
Near-Optimal Experiment Design in Linear non-Gaussian Cyclic Models

Ehsan Sharifian

EPFL, Lausanne, Switzerland
ehsan.sharifian@epfl.ch

Saber Salehkaleybar

Leiden University, Leiden, The Netherlands
s.salehkaleybar@liacs.leidenuniv.nl

Negar Kiyavash

EPFL, Lausanne, Switzerland
negar.kiyavash@epfl.ch

Abstract

We study the problem of causal structure learning from a combination of observational and interventional data generated by a linear non-Gaussian structural equation model that might contain cycles. Recent results show that using mere observational data identifies the causal graph only up to a permutation-equivalence class. We obtain a combinatorial characterization of this class by showing that each graph in an equivalence class corresponds to a perfect matching in a bipartite graph. This bipartite representation allows us to analyze how interventions modify or constrain the matchings. Specifically, we show that each atomic intervention reveals one edge of the true matching and eliminates all incompatible causal graphs. Consequently, we formalize the optimal experiment design task as an adaptive stochastic optimization problem over the set of equivalence classes with a natural reward function that quantifies how many graphs are eliminated from the equivalence class by an intervention. We show that this reward function is adaptive submodular and provide a greedy policy with a provable near-optimal performance guarantee. A key technical challenge is to efficiently estimate the reward function without having to explicitly enumerate all the graphs in the equivalence class. We propose a sampling-based estimator using random matchings and analyze its bias and concentration behavior. Our simulation results show that performing a small number of interventions guided by our stochastic optimization framework recovers the true underlying causal structure.

1 Introduction

Learning causal relationships among variables of interest in a complex system is a central goal in empirical sciences, forming the foundation for prediction, intervention, and explanation [25]. These relationships are typically represented by causal graphs in which an edge from variable X to variable Y indicates that X is a direct cause of Y . Much of the existing literature on causal structure learning assumes that the underlying causal graph is a Directed Acyclic Graph (DAG). However, many natural and engineered systems include feedback mechanisms that give rise to cycles in their causal representations. Such cyclic structures emerge in equilibrium models, low-frequency temporal sampling of dynamical systems, and various biological networks [3].

In acyclic settings, a variety of methods (such as those based on conditional independence tests) allow us to recover the skeleton and orientations of the causal graph from observational data. These methods often fail when applied to graphs with cycles. More specifically, observational data alone does not even suffice for learning the skeleton, let alone orienting the edges of the graph [23].

Interventions, i.e., actively perturbing the system and observing the resulting distributional changes can allow us to learn the graphs even with cycles. The *experiment design* problem studies how to best design interventions in order to maximize the information gained about the causal structure. Again, unlike the case of DAGs, where a body of work on experiment design exist, the work in the cyclic setting [23] is few and far in between. This is in part due to the fact that, unlike in DAGs, where an intervention on a subset of the vertices orients all the edges incident to them, in cyclic directed graphs, performing experiments in some cases, neither leads to learning the presence of edges nor orienting them [23].

In this work, we study the problem of causal structure learning from a combination of observational and interventional data in systems governed by linear non-Gaussian structural causal models (SCMs) that may contain *cycles*. Our main contributions are as follows:

- We establish that, under linear non-Gaussian assumptions, the causal graph can be identified from observational data only up to an equivalence class. We further provide a combinatorial characterization of the equivalence class, where each graph corresponds to a perfect matching in a bipartite graph (Section 4.2). Additionally, we show that the condensation graph, or Strongly Connected Components (SCCs), of the true causal graph can also be identified (4.1).
- The bipartite representation allows us to analyze how atomic interventions constrain the space of causal graphs. Specifically, we show that each intervention reveals one edge of the true matching and eliminates all incompatible graphs from the equivalence class (Section 5). Therefore, we can formulate the optimal experiment design problem as an adaptive stochastic optimization over the space of equivalence classes, using a reward function that quantifies the number of eliminated graphs following an intervention. We prove that this reward function is adaptive submodular and hence a greedy policy has provably near-optimal performance guarantees for intervention design (Section 6).
- To address the computational challenge of reward estimation, we propose a sampling-based estimator based on random matchings and provide a theoretical analysis of its bias and concentration properties (Section 7).
- Experiments show that our adaptive strategy outperforms other heuristic methods and closely matches the feedback vertex set (FVS) lower bound (see more details in Section ??).

2 Related Work

Causal discovery is concerned with learning the underlying causal graph which encodes both the existence and direction of edges among variables of interest in a system. From observational data alone, the causal graph can be recovered only up to its Markov equivalence class (MEC). [29] provided necessary and sufficient conditions for Markov equivalence of directed graphs (DGs) and proposed an algorithm for structure learning up to the MEC [28]. [24] extended the applicability of a classic algorithms for constraint-based causal discovery (i.e., FCI [32]) to cyclic DGs and showed that it can recover the structure up to the MEC in this more general setting.

