paper.pdf.
- [10] Yue Yu, Jie Chen, Tian Gao, and Mo Yu. DAG-GNN: DAG structure learning with graph neural networks. In *International Conference on Machine Learning*, pages 7154–7163. PMLR, 2019.
- [11] Ignavier Ng, AmirEmad Ghassami, and Kun Zhang. On the role of sparsity and DAG constraints for learning linear dags. Advances in Neural Information Processing Systems, 33:17943–17954, 2020.
- [12] Ignavier Ng, Shengyu Zhu, Zhuangyan Fang, Haoyang Li, Zhitang Chen, and Jun Wang. Masked gradient-based causal structure learning. In *Proceedings of the 2022 SIAM International Conference on Data Mining (SDM)*, pages 424–432. SIAM, 2022.
- [13] Xun Zheng, Chen Dan, Bryon Aragam, Pradeep Ravikumar, and Eric Xing. Learning sparse nonparametric DAGs. In Silvia Chiappa and Roberto Calandra, editors, *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics*, volume 108, pages 3414–3425, 26–28 Aug 2020.
- [14] Hao-Chih Lee, Matteo Danieletto, Riccardo Miotto, Sarah T Cherng, and Joel T Dudley. Scaling structural learning with NO-BEARS to infer causal transcriptome networks. In *Pacific Symposium on Biocomputing 2020*, pages 391–402. World Scientific, 2019.
- [15] Philippe Brouillard, Sébastien Lachapelle, Alexandre Lacoste, Simon Lacoste-Julien, and Alexandre Drouin. Differentiable causal discovery from interventional data. Advances in Neural Information Processing Systems, 33:21865–21877, 2020.
- [16] Ioannis Tsamardinos, Laura E Brown, and Constantin F Aliferis. The max-min hill-climbing bayesian network structure learning algorithm. *Machine learning*, 65(1):31–78, 2006.
- [17] Liam Solus, Yuhao Wang, Lenka Matejovicova, and Caroline Uhler. Consistency guarantees for permutation-based causal inference algorithms. *arXiv preprint arXiv:1702.03530*, 2017.
- [18] Yuhao Wang, Liam Solus, Karren Yang, and Caroline Uhler. Permutation-based causal inference algorithms with interventions. *Advances in Neural Information Processing Systems*, 30, 2017.

- [19] Jacob W. Freimer, Oren Shaked, Sahin Naqvi, Nasa Sinnott-Armstrong, Arwa Kathiria, Christian M. Garrido, Amy F. Chen, Jessica T. Cortez, William J. Greenleaf, Jonathan K. Pritchard, and Alexander Marson. Systematic discovery and perturbation of regulatory genes in human T cells reveals the architecture of immune networks. *Nature Genetics*, pages 1–12, July 2022. ISSN 1546-1718. doi: 10.1038/s41588-022-01106-y. URL https://www.nature.com/articles/s41588-022-01106-y.
- [20] Thomas Richardson. A discovery algorithm for directed cyclic graphs. In *Proceedings of the Twelfth international conference on Uncertainty in artificial intelligence*, pages 454–461, 1996.
- [21] Gustavo Lacerda, Peter Spirtes, Joseph Ramsey, and Patrik O. Hoyer. Discovering cyclic causal models by independent components analysis. In *Proceedings of the Twenty-Fourth Conference* on Uncertainty in Artificial Intelligence, UAI'08, page 366–374, Arlington, Virginia, USA, 2008. AUAI Press. ISBN 0974903949.
- [22] Jan-Christian Huetter and Philippe Rigollet. Estimation rates for sparse linear cyclic causal models. In Jonas Peters and David Sontag, editors, *Proceedings of the 36th Conference on Uncertainty in Artificial Intelligence (UAI)*, volume 124 of *Proceedings of Machine Learning Research*, pages 1169–1178. PMLR, 03–06 Aug 2020. URL https://proceedings.mlr. press/v124/huetter20a.html.
- [23] Carlos Améndola, Philipp Dettling, Mathias Drton, Federica Onori, and Jun Wu. Structure learning for cyclic linear causal models. In *Conference on Uncertainty in Artificial Intelligence*, pages 999–1008. PMLR, 2020.
- [24] Joris M Mooij and Tom Heskes. Cyclic causal discovery from continuous equilibrium data. In *Uncertainty in Artificial Intelligence*, 2013.
- [25] Mathias Drton, Christopher Fox, and Y. Samuel Wang. Computation of maximum likelihood estimates in cyclic structural equation models. *The Annals of Statistics*, 47(2):663 – 690, 2019. doi: 10.1214/17-AOS1602. URL https://doi.org/10.1214/17-AOS1602.
- [26] Antti Hyttinen, Frederick Eberhardt, and Patrik O Hoyer. Learning linear cyclic causal models with latent variables. *The Journal of Machine Learning Research*, 13(1):3387–3439, 2012.
- [27] Muralikrishnna G Sethuraman, Romain Lopez, Rahul Mohan, Faramarz Fekri, Tommaso Biancalani, and Jan-Christian Hütter. Nodags-flow: Nonlinear cyclic causal structure learning. In *International Conference on Artificial Intelligence and Statistics*, pages 6371–6387. PMLR, 2023.
- [28] Peter Spirtes. An anytime algorithm for causal inference. In Thomas S. Richardson and Tommi S. Jaakkola, editors, *Proceedings of the Eighth International Workshop on Artificial Intelligence and Statistics*, volume R3 of *Proceedings of Machine Learning Research*, pages 278–285. PMLR, 04–07 Jan 2001. URL https://proceedings.mlr.press/r3/spirtes01a.html. Reissued by PMLR on 31 March 2021.
- [29] Joris M Mooij, Sara Magliacane, and Tom Claassen. Joint causal inference from multiple contexts. *Journal of machine learning research*, 21(99):1–108, 2020.
- [30] Murat Kocaoglu, Amin Jaber, Karthikeyan Shanmugam, and Elias Bareinboim. Characterization and learning of causal graphs with latent variables from soft interventions. *Advances in Neural Information Processing Systems*, 32, 2019.
- [31] Amin Jaber, Murat Kocaoglu, Karthikeyan Shanmugam, and Elias Bareinboim. Causal discovery from soft interventions with unknown targets: Characterization and learning. *Advances in neural information processing systems*, 33:9551–9561, 2020.
- [32] Patrick Forré and Joris M Mooij. Constraint-based causal discovery for non-linear structural causal models with cycles and latent confounders. *arXiv preprint arXiv:1807.03024*, 2018.
- [33] Rohit Bhattacharya, Tushar Nagarajan, Daniel Malinsky, and Ilya Shpitser. Differentiable causal discovery under unmeasured confounding. In *International Conference on Artificial Intelligence and Statistics*, pages 2314–2322. PMLR, 2021.
- [34] Kenneth A Bollen. *Structural equations with latent variables*, volume 210. John Wiley & Sons, 1989.
- [35] Judea Pearl. *Causality*. Cambridge University Press, 2 edition, 2009. doi: 10.1017/ CBO9780511803161.

- [36] Stephan Bongers, Patrick Forré, Jonas Peters, and Joris M Mooij. Foundations of structural causal models with cycles and latent variables. *The Annals of Statistics*, 49(5):2885–2915, 2021.
- [37] Jonas Peters and Peter Bühlmann. Identifiability of gaussian structural equation models with equal error variances. *Biometrika*, 101(1):219–228, 2014.
- [38] Cheng Lu, Jianfei Chen, Chongxuan Li, Qiuhao Wang, and Jun Zhu. Implicit normalizing flows. In International Conference on Learning Representations, 2021. URL https://openreview. net/forum?id=8PS8m9oYtNy.
- [39] Charles G Broyden. A class of methods for solving nonlinear simultaneous equations. *Mathe-matics of computation*, 19(92):577–593, 1965.
- [40] Jens Behrmann, Will Grathwohl, Ricky TQ Chen, David Duvenaud, and Jörn-Henrik Jacobsen. Invertible residual networks. In *International Conference on Machine Learning*, pages 573–582. PMLR, 2019.
- [41] Romain Lopez, Jan-Christian Hütter, Jonathan K Pritchard, and Aviv Regev. Large-scale differentiable causal discovery of factor graphs. In *Advances in Neural Information Processing Systems*, 2022.
- [42] Brian C. Hall. Lie Groups, Lie Algebras, and Representations, pages 333–366. Springer New York, New York, NY, 2013. ISBN 978-1-4614-7116-5. doi: 10.1007/978-1-4614-7116-5_16. URL https://doi.org/10.1007/978-1-4614-7116-5_16.
- [43] Ricky TQ Chen, Jens Behrmann, David K Duvenaud, and Jörn-Henrik Jacobsen. Residual flows for invertible generative modeling. *Advances in Neural Information Processing Systems*, 32, 2019.
- [44] Michael F Hutchinson. A stochastic estimator of the trace of the influence matrix for Laplacian smoothing splines. *Communications in Statistics-Simulation and Computation*, 18(3):1059– 1076, 1989.
- [45] Eric Jang, Shixiang Gu, and Ben Poole. Categorical reparameterization with gumbel-softmax. In International Conference on Learning Representations, 2017. URL https://openreview. net/forum?id=rkE3y85ee.
- [46] Jerome Friedman, Trevor Hastie, and Robert Tibshirani. Sparse inverse covariance estimation with the graphical lasso. *Biostatistics*, 9(3):432–441, 2008.
- [47] Michael I Jordan. An introduction to probabilistic graphical models, 2003.
- [48] Chris J Frangieh, Johannes C Melms, Pratiksha I Thakore, Kathryn R Geiger-Schuller, Patricia Ho, Adrienne M Luoma, Brian Cleary, Livnat Jerby-Arnon, Shruti Malu, Michael S Cuoco, et al. Multimodal pooled Perturb-CITE-seq screens in patient models define mechanisms of cancer immune evasion. *Nature genetics*, 53(3):332–341, 2021.
- [49] Patrick Forré and Joris M Mooij. Markov properties for graphical models with cycles and latent variables. *arXiv preprint arXiv:1710.08775*, 2017.
- [50] Peter L Spirtes. Directed cyclic graphical representations of feedback models. *arXiv preprint arXiv:1302.4982*, 2013.
- [51] Thomas S. Richardson. A factorization criterion for acyclic directed mixed graphs. In Proceedings of the Twenty-Fifth Conference on Uncertainty in Artificial Intelligence, UAI '09, page 462–470, Arlington, Virginia, USA, 2009. AUAI Press. ISBN 9780974903958.
- [52] Steffen L Lauritzen and Finn Verner Jensen. Local computation with valuations from a commutative semigroup. Annals of Mathematics and Artificial Intelligence, 21:51–69, 1997.
- [53] Peter Spirtes and Thomas S. Richardson. A polynomial time algorithm for determining dag equivalence in the presence of latent variables and selection bias. In David Madigan and Padhraic Smyth, editors, *Proceedings of the Sixth International Workshop on Artificial Intelligence and Statistics*, volume R1 of *Proceedings of Machine Learning Research*, pages 489–500. PMLR, 04– 07 Jan 1997. URL https://proceedings.mlr.press/r1/spirtes97b.html. Reissued by PMLR on 30 March 2021.
- [54] R. Ayesha Ali, Thomas S. Richardson, and Peter Spirtes. Markov equivalence for ancestral graphs. *The Annals of Statistics*, 37(5B), October 2009. ISSN 0090-5364. doi: 10.1214/ 08-aos626. URL http://dx.doi.org/10.1214/08-A0S626.

- [55] Jiji Zhang. Causal reasoning with ancestral graphs. *Journal of Machine Learning Research*, 9 (7), 2008.
- [56] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In Yoshua Bengio and Yann LeCun, editors, 3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings, 2015. URL http://arxiv.org/abs/1412.6980.

Appendices

The appendix is structured as follows: Appendix A presents the theoretical foundations of differentiable cyclic causal discovery in the presence of unmeasured confounders, including the proof of Theorem 2, and a characterization of the equivalence class of DMGs that maximize the score function. Appendix B provides implementation details of DCCD-CONF and the baselines. Finally, Appendix C provides additional experimental results comparing DCCD-CONF with the baselines.

A Theory

In this section, we establish the theoretical foundations of differentiable cyclic causal discovery in the presence of unmeasured confounders. We begin by reviewing key definitions and results from prior work that are essential for proving Theorem 2, starting with fundamental graph terminology.

A.1 Preliminaries

Consider a directed mixed graph $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathcal{B})$. A path π between nodes *i* and *j* is a sequence $(i_0, \varepsilon_1, i_1, \dots, \varepsilon_n, i_n)$, where $\{i_0, \dots, i_n\} \subseteq \mathcal{V}$ and $\{\varepsilon_1, \dots, \varepsilon_n\} \subseteq \mathcal{E} \cup \mathcal{B}$, with $i_0 = i$ and $i_n = j$. A path is *directed* if each edge ε_k follows the form $i_{k-1} \to i_k$ for all $k \in [n]$. A cycle through node i consists of a directed path from i to some node j and an additional edge $j \rightarrow i$. For any node $i \in \mathcal{V}$, the ancestor set is defined as $\operatorname{an}_{\mathcal{G}}(i) := \{j \in \mathcal{V} \mid a \text{ directed path from } j \text{ to } i \text{ exists in } \mathcal{G}\}$, while the descendant set is given by $de_{\mathcal{G}}(i) := \{j \in \mathcal{V} \mid a \text{ directed path from } i \text{ to } j \text{ exists in } \mathcal{G}\}$. The spouse set of a node i is defined as $\operatorname{sp}_{\mathcal{G}}(i) := \{j \in \mathcal{V} \mid j \leftrightarrow i \in \mathcal{B}\}$. If i is both a spouse and an ancestor of j, this creates a *almost directed cycle*. A mixed graph is called *ancestral* if it does not contain a directed or an almost directed cycle. The strongly connected component of i, denoted $sc_{\mathcal{G}}(i)$, is the intersection of its ancestors and descendants: $sc_{\mathcal{G}}(i) = an_{\mathcal{G}}(i) \cap de_{\mathcal{G}}(i)$. The district of a node $i \in \mathcal{V}$ is defined as $\operatorname{dis}_{\mathcal{G}}(i) = \{j \mid j \leftrightarrow \cdots \leftrightarrow i \in \mathcal{G} \text{ or } i = j\}$. We can apply these definitions to subsets $\mathcal{U} \subseteq \mathcal{V}$ by taking union of over the items of the subset, for instance, $\operatorname{an}_{\mathcal{G}}(\mathcal{U}) = \bigcup_{i \in \mathcal{U}} \operatorname{an}_{\mathcal{G}}(i)$. A vertex set $\mathcal{A} \subseteq \mathcal{V}$ is said to be *barren* if $i \in \mathcal{A}$ has no descendants in \mathcal{G} that are in \mathcal{A} , however, i may have descendants in \mathcal{G} not in \mathcal{A} , that is, $\operatorname{barren}_{\mathcal{G}}(\mathcal{A}) = \{i \mid i \in \mathcal{A}; \operatorname{de}_{\mathcal{G}}(i) \cap \mathcal{A} = \{i\}\}$. A subset $\mathcal{A} \subseteq \mathcal{V}$ is ancestrally closed if \mathcal{A} contains all of its ancesters. We define $\mathcal{A}(\mathcal{G}) := \{\mathcal{A} \mid an_{\mathcal{G}}(\mathcal{A}) = \mathcal{A}\}$ as the set of ancestrally closed set of ancestrally closed sets in \mathcal{G} .

Definition A.1 (Collider). For a directed mixed graph $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathcal{B})$, a node $i_k \in \mathcal{V}$ in a path $\pi = (i_0, \varepsilon_1, i_1, \varepsilon_2, \dots, i_{n-1}, \varepsilon_n, i_n)$ is called a collider if $k \neq 0, n$ (non-endpoint) and the two edges $\varepsilon_k, \varepsilon_{k+1}$ have their heads pointed at i, i.e., the subpath $(i_{k-1}, \varepsilon_k, i_k, \varepsilon_{k+1}, i_{k+1})$ is of the form $i_{k-1} \rightarrow i_k \leftarrow i_{k+1}, i_{k-1} \leftrightarrow i_k \leftarrow i_{k+1}, i_{k-1} \rightarrow i_k \leftrightarrow i_{k+1}$. The node i_k is called a non-collider if i_k is not a collider.

Note that the end points of a walk are always non-colliders. We now define the notion of *d*-separation extended to DMGs.

Definition A.2 (*d*-separation). Let $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathcal{B})$ be a directed mixed graph and let $C \subseteq \mathcal{V}$ be a subset of nodes. A path $\pi = (i_0, \varepsilon_1, i_1, \varepsilon_2, \dots, i_{n-1}, \varepsilon_n, i_n)$ is said to be *d*-blocked given \mathcal{C} if

- 1. π contains a collider $i_k \notin \operatorname{an}_{\mathcal{G}}(C)$
- 2. π contains a non-collider $i_k \in C$.

The path π is said to be d-open given C if it is not d-blocked. Two subsets of nodes $A, B \subseteq V$ is said to be d-separated given C if all paths between a and b, where $a \in A$ and $b \in B$, is d-blocked given C, and is denoted by

$$A \stackrel{d}{\perp} B \mid C.$$

If the underlying graph is acyclic, d-separation implies conditional independence. That is, for subsets of nodes $A, B, C \subseteq \mathcal{V}$,

$$A \stackrel{d}{\perp}_{\mathcal{G}} B \mid C \implies \mathbf{X}_A \stackrel{\perp}{}_{p_{\mathcal{G}}} \mathbf{X}_B \mid \mathbf{X}_C,$$

where $\perp_{p_{\mathcal{G}}}$ denotes conditional independence, and $p_{\mathcal{G}}$ denotes the observational distribution. This is known as the *directed global Markov property* of \mathcal{G} [49]. However, in general, cyclic graphs do



Figure 4: (Left) Illustration of a directed mixed graph that disobeys directed global Markov property. (Right) The graph on the right represents the graph \mathcal{G} after the acyclification process.

not obey the directed global Markov property as shown by the counterexample below taken from [36, 50].