To address latent confounding and nonlinear mechanisms, [6] introduced σ -connection graphs (σ -CGs)—a flexible class of mixed graphs—and developed a discovery algorithm that handles both latent variables and interventional data. [10] focused on distributional equivalence for linear Gaussian models, providing necessary and sufficient conditions and proposing a score-based learning method. In the context of linear non-Gaussian models, [21] extended the Independent Component Analysis (ICA)-based approach of [31] to cyclic graphs to learn the equivalence class from the observational data.

The research on experiment design in acyclic models has focused on various aspects, including cost minimization and efficient learning under budget constraints. [5] provided worst-case bounds on the number of experiments to identify the graph. [14] introduced adaptive and exact non-adaptive algorithms for singleton interventions. [13] proposed optimal one-shot and adaptive heuristics to minimize undirected edges. [30] offered lower bounds for experiment design using separating systems, and [20] introduced a stage-wise approach in the presence of latent variables. In cost-aware settings, [19] developed an optimal algorithm for variable-cost interventions without size constraints. [12] and [33] designed approximation algorithms for trees and general graphs, respectively.

Fixed-budget design was initiated by [8] with a greedy Monte Carlo-based approach. [9] improved MEC sampling efficiency, and [7] developed an exact experiment design algorithm for tree structures. [2] proposed an exact method to enumerate DAGs post-intervention. [36, 37] showed that MEC counting and sampling can be done in polynomial time. Bayesian methods include the adaptive submodular algorithm [1] and its extension [34] to optimize both intervention targets and their assigned values.

To the best of our knowledge, only a handful of work on experiment design in cyclic graphs exists. For general SCMs, [23] proposed an experiment design approach that can learn both cyclic and acyclic graphs. They provided a lower bound on the number of experiments required to guarantee the unique identification of the causal graph in the worst case, showing that the proposed approach is order-optimal in terms of the number of experiments up to an additive logarithmic term. Our approach differs from [23] in three key aspects. First, their framework is non-adaptive, that is, all interventions are selected in advance based solely on observational data. In contrast, our method is adaptive as each experiment is selected based on the outcomes of prior interventions. Second, [23] allow intervening on multiple variables in a single experiment. Although they also considered the setup where the number of interventions per experiment is bounded, this bound should be greater than the size of the largest strongly connected component in the graph minus one. By contrast, our framework limits each experiment to a single-variable intervention, making it more practical in settings where fine-grained or limited interventions are preferred. Third, the approach in [23] is designed for general SCMs and relies on conditional independence testing to recover the causal structure. In our work, we focus specifically on linear non-Gaussian models, which enables us to use results from ICA to infer the underlying graph.

3 Preliminaries and Problem Formulation

Consider a *Structural Causal Model* (SCM) $\mathcal{C} := (\mathbf{S}, P_{\mathbf{e}})$, where \mathbf{S} consists of n structural equations and $P_{\mathbf{e}}$ is the joint distribution over exogenous noise terms [26]. In a linear SCM (LSCM), the structural equations take the following form:

$$\mathbf{x} = W\mathbf{x} + \mathbf{e}, \quad (1)$$

where $\mathbf{x} \in \mathcal{X} \subseteq \mathbb{R}^n$ denotes the vector of observable variables, $\mathbf{e} \in \mathcal{E} \subseteq \mathbb{R}^n$ represents the vector of independent exogenous noises, and $W \in \mathbb{R}^{n \times n}$ is the matrix of linear coefficients. The directed graph $G = (V, E)$ induced by the SCM captures the causal structure: node i corresponds to the variable x_i , and for any nonzero coefficient $W_{ij} \neq 0$, a directed edge $(j, i) \in E$ is drawn. The binary adjacency matrix $B_G \in \{0, 1\}^{n \times n}$ is defined as $[B_G]_{ij} = \mathbb{1}_{\{W_{ij} \neq 0\}}$, and the “free parameters” are the nonzero entries $\text{supp}(W) = \{(i, j) : W_{ij} \neq 0\}$. We impose the following assumptions:

Assumption 3.1. The model has no self-loops: $W_{ii} = 0$ for all $i \in [n]$.

Remark. This assumption can be made without loss of generality, as any self-loop can be algebraically removed by reparametrizing the structural equations. For a detailed explanation, see [22].

Assumption 3.2. The matrix $I - W$ is invertible.

Assumption 3.3. The components of \mathbf{e} are jointly independent, and at most one component is Gaussian. That is,

$$P_{\mathbf{e}} \in \mathcal{P}(\mathcal{E}) := \left\{ P_{\mathbf{e}} : P_{\mathbf{e}} = \prod_{i=1}^n P_{e_i}, \text{ with at most one } P_{e_i} \text{ Gaussian} \right\}.$$

Under these assumptions, the model admits a unique solution: $\mathbf{x} = (I - W)^{-1}\mathbf{e} = A\mathbf{e}$. The observational distribution $P_{\mathbf{x}}$ is thus given by the push-forward measure $P_{\mathbf{x}} = (I - T_W)^{-1}_{\#}(P_{\mathbf{e}})$, where T_W denotes the linear map associated with W .