Example A.3. Consider the SEM given by:

 $f_1(X,Z) = Z_1, \quad f_1(X,Z) = Z_2, \quad f_3(X,Z) = X_1X_4 + Z_3 \quad f_4(X,Z) = X_2X_3 + Z_4,$

and p_Z is the standard normal distribution. One can check that X_1 is not independent of X_2 given $\{X_3, X_4\}$. However, the X_1 and X_2 are d-separated given $\{X_3, X_4\}$ in the graph corresponding to the SEM (see Figure 4).

Forré and Mooij [49] introduced σ -separation as a generalization of d-separation to extend the directed global Markov property to cyclic graphs. This concept was motivated by applying *d*-separation to the acyclified version of the DMG. Before delving into σ -separation, we first define the acyclification procedure of a directed mixed graph, following [36].

Definition A.4 (Acyclification of a directed mixed graph). Let $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathcal{B})$ denote a directed mixed graph, the acyclification of \mathcal{G} maps \mathcal{G} to the acyclified graph $\operatorname{acy}(\mathcal{G}) = (\mathcal{V}, \hat{\mathcal{E}}, \hat{\mathcal{B}})$, where $j \to i \in \hat{\mathcal{E}}$ if and only if $j \in \operatorname{pa}_{\mathcal{G}}(\operatorname{sc}_{\mathcal{G}}(i)) \setminus \operatorname{sc}_{\mathcal{G}}(i)$, and $i \leftrightarrow j \in \hat{\mathcal{B}}$ if and only if there exists $i' \in \operatorname{sc}_{\mathcal{G}}(i)$ and $j' \in \operatorname{sc}_{\mathcal{G}}(j)$ such that i' = j' or $i' \leftrightarrow j' \in \mathcal{B}$.

It is important to note that the existence of a acylified graph for an SEM relies on the solvability of the SEM over all the strongly connected components of the DMG corresponding to the SEM. This is to say that we have a solution for $X_{sc_{\mathcal{G}}(i)}$ given $X_{pa(sc_{\mathcal{G}}(i))\setminus sc_{\mathcal{G}}(i)}$ and $Z_{sc_{\mathcal{G}}(i)}$. This is indeed the case as we assume that the forward map f_x is invertible for the all the SEMs under consideration, see Bongers et al. [36] for more details. Figure 4 illustrates the acyclification process for the graph corresponding to Example A.3.

Definition A.5 (σ -separation). Let $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathcal{B})$ be a directed mixed graph and let $C \subseteq \mathcal{V}$ be a subset of nodes. A path $\pi = (i_0, \varepsilon_1, i_1, \varepsilon_2, \dots, i_{n-1}, \varepsilon_n, i_n)$ is said to be σ -blocked given \mathcal{C} if

- 1. the first node of π , $i_0 \in C$ or its last node $i_n \in C$, or
- 2. π contains a collider $i_k \notin \operatorname{an}_{\mathcal{G}}(C)$
- 3. π contains a non-collider $i_k \in C$ that points towards a neighbor that is not in the same strongly connected component as i_k in \mathcal{G} , i.e, such that $i_{k-1} \leftarrow i_k$ in π and $i_{k-1} \notin \operatorname{sc}_{\mathcal{G}}(i_k)$, or $i_k \to i_{k+1}$ in π and $i_{k+1} \notin \operatorname{sc}_{\mathcal{G}}(i_k)$.

The path π is said to be σ -open given C if it is not σ -blocked. Two subsets of nodes $A, B \subseteq V$ is said to be σ -separated given C if all paths between a and b, where $a \in A$ and $b \in B$, is σ -blocked given C, and is denoted by

$$A \stackrel{\sigma}{\underset{\mathcal{G}}{\perp}} B \mid C.$$

Note that σ -separation reduces to *d*-separation for acyclic graphs, that is, when $sc_{\mathcal{G}}(i) = \{i\}$ for all $i \in \mathcal{V}$. The following result in [49] relates σ -separation and *d*-separation.

Proposition A.6 ([49]). Let $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathcal{B})$ be a directed mixed graph, then for $A, B, C \subseteq \mathcal{V}$,

$$A \stackrel{\sigma}{\underset{\mathcal{G}}{\perp}} B \mid C \iff A \stackrel{d}{\underset{\operatorname{acy}(\mathcal{G})}{\perp}} B \mid C.$$



Figure 5: Illustration of the augmented graph $\mathcal{G}^{\mathcal{I}}$ corresponding to the set of interventional targets $\mathcal{I} = \{\emptyset, \{X_3\}, \{X_4\}\}$. do($\{X_3\}$) and do($\{X_3\}$) corresponds to the graph obtained after hard interventions on X_3 and X_4 respectively. The augmented graph here is the union of the graphs \mathcal{G} , do($\{X_3\}$), do($\{X_4\}$) along with the context variables.

Using σ -separation we can now define the general directed global Markov property.

Definition A.7 (General directed global Markov property [49]). Let $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathcal{B})$ be a directed mixed graph and $p_{\mathcal{G}}$ denote the probability density of the observations \mathbf{X} . The probability density $p_{\mathcal{G}}$ satisfies the general directed global Markov property if for $A, B, C \subseteq \mathcal{V}$

$$A \stackrel{o}{\perp}_{\mathcal{G}} B \mid C \implies \mathbf{X}_A \underset{p_{\mathcal{G}}}{\perp} \mathbf{X}_B \mid \mathbf{X}_C,$$

that is, X_A and X_B are conditionally independent given X_C .

A.2 Joint Causal Modelling and Markov properties

In order to incorporate multiple interventional settings into a single causal modeling framework, we follow the *joint causal model* introduced by [29], where we augment the system with a set of context variables $C^{\mathcal{I}} = (C_1, \ldots, C_K)$ each corresponding to a non-empty interventional setting. In this case, $C_k = \emptyset$ for all $k = 1, \ldots, K$ corresponds to the observational setting. We construct an augmented graph, denoted by $\mathcal{G}^{\mathcal{I}}$ consisting of both the system variables X and the context variables $C^{\mathcal{I}}$, such that the $ch_{\mathcal{G}}(C_k) = I_k$, and no context variable has any parent or a spouse. Figure 5 illustrates the augmented graph for the graph from Example A.3 and the intervention sets $\mathcal{I} = \{\emptyset, \{X_3\}, \{X_4\}\}$. The new system containing both the observed variables and the context variables is called the *meta system*. Finally, given a set of interventional targets $I_k \subseteq \mathcal{V}$ and the corresponding context variable C_k , the structural equations governing the observations X has the following form:

$$\tilde{F}_i(\boldsymbol{X}_{\mathrm{pa}_{\mathcal{G}}(i)}, \boldsymbol{C}_{\mathrm{pa}_{\mathcal{G}}^{\mathcal{I}}(i)}^{\mathcal{I}}, Z_i) = \begin{cases} (\boldsymbol{C}_k)_i, & \text{if } X_i \in I_k, \\ F_i(\boldsymbol{X}_{\mathrm{pa}_{\mathcal{G}}(i)}, Z_i), & \text{if } X_i \notin I_k. \end{cases}$$

We call the distribution over the context variables $p(\mathbf{C}^{\mathcal{I}})$ the *context distributions* and as noted by Mooij et al. [29], the behavior of the system is usually invariant to the context distribution. We assume access to the context distribution as the interventional settings are known apriori. Note that, the observational distribution corresponds to $p_{\mathcal{G}^{\mathcal{I}}}(\mathbf{X} \mid \mathbf{C}_1 = \cdots = \mathbf{C}_K = \emptyset)$. Similarly, the interventional distribution for the interventional setting I_k corresponds to $p_{\mathcal{G}^{\mathcal{I}}}(\mathbf{X} \mid \mathbf{C}_k = \boldsymbol{\xi}_{I_k}, \mathbf{C}_{-k} = \emptyset)$, i.e.,

$$p_{\mathcal{G}^{\mathcal{I}}}(\boldsymbol{X} \mid \boldsymbol{C}_{k} = \boldsymbol{\xi}_{I_{k}}, \boldsymbol{C}_{-k} = \emptyset) = p_{\operatorname{do}(I_{k})(\mathcal{G})}(\boldsymbol{X})$$

Furthermore,

$$p_{\mathcal{G}^{\mathcal{I}}}(\boldsymbol{C}^{\mathcal{I}}, \boldsymbol{X}) = p_{\mathcal{G}^{\mathcal{I}}}(\boldsymbol{C}^{\mathcal{I}}) p_{\mathcal{G}^{\mathcal{I}}}(\boldsymbol{X} \mid \boldsymbol{C}^{\mathcal{I}}).$$
(18)

Recall that for the interventional setting I_k , the probability density function governing the observations X is given by (6), which we repeat here for convenience

$$p_{\mathrm{do}(I_k)(\mathcal{G})}(\boldsymbol{X}) = p_I(\boldsymbol{C})p_Z\Big(\Big[\mathbf{f}_x^{(I_k)}(\boldsymbol{X})\Big]_{\mathcal{U}_k}\Big)\Big|\det\Big(\mathbf{J}_{\mathbf{f}_x^{(I_k)}}(\boldsymbol{X})\Big)\Big|$$

Definition A.8. Let $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathcal{B})$ be a directed mixed graph, and $\mathcal{I} = \{I_k\}_{k=0}^K$ with $I_0 = \emptyset$ be a family of interventional targets. Let $\mathcal{M}_{\mathcal{I}}(\mathcal{G})$ denote the set of positive densities $p_{\mathcal{G}^{\mathcal{I}}} : \mathbb{R}^{2d} \to \mathbb{R}$ such that $p_{\mathcal{G}^{\mathcal{I}}}$ is given by (18) for all $\mathbf{F} : \mathbb{R}^{2d} \to \mathbb{R}^d$, with $F_i(\mathbf{X}, \mathbf{Z}) = F_i(\mathbf{X}_{pa_{\mathcal{G}}}(i), Z_i)$, such that the resulting forward map \mathbf{f}_x is unique and invertible, and $\mathbf{\Sigma}_Z \succ 0$ and $(\mathbf{\Sigma}_Z)_{ij} \neq 0$ if and only if $i \leftrightarrow j \in \mathcal{B}$.