We consider *perfect interventions*, in which the causal mechanism for a variable x_i is replaced with an exogenous noise \tilde{e}_i , removing all incoming edges to x_i . The interventional SCM is given by:

$$\mathbf{x} = W^{(i)}\mathbf{x} + \mathbf{e}^{(i)}, \quad (2)$$

where the i -th row of $W^{(i)}$ is zeroed out and the noise in the i -th row of \mathbf{e} is replaced with the new independent noise \tilde{e}_i . We denote the new noise vector by $\mathbf{e}^{(i)}$. While $(I - W^{(i)})$ might not always

be invertible, we assume that the probability it becomes singular is zero under mild randomness conditions on the rows of W . Thus, we treat $(I - W^{(i)})$ as invertible so that the interventional distribution $P(\mathbf{x}|\text{do}(x_i)) = (I - T_{W^{(i)}})^{-1}_{\#}(P_{\mathbf{e}^{(i)}})$ is well defined.

In interventional structure learning, we perform K distinct experiments to reduce uncertainty over the underlying causal graph. Each experiment involves a *perfect intervention* on a single variable. We denote the set of intervention targets by $\mathcal{I} = \{i_1, i_2, \dots, i_K\} \subseteq [n]$, where $i_j \in [n]$ indicates that the j -th experiment intervenes on variable x_{i_j} . In the single-intervention setting we consider, each experiment replaces the structural assignment of the target variable with a new exogenous noise term (cf. Equation 2). As a result, data from the j -th experiment is drawn from the interventional distribution $P(\mathbf{x}|\text{do}(x_{i_j}))$.

Each such experiment provides partial information about the true structural matrix W by revealing the causal parents of the intervened variable. The overall objective is to combine observational data with strategically designed interventions to eliminate incorrect graph candidates and recover the true causal structure with minimal experimentation.

Before formulating an optimization strategy for experiment design, we first investigate the extent to which the observational distribution constrains the underlying causal graph and matrix W . This characterizes the *observational equivalence class* of graphs, that is, the set of causal structures that cannot be distinguished based on observational data alone.

4 Distribution Equivalent Graphs

We aim to characterize the class of linear structural causal models (LSCMs) and their associated graphs that generate the same distribution over observable variables. There are two related but distinct notions in the literature:

- **Distribution-entailment equivalence**, which considers whether two parameterized LSCMs yield the same observational distribution.
- **Distribution equivalence**, which concerns whether two graph structures admit the same set of observable distributions across all compatible LSCMs.

These concepts are formally defined and compared in Appendix A. In our setting—linear SCMs with non-Gaussian noise—the two notions coincide due to identifiability guarantees from ICA. For a detailed discussion, see also Lacerda et al. [22].

In the remainder of this section, we focus on a graph operation that connects distribution-equivalent graphs.

Definition 4.1 (Cycle Reversion). Let C be a cycle in the graph G . The **cycle reversion (CR)** operation involves swapping the rows of each member of C in $I + B_G$ with the row corresponding to its subsequent node in the cycle. This operation reverses the direction of the cycle C . Moreover, any edge from a node outside C that was originally a parent of some $X \in C$ will now point to the predecessor of X in the original cycle C (see Figure 1a).

Propositions A.3 and A.4 (see Appendix A) shows that distribution-equivalent graphs can be transformed into one another through a sequence of cycle reversions.

4.1 Strongly Connected Components

In this section, we examine an important property that remains invariant under cycle reversion. To establish this, we first define strongly connected components (SCCs) in directed graphs.

Definition 4.2 (Strongly Connected Component (SCC)). Two vertices u and v in a directed graph G are said to be strongly connected if there are directed paths from u to v and from v to u . This defines an equivalence relation on the vertex set, whose equivalence classes are called *strongly connected components* (SCCs). Each SCC is a maximal subgraph in which every pair of vertices is strongly connected, and no additional vertex from G can be added without violating this property. The collection of SCCs forms a partition of the vertex set of G . See [4, Section 22.5]




Figure 1: (a) Graph representation of cycle reversion. (b) The figure highlights representative nodes from an SCC to show that their reachability is maintained, although the full SCC is not depicted.

Definition 4.3 (Condensation Graph). The *condensation graph* of a directed graph G is obtained by contracting each SCC into a single vertex and mapping all edges between SCCs to a single edge. The resulting graph is a directed acyclic graph (DAG).

Theorem 4.4.¹ *Strongly connected components remain unchanged in distribution-equivalent graphs.*

Figure 1b illustrates the core idea behind the proof by showing how cycle reversion operations preserve strongly connected components. The following corollary is a direct consequence of the above theorem and the definition of the condensation graph in 4.3.

Corollary 4.5. *The condensation graphs of all distribution-equivalent graphs are identical.*

Additionally, note that for a subset of vertices in the condensation graph that correspond to single vertices in the original graph, their corresponding row in the recovered matrix via *ICA* can be identified. This is because cycles only occur within an SCC, and if an SCC consists of a single vertex, no cycle passes through it, eliminating any *CR* ambiguity. Therefore, the corresponding row can be determined. Furthermore, if the original graph is acyclic, all its SCCs contain only one vertex, meaning that row permutation ambiguity is entirely resolved. As a result, the original data-generating graph can be inferred purely from observational data, a result previously established in Linear Non-Gaussian Acyclic Models (LiNGAM) in [31].

4.2 Matching in Bipartite Graphs

Graphs in a distribution-entailment equivalence class correspond to perfect matchings in a bipartite graph. Given $I - W_{\text{ICA}}$, recovered via *ICA* from observational data, any equivalent graph can be written as $P_\pi D(I - W_{\text{ICA}})$, where P_π is a row permutation matrix such that its diagonal entries remain nonzero.

Define a bipartite graph $G_b = (\{r_1, \dots, r_n\}, [n], E)$, where each node r_i represents the i -th row of the matrix $I - W_{\text{ICA}}$, and an edge $(r_i, j) \in E$ exists if and only if $W_{\text{ICA}}[i, j] \neq 0$. Each valid permutation π defines a perfect matching $r_i \mapsto \pi(i)$, and vice versa. Thus, distribution-equivalent graphs correspond to row permutations consistent with matchings in G_b (see an example in Appendix A.4).

This bipartite view provides a combinatorial handle on equivalence classes and forms the foundation for analyzing how interventions restrict the matching space.

5 Interventional Distribution

We now investigate how interventional data, when combined with observational data, provides additional knowledge about the causal structure. In particular, we show that an intervention on a single variable enables the identification of its causal parents.

¹All proofs are provided in the appendix.

Proposition 5.1. *Given the observational distribution generated by the model in (1) and the interventional distribution from model (2), the i -th row of the matrix W can be recovered. Consequently, the parents of x_i and their causal effects are identifiable.*

5.1 Recovering an Edge of the True Matching in the Bipartite Graph

Proposition 5.1 has a natural interpretation in the bipartite graph of 4.2. Recall that each perfect matching in the bipartite graph $G_b = (\{r_1, \dots, r_n\}, [n], E)$, with edges $(r_i, j) \in E$ iff $W_{\text{ICA}}[i, j] \neq 0$, corresponds to a valid permutation matrix P_π defining a distribution-equivalent graph.

Intervening on variable x_i , identifies the row corresponding to x_i . In the bipartite graph, this identifies the matching edge $(r_{\pi^{-1}(i)}, i)$ corresponding to the true permutation. This eliminates ambiguity for row $r_{\pi^{-1}(i)}$ and restricts the possible options for matching the remaining vertices in the bipartite graph (see an example in Appendix B.2).

6 Adaptive Experiment Design

In previous sections, we showed that observational data alone allows us to recover the causal structure only up to an equivalence class of graphs, denoted by \mathcal{G}_{eq} . Each member of this class corresponds to a different row permutation of the matrix $I - W_{\text{ICA}}$, which is itself associated with a valid matching in the bipartite graph introduced in Section 4.2.

To uniquely identify the causal structure, we need to perform interventions. As demonstrated in Proposition 5.1, an intervention on variable x_i reveals the row corresponding to x_i , thereby identifies the correct matching edge $(r_{\pi^{-1}(i)}, i)$ in the bipartite graph. Each such intervention removes all graphs from \mathcal{G}_{eq} that are incompatible with the identified edge, shrinking the equivalence class.

We now define an **adaptive** experiment design framework tailored to this setting. Our objective is to strategically select a sequence of intervention targets to maximally reduce the size of the equivalence class \mathcal{G}_{eq} , with the goal of identifying the true causal graph G^* using as few interventions as possible.

Let $\mathcal{I} = \{i_1, i_2, \dots, i_K\} \subseteq [n]$ be the set of selected variables for intervention. For each $i_j \in \mathcal{I}$, we collect the interventional data from the distribution $\mathbf{x}^{(i_j)} \sim P(\mathbf{x}|\text{do}(x_{i_j}))$, and perform ICA to estimate the corresponding interventional matrix $I - W^{(i_j)}$. By comparing this estimate with the observational estimate $I - W_{\text{ICA}}$, we recover the row of x_{i_j} and hence the corresponding matching edge. We then eliminate all graphs in \mathcal{G}_{eq} that are incompatible with this edge.

Let $\Omega := \mathcal{G}_{\text{eq}}$ denote the initial equivalence class. For a given intervention set $\mathcal{I} \subseteq [n]$ and a true graph $\phi \in \Omega$, let $\Omega^{(\mathcal{I}, \phi)} \subseteq \Omega$ denote the set of all graphs that are compatible with the result of interventions on variables in \mathcal{I} . The objective is to find the set \mathcal{I} that minimizes the size of the reduced equivalence class $|\Omega^{(\mathcal{I}, \phi)}|$. In the absence of a prior on the true graph $\phi \in \Omega$, we consider a uniform prior over Ω . This leads to an expected-size criterion for selecting the next intervention.