Proposition A.9. For a directed mixed graph $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathcal{B})$ and a family of interventional targets $\mathcal{I} = \{I_k\}_{k=0}^K$ such that $I_0 = \emptyset$, let $p \in \mathcal{M}_{\mathcal{I}}(\mathcal{G})$, then p satisfies the general directed global Markov property relative to $\mathcal{G}^{\mathcal{I}}$.

Proof. For an DMG \mathcal{G} and a choice of $\mathbf{F} : \mathbb{R}^{2d} \to \mathbb{R}^d$ such that \mathbf{f}_x is unique and invertible and $\Sigma \succ 0$, the structural equations are uniquely solvable with respect to each strongly connected component of \mathcal{G} . Morover, the addition of context variables in the augmented graph does not introduce any new cycles. Therefore the meta system forms a simple SCM. Thus, from Theorem A.21 in [36], the distribution $p_{\mathcal{G}^{\mathcal{I}}}$ is unique and it satisfies the general directed global Markov property.

We now define the notion of interventional Markov equivalence class for DMGs based on the set of distribution induced by them.

Definition A.10 (\mathcal{I} -Markov Equivalence Class). Two directed mixed graphs \mathcal{G}_1 and \mathcal{G}_2 are \mathcal{I} -Markov equivalent if and only if $\mathcal{M}_{\mathcal{I}}(\mathcal{G}_1) = \mathcal{M}_{\mathcal{I}}(\mathcal{G}_2)$, denoted as $\mathcal{G}_1 \equiv_{\mathcal{I}} \mathcal{G}_2$. The set of all directed mixed graphs that are \mathcal{I} -Markov equivalent to \mathcal{G}_1 is the \mathcal{I} -Markov equivalence class of \mathcal{G}_1 , denoted as \mathcal{I} -MEC(\mathcal{G}_1).

From Proposition A.6, for a DMG \mathcal{G} and a family of interventional targets $\mathcal{I} = \{I_k\}_{k=0}^K$, any σ -separation statement in $\mathcal{G}^{\mathcal{I}}$ translates to a *d*-separation statement in the acyclified graph $\operatorname{acy}(\mathcal{G}^{\mathcal{I}})$. Consequently, the acyclified graph $\operatorname{acy}(\mathcal{G}^{\mathcal{I}})$ is equivalent to the augmented graph $\mathcal{G}^{\mathcal{I}}$. Furthermore, by the results of [51], $p_{\mathcal{G}^{\mathcal{I}}}$ admits a factorization, as formalized in the theorem below.

Theorem A.11 ([51]). A probability distribution p obeys the directed Markov property for an acyclic directed mixed graph G if and only if for every $A \in \mathcal{A}(G)$,

$$p(\boldsymbol{X}_{A}) = \prod_{H \in [A]_{\mathcal{G}}} p(\boldsymbol{X}_{H} \mid \boldsymbol{X}_{tail(H)})$$
(19)

where $[A]_{\mathcal{G}}$ denotes a partition of A into sets $\{H_1, \ldots, H_k\}$.

Each term in the factorization above is of the form $p(X_H | X_T)$, $H, T \subseteq \mathcal{V}$, and $H \cap T = \emptyset$. Following Richardson [51], Lauritzen and Jensen [52] we refer to H as the *head* of the term $p(X_H | X_T)$, and T as the *tail*. An ordered pair of sets (H, T) form the head and tail of the factor associated with \mathcal{G} if and only if all of the following conditions hold:

- 1. $H = \operatorname{barren}_{\mathcal{G}}(\operatorname{an}_{\mathcal{G}}(H)),$
- 2. If every nodes $h \in H$ is connected via a path in the graph obtained by removing all the directed edges in the graph \mathcal{G} when restricted to the nodes $an_{\mathcal{G}}(H)$, and
- 3. $T = (\operatorname{dis}_{\operatorname{an}_{\mathcal{G}}(H)} \setminus H) \cup \operatorname{pa}_{\mathcal{G}}(\operatorname{dis}_{\operatorname{an}_{\mathcal{G}}(H)}).$

Proposition A.12. Let $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathcal{B})$ be a directed mixed graph and $\mathcal{I} = \{I_k\}_{k=0}^K$ be a family of interventional targets. The set of interventional distributions $p_{\mathcal{G}^{\mathcal{I}}} \in \mathcal{M}_{\mathcal{I}}(\mathcal{G})$ if and only if $p_{\mathcal{G}^{\mathcal{I}}}$ admits a factorization of the form given by (19).

Proof. Since any $p_{\mathcal{G}^{\mathcal{I}}} \in \mathcal{M}_{\mathcal{I}}(\mathcal{G})$ is also Markov to $acy(\mathcal{G}^{\mathcal{I}})$, the proposition above is a direct implication of applying of Theorem 19 on $acy(\mathcal{G}^{\mathcal{I}})$.

A.3 Proof of Theorem 2

We now present the main result of this paper. Recall the score function introduced in Section 3.2,

$$\mathcal{S}_{\mathcal{I}}(\mathcal{G}) := \sup_{\boldsymbol{\phi}} \sum_{k=1}^{K} \mathop{\mathbb{E}}_{\boldsymbol{X} \sim p^{(k)}} \log p_{\operatorname{do}(I_k)(\mathcal{G})}(\boldsymbol{X}) - \lambda |\mathcal{G}|,$$

where $p^{(k)}$ is the data-generating distribution for $I_k \in \mathcal{I}$, and $\phi = \{\theta, \Sigma_Z\}$ represents the set of all model parameters. In the context of the meta system, since we assume access to the context distribution, the score function above is equivalent to the following score:

$$\mathcal{S}_{\mathcal{I}}(\mathcal{G}) := \sup_{\phi} \mathop{\mathbb{E}}_{(\boldsymbol{X}, \boldsymbol{C}) \sim p_{\mathcal{I}}^{\star}} \log p_{\mathcal{G}^{\mathcal{I}}}(\boldsymbol{X}, \boldsymbol{C} \mid \boldsymbol{\phi}) - \lambda |\mathcal{G}|,$$

where $p_{\mathcal{G}^{\mathcal{I}}}(\boldsymbol{X}, \boldsymbol{C} \mid \boldsymbol{\phi})$ is given by (18) for a specific choice of $\boldsymbol{\phi}$, and $p_{\mathcal{I}}^*$ denotes the joint groundtruth distribution for the observed and the context variables. We define $\mathcal{P}_{\mathcal{I}}(\mathcal{G})$ as the set of all distributions $p_{\mathcal{G}^{\mathcal{I}}}(\boldsymbol{X}, \boldsymbol{C} \mid \boldsymbol{\phi})$ that can be expressed by the model specified by equations (5) and (8). That is,

$$\mathcal{P}_{\mathcal{I}}(\mathcal{G}) := \{ p \mid \exists \phi \text{ s.t } p = p_{\mathcal{G}^{\mathcal{I}}}(\cdot \mid \phi) \}.$$
(20)

From the above definition it is clear that $\mathcal{P}_{\mathcal{I}}(\mathcal{G}) \subseteq \mathcal{M}_{\mathcal{I}}(\mathcal{G})$. Theorem 2 relies on the following set of assumptions. The first one ensures that the model is capable of representing the ground truth distribution.

Assumption A.13 (Sufficient Capacity). *The joint ground truth distribution* $p_{\mathcal{I}}^*$ *is such that* $p_{\mathcal{I}}^* \in \mathcal{P}_{\mathcal{I}}(\mathcal{G}^*)$ *, where* \mathcal{G}^* *is the ground truth graph.*

In other words, there exists a ϕ such that $p_{\mathcal{I}}^* = p_{\mathcal{G}^{\mathcal{I}}}(\cdot | \phi)$. The second assumption generalizes the notion of faithfulness assumption to the interventional setting.