6.1 Adaptive Stochastic Optimization

Our aim is to eliminate as many candidate graphs as possible *adaptively*, by choosing each intervention based on the outcomes observed so far. Let the unknown true graph be a random variable drawn uniformly from the set of observationally equivalent graphs, i.e., $\Phi \sim \text{Unif}(\Omega)$.

Following the *Adaptive Stochastic Optimization* framework [11], we define:

- A *policy* π is an adaptive strategy that, at each step t , selects an intervention $I_t = \pi((i_1, O_1), \dots, (i_{t-1}, O_{t-1}))$, based on the history of past interventions i_1, \dots, i_{t-1} and their observed outcomes O_1, \dots, O_{t-1} . Here, each outcome O_s corresponds to the information gained from intervening on variable i_s —for example, the index $\pi^{-1}(i_s)$ of the true row recovered via ICA, or equivalently, the bipartite matching edge $(r_{\pi^{-1}(i_s)}, i_s)$ that was identified.
- For a given realization $\phi \in \Omega$ and policy π , let $\mathcal{I}(\pi, \phi) = \{i_1, i_2, \dots, i_K\}$ be the (random) set of variables on which π intervenes before terminating (subject to a budget of K interventions).

- Define the *reward function* as

$$f(\mathcal{I}, \phi) := |\Omega| - |\Omega^{(\mathcal{I}, \phi)}|, \quad (3)$$

where $\mathcal{I} \subseteq [n]$ denotes the set of intervened variables, and $\Omega^{(\mathcal{I}, \phi)} \subseteq \Omega$ is the set of graphs that agree with ϕ on the outcomes of all interventions in \mathcal{I} .

The *utility* of a policy π is the expected reward:

$$F(\pi) = \mathbb{E}_{\Phi \sim \text{Unif}(\Omega)} [f(\mathcal{I}(\pi, \Phi), \Phi)]. \quad (4)$$

Our goal is to find an adaptive policy π , subject to a budget of K interventions, that maximizes $F(\pi)$.

6.2 Adaptive Monotonicity and Submodularity

We now review some relevant definitions from the adaptive submodularity framework [11], and show that our reward function satisfies these properties.

Definition 6.1 (Universe and Random Realization). Let Ω be the finite set of all possible true graphs (realizations) consistent with the observational data. We treat the true graph as a random variable $\Phi \sim \text{Unif}(\Omega)$.

Definition 6.2 (Partial Realization). Let $\mathcal{E} = \{1, 2, \dots, n\}$ denote the set of variables eligible for intervention, which, in most cases, coincides with the full set of observed variables, and let O denote the set of possible outcomes from each intervention. A *partial realization* ψ is a function

$$\psi : \text{dom}(\psi) \rightarrow O, \quad \text{dom}(\psi) \subseteq \mathcal{E},$$

where $\psi(i) \in O$ records the observed outcome (e.g., recovered matching edge) for intervention i . We write $\Phi \sim \psi$ to denote the posterior distribution over realizations conditioned on consistency with all observations in ψ , i.e.,

$$\Pr[\Phi = \phi | \Phi \sim \psi] \propto \begin{cases} 1 & \text{if } \phi \text{ agrees with } \psi, \\ 0 & \text{otherwise.} \end{cases}$$

Definition 6.3 (Conditional Expected Marginal Benefit). Let $f : 2^{\mathcal{E}} \times \Omega \rightarrow \mathbb{R}_{\geq 0}$ be a reward function and let ψ be a partial realization. For any $v \notin \text{dom}(\psi)$, define the *conditional expected marginal benefit* as

$$\Delta(v|\psi) := \mathbb{E} [f(\text{dom}(\psi) \cup \{v\}, \Phi) - f(\text{dom}(\psi), \Phi) | \Phi \sim \psi]. \quad (5)$$

Definition 6.4 (Adaptive Monotonicity). A reward function f is *adaptive monotone* if

$$\Delta(v|\psi) \geq 0 \quad (6)$$

for every partial realization ψ and item $v \notin \text{dom}(\psi)$.

Definition 6.5 (Adaptive Submodularity). A reward function f is *adaptive submodular* if, for all partial realizations $\psi \subseteq \psi'$ and all $v \notin \text{dom}(\psi')$,

$$\Delta(v|\psi) \geq \Delta(v|\psi'). \quad (7)$$

Theorem 6.6. Let $f(\mathcal{I}, \phi)$ be the number of graphs eliminated after performing interventions \mathcal{I} under realization ϕ , as defined in Equation (3). Then f is adaptive monotone and adaptive submodular with respect to the uniform prior distribution over realizations.

Let $F(\pi) = \mathbb{E}_{\Phi \sim \text{Unif}(\Omega)} [f(\mathcal{I}(\pi, \Phi), \Phi)]$ denote the expected number of eliminated graphs under an adaptive policy π , as defined in Equation (4). Then the adaptive greedy policy achieves a $(1 - 1/e)$ -approximation to the optimal value of F .