Assumption A.14 (\mathcal{I} - σ -faithfulness). Let $\mathbf{V} = (\mathbf{X}, \mathbf{C}^{\mathcal{I}})$, for any subset of nodes $A, B, C \subseteq \mathcal{V} \cup \mathbf{C}^{\mathcal{I}}$, and $I_k \in \mathcal{I}$

$$A \not \downarrow_{\mathcal{G}^{\mathcal{I}}} B \mid C \implies V_A \not \downarrow_{p_{\mathcal{G}^{\mathcal{I}}}} V_B \mid V_C$$

The above assumption implies that any conditional independency observed in the data must imply a σ -separation in the corresponding interventional ground truth graph.

Assumption A.15 (Finite differential entropy). For $\mathcal{I} = \{I_k\}_{k=0}^K$,

$$|\mathbb{E}_{p_{\mathcal{I}}^*} \log p_{\mathcal{I}}^*(\boldsymbol{X}, \boldsymbol{C})| < \infty.$$

The above assumption ensures that the hypothetical scenario where $S(\mathcal{G}^*)$ and $S(\mathcal{G})$ are both infinity is avoided. This is formalized in the lemma below taken from [15].

Lemma A.16 (Finiteness of the score function [15]). Under assumptions A.13 and A.15, $|S_{\mathcal{I}}(\mathcal{G})| < \infty$.

From the results of [15], we can now express the difference in score function between \mathcal{G}^* and \mathcal{G} as the minimization of KL divergence plus the difference in the regularization terms.

Lemma A.17 (Rewritting the score function [15]). Under assumptions A.13 and A.15, we have

$$\mathcal{S}(\mathcal{G}^*) - \mathcal{S}(\mathcal{G}) = \inf_{\phi} D_{KL}(p_{\mathcal{I}}^* || p_{\mathcal{G}^{\mathcal{I}}}(\cdot \mid \phi)) + \lambda(|\mathcal{G}| - |\mathcal{G}^*|).$$

We will now prove the following technical lemma (adapted from [15]) which we will be used in proving Theorem 2.

Lemma A.18. Let $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathcal{B})$ be a directed mixed graph, for a set of interventional targets $\mathcal{I} = \{I_k\}_{k=0}^K$, and $p^* \notin \mathcal{M}_{\mathcal{I}}(\mathcal{G})$, then

$$\inf_{p\in\mathcal{M}_{\mathcal{I}}(\mathcal{G}))} D(p^* \| p) > 0.$$

Proof. Let $V = (X, C^{\mathcal{I}})$, from theorem A.11, any $p \in \mathcal{M}_{\mathcal{I}}(\mathcal{G})$) admits a factorization of the form

$$p(\boldsymbol{V}) = \prod_{H \in [\mathcal{V}]_{\mathcal{G}}} p(\boldsymbol{V}_H \mid \boldsymbol{V}_T)$$

Let us define a new distribution \hat{p} as follows:

$$\hat{p}(\boldsymbol{V}) = \prod_{H \in [\mathcal{V}]_{\mathcal{G}}} \hat{p}(\boldsymbol{V}_H \mid \boldsymbol{V}_T),$$

where

$$\hat{p}(\boldsymbol{V}_H \mid \boldsymbol{V}_T) = \frac{p^*(\boldsymbol{V}_H, \boldsymbol{V}_T)}{p^*(\boldsymbol{V}_T)}.$$

From proposition A.12, we see that $\hat{p} \in \mathcal{M}_{\mathcal{I}}(\mathcal{G})$ and hence $p \neq \hat{p}$. We will show that

$$\hat{p} = \arg\min_{p \in \mathcal{M}_{\mathcal{I}}(\mathcal{G}))} D_{KL}(p^* \| p)$$

For an arbitrary $p \in \mathcal{M}_{\mathcal{I}}(\mathcal{G})$), consider the following:

$$\mathbb{E}_{p^*} \log \frac{\hat{p}(\mathbf{V})}{p(\mathbf{V})} = \mathbb{E}_{p^*} \sum_{H \in [\mathcal{V}]_{\mathcal{G}}} \log \frac{p^*(\mathbf{V}_H \mid \mathbf{V}_T)}{p(\mathbf{V}_H \mid \mathbf{V}_T)}$$
(21)

$$= \sum_{H \in [\mathcal{V}]_{\mathcal{G}}} \mathbb{E}_{p^*} \log \frac{p^*(\mathcal{V}_H \mid \mathcal{V}_T)}{p(\mathcal{V}_H \mid \mathcal{V}_T)}.$$
(22)

In the equation above, we leverage the linearity of expectation, which holds under Assumption A.15, ensuring that we don't sum infinities of opposite signs. We now show that each term in the right hand side of the above equality is an expectation of KL divergence which is always in $[0, \infty)$.

$$\mathbb{E}_{p^*}\log\frac{p^*(\boldsymbol{V}_H \mid \boldsymbol{V}_T)}{p(\boldsymbol{V}_H \mid \boldsymbol{V}_T)} = \int p^*(\boldsymbol{V}_T) \int p^*(\boldsymbol{V}_H \mid \boldsymbol{V}_T)\log\frac{p^*(\boldsymbol{V}_H \mid \boldsymbol{V}_T)}{p(\boldsymbol{V}_H \mid \boldsymbol{V}_T)} d\boldsymbol{V}_H d\boldsymbol{V}_T$$
(23)

$$= \int p^*(\boldsymbol{V}_T) D_{KL} \big(p^*(\cdot \mid \boldsymbol{V}_T) \| p(\cdot \| \boldsymbol{V}_T) \big) d\boldsymbol{V}_T.$$
(24)

Thus, $\mathbb{E}_{p^*} \log \frac{\hat{p}(V)}{p(V)} \in [0,\infty).$

We now show that $\hat{p} = \arg \min_{p \in \mathcal{M}(\operatorname{do}(I_k)(\mathcal{G}))} D_{KL}(p^* || p)$:

$$D_{KL}(p^* \| p) = \mathbb{E}_{p^*} \log \frac{p^*(\mathbf{V})}{\hat{p}(\mathbf{V})} \frac{\hat{p}(\mathbf{V})}{p(\mathbf{V})}$$
(25)

$$= \mathbb{E}_{p^*} \log \frac{p^*(\boldsymbol{V})}{\hat{p}(\boldsymbol{V})} + \mathbb{E}_{p^*} \log \frac{\hat{p}(\boldsymbol{V})}{p(\boldsymbol{V})}$$
(26)

$$= D_{KL}(p^* \| \hat{p}) + \mathbb{E}_{p^*} \log \frac{\hat{p}(\mathbf{V})}{p(\mathbf{V})}$$

$$\tag{27}$$

$$\geq D_{KL}(p^* \| \hat{p}) > 0.$$
⁽²⁸⁾

Since the expectations in (26) are both in $[0, \infty)$, splitting the expectation is valid. The very last inequality holds since $p^* \neq \hat{p}$. Thus,

$$\inf_{p \in \mathcal{M}_{\mathcal{I}}(\mathcal{G}))} D(p^* \| p) \ge D_{KL}(p^* \| \hat{p}) > 0.$$

We are now ready to prove Theorem 2. Recall,

Theorem 2. Let $\mathcal{I} = \{I_k\}_{k=0}^K$ be a family of interventional targets, let \mathcal{G}^* denote the ground truth directed mixed graph, $p^{(k)}$ denote the data generating distribution for I_k , and $\hat{\mathcal{G}} := \arg \max_{\mathcal{G}} \mathcal{S}(\mathcal{G})$. Then, under the Assumptions 1, A.13, A.14, and A.15, and for a suitably chosen $\lambda > 0$, we have that $\hat{\mathcal{G}} \equiv_{\mathcal{I}} \mathcal{G}^*$. That is, $\hat{\mathcal{G}}$ is \mathcal{I} -Markov equivalent to \mathcal{G}^* .

Proof. It is sufficient to show that for $\mathcal{G} \notin \mathcal{I}$ -MEC(\mathcal{G}^*), the score function of $\hat{\mathcal{G}}$ is strictly lower than the score function of \mathcal{G}^* , i.e., $\mathcal{S}(\mathcal{G}^*) > \mathcal{S}(\mathcal{G})$. Since $\mathcal{G} \notin \mathcal{I}$ -MEC(\mathcal{G}^*) and $p_{\mathcal{I}}^* \in \mathcal{M}_{\mathcal{I}}(\mathcal{G}^*)$ (by Assumption A.13), there must exist subsets of nodes $A, B, C \subseteq \mathcal{V} \cup C^{\mathcal{I}}$ such that

$$A \stackrel{\sigma}{\underset{\mathcal{G}}{\perp}} B \mid C \text{ and } A \stackrel{\sigma}{\underset{\mathcal{G}^{*}}{\not=}} B \mid C.$$

If no such subsets exist, then \mathcal{G} and \mathcal{G}^* impose the same σ -separation constraints and thus induce the same set of distributions. This would imply that $\mathcal{G} \in \mathcal{I}$ -MEC(\mathcal{G}^*), contradicting our assumption. Since $p_{\mathcal{I}}^* \in \mathcal{M}_{\mathcal{I}}(\mathcal{G}^*)$), it must be true that $V_A \not\perp_{p^{(k)}} V_B \mid V_C$ (Assumption A.14). Therefore $p\mathcal{I}^*$ doesn't satisfy the general directed Markov property with respect to $\mathcal{G}^{\mathcal{I}}$ and hence $p_{\mathcal{I}}^* \notin \mathcal{M}_{\mathcal{I}}(\mathcal{G})$. For convenience, let

$$\eta(\mathcal{G}) := \inf_{\mathcal{I}} D_{KL}(p_{\mathcal{I}}^* \| p_{\mathcal{G}^{\mathcal{I}}}(\cdot \mid \boldsymbol{\phi}))$$

Note that

$$\eta(\mathcal{G}) = \inf_{\phi} D_{KL}(p_{\mathcal{I}}^* \| p_{\mathcal{G}^{\mathcal{I}}}(\cdot \mid \phi)) \ge \inf_{p \in \mathcal{M}_{\mathcal{I}}(\mathcal{G})} D_{KL}(p^{(k)} \| p) > 0$$

where we use Lemma A.18 for the final inequality. Thus, from Lemma A.17

$$\mathcal{S}(\mathcal{G}^*) - \mathcal{S}(\mathcal{G}) = \eta(\mathcal{G}) + \lambda(|\mathcal{G}| - |\mathcal{G}^*|)$$
⁽²⁹⁾

Following [15], we now show that by choosing λ sufficiently small, the above equation is stictly positive. Note that if $|\mathcal{G}| \ge |\mathcal{G}^*|$ then $\mathcal{S}(\mathcal{G}^*) - \mathcal{S}(\mathcal{G}) > 0$. Let $\mathbb{G}^+ := \{\mathcal{G} \mid |\mathcal{G}| < |\mathcal{G}^*|\}$. Choosing λ such that $0 < \lambda < \min_{\mathcal{G} \in \mathbb{G}^+} \frac{\eta(\mathcal{G})}{|\mathcal{G}^*| - |\mathcal{G}|}$ we see that:

$$\lambda < \min_{\mathcal{G} \in \mathbb{G}^+} \frac{\eta(\mathcal{G})}{|\mathcal{G}^*| - |\mathcal{G}|} \tag{30}$$

$$\iff \lambda < \frac{\eta(\mathcal{G})}{|\mathcal{G}^*| - |\mathcal{G}|} \quad \forall \mathcal{G} \in \mathbb{G}^+$$
(31)

$$\iff \lambda(|\mathcal{G}^*| - |\mathcal{G}|) < \eta(\mathcal{G}) \quad \forall \mathcal{G} \in \mathbb{G}^+$$
(32)

$$\iff 0 < \eta(\mathcal{G}) + \lambda(|\mathcal{G}| - |\mathcal{G}^*|) = \mathcal{S}(\mathcal{G}^*) - \mathcal{S}(\mathcal{G}) \quad \forall \mathcal{G} \in \mathbb{G}^+.$$
(33)

Thus, every graph outside of the general directed Markov equivalence class of $(\mathcal{G}^*)^{\mathcal{I}}$ has a strictly lower score.

A.4 Characterization of Equivalence Class

Let $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathcal{B})$ be a directed mixed graph, and consider a family of interventional targets $\mathcal{I} = I_{k_{k=0}}^{K}$ with $I_0 = \emptyset$. From Proposition A.6, for any graph $\mathcal{G}_1 \in \mathcal{I}$ -MEC(\mathcal{G}), the corresponding augmented graph $\mathcal{G}_1^{\mathcal{I}}$ is equivalent to the acyclification $acy(\mathcal{G}^{\mathcal{I}})$ of $\mathcal{G}^{\mathcal{I}}$. Several prior works have studied the characterization of equivalence classes of acyclic directed mixed graphs (ADMGs), including [53, 30, 54]. We now provide a graphical notion of the \mathcal{I} -Markov equivalence class of a DMG \mathcal{G} . A graph \mathcal{G} is said to be *maximal* if there exists no inducing path (relative to the empty set) between any two non-adjacent nodes. An *inducing path* relative to a subset L is a path on which every non-endpoint node $i \in L$ is a collider on the path and every collider is an ancestor of an endpoint of the path. A *Maximal Ancestral Graph* (MAG) is one that is both ancestral and maximal. Given an ADMG $acy(\mathcal{G}^{\mathcal{I}})$, it is possible to construct a MAG over the variable set $V = (X, C^{\mathcal{I}})$ that preserves both the independence structure and ancestral relationships encoded in $acy(\mathcal{G}^{\mathcal{I}})$; see [55] for details. We denote $MAG(acy(\mathcal{G}^{\mathcal{I}}))$ to mean MAG that is constructed from $acy(\mathcal{G}^{\mathcal{I}})$. Therefore all the independencies encoded in $acy(\mathcal{G}^{\mathcal{I}})$ is also present in $MAG(acy(\mathcal{G}^{\mathcal{I}}))$. Before present the condition for MAG equivalence, we introduce the notion of discriminating path. A path $\pi = (i_0, \varepsilon_1, \ldots, i_{n-1}, \varepsilon_n, i_n)$ in $acy(\mathcal{G}^{\mathcal{I}})$ is called a *discriminating path* for i_{n-1} if (1) π includes at lest three edges; (2) i_{n-1} is a non-endpoint node on π , and is adjacent to i_n on π ; and (3) i_0 and i_n are not adjacent, and every node in between i_0 and i_{n-1} is a collider on π and is a parent of i_n . The following theorem from Spirtes and Richardson [53] characterizes the equivalence of MAGs.

Theorem A.19 (Spirtes and Richardson [53]). *Two MAGs* G_1 and G_2 are Markov equivalent if and only if:

- *1.* G_1 and G_2 have the same skeleton;
- 2. G_1 and G_2 have the same unshielded colliders; and
- 3. if π forms a discriminating path for *i* in \mathcal{G}_1 and \mathcal{G}_2 , then *i* is a collider on π if and only it is a collider on π in \mathcal{G}_2 .

Therefore, from Proposition A.6 and Theorem A.19, two DMGs \mathcal{G}_1 and \mathcal{G}_2 are equivalent if and only if $MAG(\operatorname{acy}(\mathcal{G}_1^{\mathcal{I}}))$ and $MAG(\operatorname{acy}(\mathcal{G}_2^{\mathcal{I}}))$ satisfying the conditions of Theorem 2, i.e., $MAG(\operatorname{acy}(\mathcal{G}_1^{\mathcal{I}}))$ and $MAG(\operatorname{acy}(\mathcal{G}_2^{\mathcal{I}}))$: (i) have the same skeleton, (ii) same unshielded colliders, and (iii) same discriminating paths with consistent colliders.

B Implementation Details

In this section we provide the implementation details of DCCD-CONF and the baseline models along with the details of the experimental setup.

B.1 Hutchinson trace estimator for computing log determinant of the Jacobian

Computing the log determinant of the Jacobian matrix present in (6) poses a significant challenge. However, following Behrmann et al. [40], in Section 3.3.1 showed that the log-determinant of the Jacobian can be estimated using the following estimator

$$\log \left| \det \left(\mathbf{J}_{\mathbf{f}_{x}^{(I_{k})}}(\boldsymbol{X}) \right) \right| = \mathbb{E}_{n \sim p_{\mathbb{N}}(N)} \left[\sum_{m=1}^{n} \frac{(-1)^{m+1}}{m} \cdot \frac{\operatorname{Tr} \left\{ \mathbf{J}_{\mathbf{U}_{k}\mathbf{g}_{x}}^{m}(\boldsymbol{X}) \right\} - \operatorname{Tr} \left\{ \mathbf{J}_{\mathbf{U}_{k}\mathbf{g}_{z}}^{m}(\boldsymbol{Z}) \right\}}{p_{\mathbb{N}}(\ell \geq m)} \right].$$

The estimator above still has a major drawback: computing the $Tr(\mathbf{J}_{\mathbf{U}_k\mathbf{g}})$ still requires $\mathcal{O}(d^2)$ gradient calls to compute exactly. Fortunately, *Hutchinson trace estimator* [44] can be used to stochastically approximate the trace of the Jacobian matrix. This then results in the following estimator that can be computed efficiently via reverse-mode automatic differentiation

$$\log \left| \det \left(\mathbf{J}_{\mathbf{f}_{x}^{(I_{k})}}(\boldsymbol{X}) \right) \right| = \underset{n \sim p_{\mathbb{N}}(N), \boldsymbol{V} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})}{\mathbb{E}} \left[\sum_{m=1}^{n} \frac{(-1)^{m+1}}{m} \times \frac{\boldsymbol{V}^{\top} \{ \mathbf{J}_{\mathbf{U}_{k}\mathbf{g}_{x}}^{m}(\boldsymbol{X}) \} \boldsymbol{V} - \boldsymbol{V}^{\top} \{ \mathbf{J}_{\mathbf{U}_{k}\mathbf{g}_{z}}^{m}(\boldsymbol{Z}) \} \boldsymbol{V}}{p_{\mathbb{N}}(\ell \geq m)} \right].$$
(34)