7 Estimating the Reward Function by Sampling

Based on Theorem 6.6, at each step (after observing the results of previous interventions), we greedily intervene on the variable that maximally reduces the size of the remaining equivalence class. Thus, it is sufficient to compute the conditional expected marginal benefit $\Delta(v|\psi)$ for each candidate variable $v \notin \text{dom}(\psi)$.

Let ϕ be any realization consistent with ψ , and define $N = |\Omega^{(\text{dom}(\psi), \phi)}|$. In the bipartite-matching representation of the problem, the column vertex v in the bipartite graph may be matched to different rows across the candidate graphs in $\Omega^{(\text{dom}(\psi), \phi)}$. Concretely, recall that each graph in this equivalence class corresponds to a perfect matching in the bipartite graph G_b , whose left nodes are rows r_1, \dots, r_n , and whose right nodes are column indices $1, \dots, n$. Before intervening on variable v , the right-hand node v may be matched to different row nodes r_{z_1}, \dots, r_{z_k} across the N graphs. Denote these distinct row indices by z_1, z_2, \dots, z_k . Therefore, the set of candidate edges for node v is

$$\{(r_{z_i}, v) : i = 1, \dots, k\}.$$

Each graph $g \in \Omega^{(\text{dom}(\psi), \phi)}$ selects exactly one of these edges in its matching. Let

$$n_i = |\{g \in \Omega^{(\text{dom}(\psi), \phi)} : g \text{ matches } v \text{ to } r_{z_i}\}|, \quad \sum_{i=1}^k n_i = N.$$

Therefore, prior to an intervention, the probability that v is matched to row r_{z_i} is $p_i := n_i/N$. Once we intervene on v and recover its true row, we identify the correct edge (r_{z_i}, v) and eliminate all graphs that do not contain this edge. If edge (r_{z_i}, v) is observed, exactly n_i graphs survive, and the number of eliminated graphs is $N - n_i$. Hence, the conditional expected marginal benefit under ψ is:

$$\Delta(v|\psi) = \sum_{i=1}^k p_i(N - n_i) = N \sum_{i=1}^k p_i(1 - p_i). \quad (8)$$

To compute $\Delta(v|\psi)$ via Equation (8), we need to know the full equivalence class $\Omega^{(\text{dom}(\psi), \phi)}$ (This is necessary for evaluating n_i and N). Since explicitly enumerating all graphs in this class is computationally infeasible due to the exponential number of perfect matchings, we estimate $\Delta(v|\psi)$ by sampling from $\Omega^{(\text{dom}(\psi), \phi)}$.

Let $\mathbf{p} = [p_1, \dots, p_k]^T$ and define the normalized marginal benefit as $L(\mathbf{p}) := \sum_{i=1}^k p_i(1 - p_i)$. Although \mathbf{p} is not directly accessible, we assume a **sampling oracle** \mathcal{S} that returns i.i.d. samples from \mathbf{p} . This corresponds to sampling a perfect matching uniformly at random and observing the edge incident to v . A single batch of such samples allows us to estimate $L(\mathbf{p})$ for all intervention candidates.

The existence of a polynomial-time approximate uniform sampler is ensured by the FPRAS of Jerrum, Sinclair, and Vigoda [16]. While their method offers theoretical guarantees, we also consider faster heuristic alternatives in practice. For the sake of analysis, it suffices to assume access to such a sampler.

Given samples $X_1, \dots, X_M \sim \mathbf{p}$, where each $X_j \in [k]$ indicates the row matched to v , we construct an empirical estimate $\hat{\mathbf{p}}$ and use it to approximate $\Delta(v|\psi)$.

7.1 Statistical Accuracy of Sampling-Based Estimation

We now quantify the accuracy of estimating the marginal benefit $\Delta(v|\psi) = NL(\mathbf{p}) = N \sum_{i=1}^k p_i(1 - p_i)$ via empirical sampling.

Let $\mathbf{p} = (p_1, \dots, p_k)$ be the unknown categorical distribution over candidate rows (i.e., the edges connected to v in the bipartite graph). Let $X_1, \dots, X_M \stackrel{\text{iid}}{\sim} \mathbf{p}$, and define the empirical distribution:

$$\hat{\mathbf{p}} = (\hat{p}_1, \dots, \hat{p}_k), \quad \hat{p}_i := \frac{1}{M} \sum_{j=1}^M \mathbf{1}\{X_j = i\}. \quad (9)$$

We use the estimator $L(\hat{\mathbf{p}}) = \sum_{i=1}^k \hat{p}_i(1 - \hat{p}_i)$. The following theorem provides an estimation error bound on $|L(\hat{\mathbf{p}}) - L(\mathbf{p})|$.

Theorem 7.1. *Given M i.i.d samples $X_1, \dots, X_M \stackrel{\text{iid}}{\sim} \mathbf{p}$, for the estimator $L(\hat{\mathbf{p}})$, we have:*

$$\mathbb{P} \left[|L(\hat{\mathbf{p}}) - L(\mathbf{p})| \geq \frac{1}{M} + \sqrt{\frac{2}{M} \log \left(\frac{2}{\epsilon} \right)} \right] \leq \epsilon. \quad (10)$$

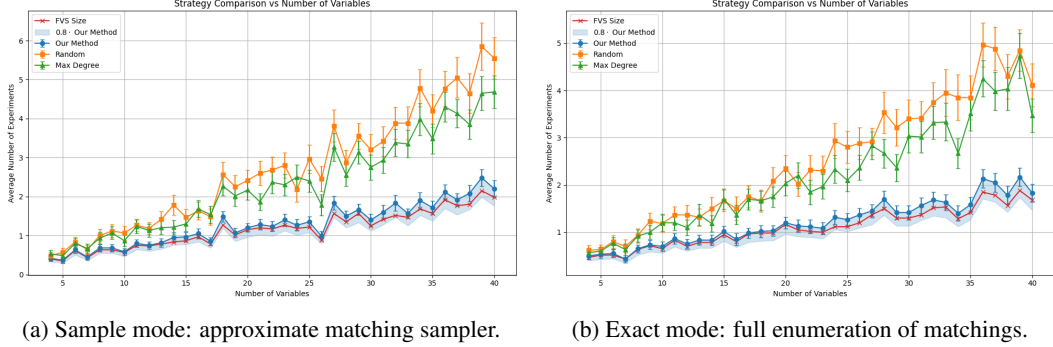


Figure 2: Comparison of intervention strategies assuming an ideal ICA oracle. Our method (Adaptive) consistently performs close to the feedback vertex set (FVS) lower bound.

In particular, picking $\epsilon = \frac{1}{M}$ implies that with probability at least $1 - \frac{1}{M}$, the error satisfies

$$|L(\hat{\mathbf{p}}) - L(\mathbf{p})| \leq \frac{1}{M} + \sqrt{\frac{2 \log(2M)}{M}} = \mathcal{O}\left(\sqrt{\frac{\log M}{M}}\right) = \tilde{\mathcal{O}}\left(\sqrt{\frac{1}{M}}\right).$$

8 Experimental Results

We evaluate the performance of our adaptive experiment design method for identifying causal graphs under the linear non-Gaussian structural causal model assumption. Experiments are conducted on synthetic data generated from Erdős–Rényi random directed graphs, where each possible directed edge (excluding self-loops) is included independently with a fixed probability. To ensure identifiability, we enforce that the resulting matrix $I - W$ is invertible. We measure the average number of interventions required for full causal identification.

8.1 Performance with an Ideal ICA Oracle

To evaluate the intervention strategy on its own, isolating it from statistical estimation errors, we assume access to an ideal ICA procedure that returns a row-permuted and scaled version of the true matrix $I - W$. The adaptive method uses a bipartite representation of the graph and samples perfect matchings using two modes: *exact*, where all matchings are enumerated, and *sample*, where a fast greedy heuristic is used (E.2). We compare our adaptive strategy against two baselines: *Random*, which selects a target uniformly at random, and *Max Degree*, which chooses the node with the highest degree in the bipartite graph.

As a theoretical baseline, we compute the *feedback vertex set (FVS)* of each true graph. The FVS size represents a fundamental lower bound on the number of interventions required to break all cycles. Although computing the FVS is NP-hard, we solve it exactly to obtain the best possible benchmark.

Figures 2a and 2b show that our adaptive method consistently outperforms the baselines and operates remarkably close to the intractable FVS lower bound, despite having no knowledge of the true graph structure. This demonstrates the near-optimality of our greedy approach in an ideal setting. For more details, please refer to E.

8.2 Robustness in a Practical Setting with Finite-Sample ICA

To assess the practical viability of our method, we evaluate its performance in a more realistic setting without an ideal ICA oracle. Instead, we apply the FastICA algorithm [15] to finite samples generated from both observational and interventional distributions. This introduces estimation noise, which can corrupt the recovered matrices. To handle these challenges, we introduce two key algorithmic modifications for robustness: *adaptive thresholding* of matrix entries and a *safe matching procedure* to prevent incorrect row assignments. A detailed description of these modifications is provided in Appendix E.3.

In this setting, our adaptive strategy continues to exhibit a performance trend similar to the ideal case, consistently outperforming the Random and Max Degree baselines and remaining remarkably close to the FVS lower bound (see Figure 6 in Appendix E.3), confirming its effectiveness even under estimation noise of ICA.

Furthermore, to analyze the quality of the recovered structure, we measured the relative error ($\varepsilon_{\text{rel}} = \frac{\|\widehat{W} - W^*\|_F}{\|W^*\|_F}$) between the estimated matrix \widehat{W} and the ground-truth W^* . Our result, presented in Appendix E.3, shows that the algorithm achieves a high-fidelity recovery of the true causal structure in the vast majority of runs.