B.2 Parameter update via score maximization

As described in Section 3, the model parameters are updated in two stages. In the first stage, the parameters of the neural networks and the Gumbel-Softmax distribution, used to sample adjacency matrices, are updated via backpropagation using stochastic gradient descent. Since (5) forms an implicit block of an implicit normalizing flow, following, we directly estimate the gradients of $\hat{S}(\mathbf{B})$ with respect to x and $\phi = (\theta, \mathbf{B})$. The gradient computation involves two terms: $\frac{\partial}{\partial(\cdot)} \log \det(\mathbf{I} + \mathbf{J}_{\mathbf{g}_x}(x, \phi))$ and $\frac{\partial \hat{S}}{\partial z} \frac{\partial z}{\partial(\cdot)}$, where (\cdot) is a placeholder for x and ϕ . From [38, 40], we use the following unbiased estimators for the gradients:

$$\frac{\partial \log \det(\mathbf{I} + \mathbf{J}_{\mathbf{g}_x}(\boldsymbol{x}, \boldsymbol{\phi}))}{\partial(\cdot)} = \mathbb{E}_{n \sim p(N), \boldsymbol{V} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})} \left[\left(\sum_{k=0}^n \frac{(-1)^k}{P(N \ge k)} \boldsymbol{V}^\top \mathbf{J}^k \right) \frac{\partial \mathbf{J}_{\mathbf{g}_x}(\boldsymbol{x}, \boldsymbol{\phi})}{\partial(\cdot)} \boldsymbol{V} \right].$$
(35)

On the other hand, $\frac{\partial \hat{S}(\mathbf{B})}{\partial z} \frac{\partial z}{\partial (\cdot)}$ can be computed according to the *implicit function theorem* as follows:

$$\frac{\partial \hat{\mathcal{S}}(\mathbf{B})}{\partial \boldsymbol{z}} \frac{\partial \boldsymbol{z}}{\partial (\cdot)} = \frac{\hat{\mathcal{S}}(\mathbf{B})}{\partial \boldsymbol{z}} \mathbf{J}_{\mathbf{G}_{z}}^{-1}(\boldsymbol{z}) \frac{\mathbf{G}(\boldsymbol{x}, \boldsymbol{z}, \boldsymbol{\phi})}{\partial (\cdot)},$$
(36)

where $\mathbf{G}_z(z) = \mathbf{g}_z(z, \phi) + z$, and recall that $\mathbf{G}(x, z, \phi) = \mathbf{g}_x(x, \phi) + x + \mathbf{g}_z(z, \phi) + z$. See [38] for more details. The procedure SGUPDATE shown in Algorithm 1 performs the gradient computation in (35) and (36).

In the second stage, the entries of the covariance matrix of the endogenous noise distribution are updated column-wise by solving a sequence of Lasso optimization problems. The complete parameter update procedure is summarized in Algorithm 1.

B.3 DCCD-CONF and the baselines code details

DCCD-CONF. We implemented our framework using the libraries Pytorch and Scikit-learn in Python and the code used in running the experiments can be found in the codes folder within the supplementary materials. We plan to make the code publicly available on GitHub upon publication of the paper.

Algorithm 1 PARAMETER UPDATE

- **Require:** Family of interventional targets $\mathcal{I} = \{I_k\}_{k=1}^K$, interventional dataset $\{x^{(i,k)}\}_{i=1,k=1}^{N_k,K}$, regularization coefficients λ and ρ .
- **Ensure:** Learned neural network parameters $\hat{\theta}$, graph structure parameters $\hat{\beta}$, confounder-noise distribution parameters $\hat{\Sigma}_Z$. 1: Initialize the parameters: $\theta^{(0)} \sim p_{\theta}(\theta)$, $\mathcal{B}^{(0)} \sim p_{\mathcal{B}}(\mathcal{B})$, and $\Sigma_Z = \mathbf{I}$
- 2: Iteration counter: t = 0
- 3: while NOT CONVERGED do
- 4: for k = 1 to K do
- $t \leftarrow t + 1$ 5:
- $\mathbf{W} \leftarrow (\mathbf{\Sigma}_Z^{(t)})_{\mathcal{U}_k,\mathcal{U}_k}$ 6:

Compute score function $\tilde{\mathcal{L}}(\mathcal{B}^{(t)}, \boldsymbol{\theta}^{(t)}, \mathbf{W}, I_k)$ 7:

- $\mathcal{B}^{(t+1)}, \boldsymbol{\theta}^{(t+1)} \leftarrow \text{SGUPDATE}(\tilde{\mathcal{L}}, \mathcal{B}^{(t)}, \boldsymbol{\theta}^{(t)})$ 8:
- 9: for j = 1 to d do
- Push j-th row and column in W to the end 10:

 $\boldsymbol{\beta} \leftarrow \texttt{lasso}(\mathbf{W}_{11}, \boldsymbol{s}_{12}, \rho)$ 11:

 $\boldsymbol{w}_{12} \leftarrow \mathbf{W}_{11} \boldsymbol{\beta}$ 12:

end for $(\boldsymbol{\Sigma}_{Z}^{(t+1)})_{\mathcal{U}_{k},\mathcal{U}_{k}} \leftarrow \mathbf{W}$ 13:

14:

end for 15:

16: end while

17: $\hat{\boldsymbol{\theta}}, \hat{\mathcal{B}}, \hat{\boldsymbol{\Sigma}}_{Z} \leftarrow \boldsymbol{\theta}^{(t)}, \mathcal{B}^{(t)}, \boldsymbol{\Sigma}_{Z}^{(t)}$

return $\hat{\theta}, \hat{\mathcal{B}}, \hat{\Sigma}_Z$

Starting with an initialization of the model parameters $(\theta^{(0)}, \mathbf{B}^{(0)}, \boldsymbol{\Sigma}_Z^{(0)})$, we iteratively alternate between maximizing the score function with respect to $(\theta^{(t)}, \mathbf{B}^{(t)})$ and $\Sigma_Z^{(t)}$, as described in Algorithm 1. Standard stochastic gradient updates are used for $(\theta^{(t)}, \mathbf{B}^{(t)})$, while coordinate gradient descent, implemented via the Scikit-learn library, is applied to $\Sigma_Z^{(t)}$. For modeling the causal function \mathbf{g}_x , we follow the setup of Sethuraman et al. [27], employing neural networks (NNs) with dependency masks parameterized by a Gumbel-softmax distribution. The log-determinant of the Jacobian is computed using a power series expansion combined with the Hutchinson trace estimator. To mitigate bias from truncating the power series expansion, the number of terms is sampled from a Poisson distribution, as detailed in Section 3.3 and Appendix B.1. The final objective is optimized using the Adam optimizer [56].

The learning rate in all our experiments was set to 10^{-2} . The neural network models used in our experiments contained one multi-layer perceptron layer. No nonlinearities were added to the neural networks for the linear SEM experiments. We used tanh activation for the nonlinear SEM experiments and for the experiments on the perturb-CITE-seq data set. The graph sparsity regularization constant λ was set to 10^{-3} for all the experiments. The sparsity inducing regularization constant for the inverse covariance matrix of the confounder distribution, ρ , was set to 10^{-2} in all the experiments. The models were trained and evaluated on NVIDIA RTX6000 GPUs.

Baselines. For NODAGS-Flow, we used the code provided by authors [27] available at https: //github.com/Genentech/nodags-flows. The default values were set for the hyperparameters. We implemented the LLC algorithm based on the details provided in [26]. The implementation can be found within the codes/baselines folder in the supplementary materials. For FCI, we used the implementation that is available in the causallearn python library (https://github. com/py-why/causal-learn). For DCDI, we used the codebase provided by the authors [15], available at https://github.com/slachapelle/dcdi. The default hyperparameters were used while training and evaluating the model.

B.4 Experimental setup

In this section, we describe how the data sets were generated for the various experiments conducted.

B.4.1 Synthetic Experiments

We begin by sampling a directed graph using the Erdős-Rényi (ER) random graph model with an edge density of 2 unless specified otherwise, which determines the directed edges in the DMG \mathcal{G} . Next, we generate a random matrix and project it onto the space of positive definite matrices to obtain the confounder covariance matrix Σ_Z , setting the maximum exogenous noise standard deviation to 0.5 unless specified otherwise. The nonzero off-diagonal entries of Σ_Z correspond to the bidirectional edges in \mathcal{G} . For the linear SEM, edge weights are sampled uniformly from $\text{Unif}((-0.9, -0.2) \cup (0.2, 0.9))$. In all experiments except those on non-contractive SEMs, the edge weight matrix is rescaled to ensure a Lipschitz constant of less than one. For nonlinear SEMs, we apply a tanh nonlinearity to the linear system defined by the edge weights, i.e.,

$$\boldsymbol{x} = \tanh(\mathbf{W}^{\top}\boldsymbol{x} + \boldsymbol{z}),$$

where W is the weighted adjacency matrix.

Impact of Confounder Count In this experiment, the number of observed nodes in the graph is fixed at d = 10. Training data consists of combination of observational data and single-node interventions over all nodes, i.e., $\mathcal{I} = \emptyset \cup \{\{i\} \mid i \in [d]\}$, with 1000 samples per interventional setting. The confounder ratio (number of confounders divided by the number of nodes) is varied from 0.2 to 0.8.

Scaling with Number of Nodes Here, we fix the confounder ratio at 0.4. The total number of nodes in the graph is varied from 10 to 80. As in the previous setup, training data consists of combination of observational data and single-node interventions across all nodes, with 1000 samples per interventional setting.

Scaling with Interventions We fix d = 10 and set the confounder ratio to 0.4. The number of interventions during training varies from 0 to d. Zero interventions corresponds to observational data. When fewer than d interventions are provided, the intervened nodes are selected arbitrarily. Each interventional setting consists of 1000 samples.

Non-Contractive SEM In this case, we explicitly enforce a non-contractive causal mechanism **F** by rescaling edge weights to ensure that the Lipschitz constant of the edge weight matrix exceeds one. We set d = 10 and provide observational data and single-node interventions across all nodes, with 1000 samples per intervention. The confounder ratio varies between 0.2 and 0.8.

Scaling with Training Samples To examine the sample requirements of DCCD-CONF, we set the confounder ratio to 0.4. Training data consists of observational data and single-node interventions over all nodes, while the number of samples per intervention is varied from 500 to 2500.

Scaling with outgoing edge density In this case, the outgoing edge density of the ER random graphs is varied from 1 to 4. The confounder ratio is set to 0.4 and the number of nodes d = 10. The training data consists of observational data and single node experiments over all the nodes in the graph.

Scalinf with noise standard deviation In this setting, we vary the maximum noise standard deviation between 0.2 and 0.8. The confounder ratio is set to 0.4 and the number of nodes d = 10. The training data consists of observational data and single node experiments over all the nodes in the graph.

Evaluation Metrics Across all experiments, we use Structural Hamming Distance (SHD) to evaluate the accuracy of the estimated directed edges relative to the ground truth. SHD measures the number of modifications (edge additions, reversals, and deletions) required to match the estimated graph to the ground truth. For DCCD-CONF and NODAGS-Flow we fix a threshold value of 0.8 for the estimated adjacency matrix. The recovery of bidirectional edges is assessed using the F1 score, which is defined as:

F1 score = $2 \frac{\text{precision} \times \text{recall}}{\text{precision} + \text{recall}}$



Figure 6: Performance comparison between DCCD-CONF and baseline methods on causal graph recovery and confounder identification, evaluated across varying model parameters. Each subplot details a specific experimental setting.

where

precision
$$= \frac{TP}{TP + FP}$$
, recall $= \frac{TP}{TP + FN}$

and TP, FP, FN denote true positives, false positives, and false negatives, respectively. We use a threshold of 0.01 for the estimated covariance matrix to identify the bidirectional edges.

Additionally, we also measure the performance of DCCD-CONF and the baselines using *Area Under Precision-Recall Curve* (AUPRC) as the error metric. AUPRC computes the area under the precision-recall curve evaluated at various threshold values (the higher the better).

B.4.2 Gene Perturbation Data set

1

The dataset was obtained from the Single Cell Portal of the Broad Institute (accession code SCP1064). Following the experimental setup of Sethuraman et al. [27], we filtered out cells with fewer than 500 expressed genes and removed genes expressed in fewer than 500 cells. Due to computational constraints, we selected a subset of 61 perturbed genes (Table 2) from the full genome. The three experimental conditions—co-culture, IFN- γ , and control—were partitioned into separate datasets, and models were trained and evaluated on each condition independently.

Table 2: The list of chosen genes from Perturb-CITE-seq dataset [48].

ACSL3	ACTA2	B2M	CCND1	CD274	CD58	CD59	CDK4	CDK6
CDKN1A	CKS1B	CST3	CTPS1	DNMT1	EIF3K	EVA1A	FKBP4	FOS
GSEC	GSN	HASPIN	HLA-A	HLA-B	HLA-C	HLA-E	IFNGR1	IFNGR2
ILF2	IRF3	JAK1	JAK2	LAMP2	LGALS3	MRPL47	MYC	P2RX4
PABPC1	PAICS	PET100	PTMA	PUF60	RNASEH2A	RRS1	SAT1	SEC11C
SINHCAF	SMAD4	SOX4	SP100	SSR2	STAT1	STOM	TGFB1	TIMP2
TM4SF1	TMED10	TMEM173	TOP1MT	TPRKB	TXNDC17	VDAC2		

C Additional Results

C.1 Additional synthetic experiments

Experiments on Non-contractive DAGs. We evaluated the performance of DCCD-CONF and baseline methods on non-contractive SEMs, where the ground truth DMG is acyclic. While DCCD-CONF assumes a contractive causal mechanism, we adapt it for non-contractive settings using the



Figure 7: Comparison of DCCD-CONF and baseline methods on causal graph recovery and confounder identification, measured using AUPRC across varying model parameters.

preconditioning trick proposed by Sethuraman et al. [27]. This approach introduces a learnable diagonal preconditioning matrix Λ , transforming the causal mechanism as follows:

$$\hat{\mathbf{g}}_x = \mathbf{\Lambda}^{-1} \circ \mathbf{g}_x \circ \mathbf{\Lambda},$$

where g_x , as defined in (8), remains contractive (see Sethuraman et al. [27] for details). We vary the confounder ratio, and the results are summarized in Figure 6a. As shown in Figure, DCCD-CONF effectively learns the ADMG even in non-contractive SEM settings, demonstrating competitive performance against the baselines.

Performance comparison vs. sample size We also assess the sample requirements of DCCD-CONF. Figure 6b summarizes the results obtained by varying the number of samples per intervention. As shown in the figure, when the confounder ratio is 0.4, DCCD-CONF achieves low SHD even with 500 samples per intervention and attains near-perfect accuracy from 2000 samples onward. However, performance declines slightly as the confounder ratio increases.

Performance comparison vs. max endogenous noise st. deviation In this setting, we compare DCCD-CONF with the baseline by varying the maximum standard deviation of the endogenous noise terms between 0.2 and 0.8, the results are summarized in Figure 6c. DCCD-CONF outperforms the baselines for all noise standard deviations. However, the performance of the models does deteriorate slightly as the noise standard deviation increases.

Performance comparison vs. outgoing edge density In this case, the expected number of outgoing edges from each node is varied between 1 and 4. This affects the sparsity of the resulting graph. The results are summarized in Figure 6d. As seen from Figure 6d, DCCD-CONF still outperforms the baselines, even though the performance of all the models worsens as the edge density increases.



Figure 8: Performance comparison between DCCD-CONF and FCI-JCI with respect to modified SHD as the error metric. The number of observed nodes was set to d = 5. For the left plot, all single node interventions along with the observational data were provided as training data. For the right plot, observational data and interventions over two of the nodes (randomly chosen) were used as training data.

C.2 Additional performance metrics

In addition to SHD for directed edge recovery and F1 score for bi-directional edge recovery. We also AUPRC to compare DCCD-CONF and the baseline for all of the experimental settings stated in Appendix B.4.1. The results are summarized in Figure 7. Overall, DCCD-CONF performs better than LLC on nonlinear SEMs across all the settings, while achieving perfect AUPRC scores in several cases.

C.3 Comparison with FCI-JCI

We also compare the performance of DCCD-CONF with FCI-JCI [29], which is an extension of FCI algorithm that is capable of handling multiple contexts (in this case interventional settings). FCI-JCI outputs a Partial Ancestral Graph (PAG), which is a graph structure that represents the equivalence class of MAGs. We define a modified SHD score in order to check if the DMG estimated by DCCD-CONF belongs to the same equivalence class of the ground truth DMG. To that end, we convert the ground-truth DMG and the estimated DMG to their augmented DMGs and then construct the MAG of the acyclified version of the augmented DMGs. The modified SHD score then computes the discrepancies in the conditions of Theorem A.19, i.e., we count: (i) the number of extra edges (N_1) using the skeletons of the estimated MAG and ground truth MAG, (2) number of mismatched unshielded colliders (N_2) , and (3) discrepancies in the discriminating paths (N_3) . Similarly, for FCI-JCI, we count the disagreements between the ground-truth MAG and the estimated PAG, i.e., mismatch in skeleton (N_1) , mismatch in unshielded colliders (N_2) , invariant edge orientation discrepancies (N₃). Finally, the modified SHD = $N_1 + N_2 + N_3$. We compare DCCD-CONF and JCI-FCI over two different settings: (i) the training data consists of observational data and single node interventions over all the nodes in the graph, and (ii) the training data consists of observational data and interventions over 2 nodes (randomly chosen) in the graph. Due to the complexity of computing the modified SHD (as it involves iterating over the discriminating paths and inducing paths) we fix the number of observed nodes to be d = 5. In the both the cases, the confounder ratio is varied between 0.2 and 0.8, and nonlinear SEM is used to generate the data. The results are summarized in Figure 8.

As seen from Figure 8, DCCD-CONF outperforms FCI-JCI in the both the settings. However, the performance does decrease as the number of training interventions reduces. We attribute this to the increase sample requirements as the number of training interventions goes down.