9 Extensions

Our core framework focuses on adaptive, single-variable, perfect interventions in linear non-Gaussian models. Here, we briefly discuss how the proposed method can be extended to more general settings. We defer detailed derivations and proofs to Appendix F.

9.1 Multi-Node Interventions

The proposed method can be generalized to handle simultaneous interventions on a set of variables $E = \{i_1, \dots, i_t\}$. Such an intervention is always informative, as it localizes the row-permutation ambiguity to the intervened set E and effectively disambiguates it from the rest of the graph. Furthermore, if the subgraph induced by the variables in E is acyclic, the internal cycle reversion ambiguity within E is fully resolved. This allows for the unique recovery of the true row for each variable in E , enabling the safe parallelization of experiments to accelerate the causal discovery process. We provide a detailed analysis in Appendix F.1.

9.2 Imperfect Interventions

Our approach is also robust to imperfect interventions. The core requirement for identifying a variable’s corresponding row is that the intervention sufficiently perturbs its causal mechanism. As long as the post-intervention row is distinguishable from the set of observational rows after ICA estimation, the intervention is informative. This condition holds for some realistic noisy or incomplete intervention types, enhancing the practical applicability of our method. Further details are discussed in Appendix F.2.

9.3 Generalization to Non-linear Models

While we focused on the linear case, our experimental design strategy could be extended to non-linear structural causal models. Recent advances in non-linear ICA can recover the Jacobian of the model’s inverse function, whose sparsity pattern reveals the underlying causal graph structure up to the same permutation ambiguity found in the linear setting [27]. Our bipartite matching formulation and reward function could then be applied to resolve this ambiguity. We outline this potential generalization in Appendix F.3.

10 Conclusion

We introduced a framework for causal structure learning in linear non-Gaussian models with cycles, leveraging a combinatorial characterization of equivalence classes via bipartite matchings. By formalizing experiment design as an adaptive submodular optimization problem, we developed a near-optimal greedy policy that incrementally resolves causal ambiguity through targeted interventions. Our sampling-based estimator enables practical implementation without exhaustive enumeration. Empirical results confirm that our method recovers the true graph with few interventions, often matching the feedback vertex set lower bound, despite having no access to the true structure.

References

- [1] Raj Agrawal, Chandler Squires, Karren Yang, Karthikeyan Shanmugam, and Caroline Uhler. Abcd-strategy: Budgeted experimental design for targeted causal structure discovery. In

- International Conference on Artificial Intelligence and Statistics*, pages 3400–3409. PMLR, 2019.
- [2] Ali AhmadiTeshnizi, Saber Salehkaleybar, and Negar Kiyavash. Lazyiter: a fast algorithm for counting markov equivalent dags and designing experiments. In *International Conference on Machine Learning*, pages 125–133. PMLR, 2020.
 - [3] 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.
 - [4] Thomas H Cormen, Charles E Leiserson, Ronald L Rivest, and Clifford Stein. *Introduction to algorithms*. MIT press, 2022.
 - [5] Frederick Eberhardt, Clark Glymour, and Richard Scheines. On the number of experiments sufficient and in the worst case necessary to identify all causal relations among n variables. In *Conference on Uncertainty and Artificial Intelligence*, pages 178–184, 2005.
 - [6] Patrick Forré and Joris M Mooij. Constraint-based causal discovery for non-linear structural causal models with cycles and latent confounders. In *Conference on Uncertainty in Artificial Intelligence*, 2018.
 - [7] AmirEmad Ghassami, Saber Salehkaleybar, and Negar Kiyavash. Interventional experiment design for causal structure learning. *arXiv preprint arXiv:1910.05651*, 2019.
 - [8] AmirEmad Ghassami, Saber Salehkaleybar, Negar Kiyavash, and Elias Bareinboim. Budgeted experiment design for causal structure learning. In *International Conference on Machine Learning*, pages 1724–1733. PMLR, 2018.
 - [9] AmirEmad Ghassami, Saber Salehkaleybar, Negar Kiyavash, and Kun Zhang. Counting and sampling from markov equivalent dags using clique trees. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 3664–3671, 2019.
 - [10] AmirEmad Ghassami, Alan Yang, Negar Kiyavash, and Kun Zhang. Characterizing distribution equivalence and structure learning for cyclic and acyclic directed graphs. In *International Conference on Machine Learning*, pages 3494–3504. PMLR, 2020.
 - [11] Daniel Golovin and Andreas Krause. Adaptive submodularity: Theory and applications in active learning and stochastic optimization. *Journal of Artificial Intelligence Research*, 42:427–486, 2011.
 - [12] Kristjan Greenewald, Dmitriy Katz, Karthikeyan Shanmugam, Sara Magliacane, Murat Kocaoglu, Enric Boix Adsera, and Guy Bresler. Sample efficient active learning of causal trees. *Advances in Neural Information Processing Systems*, 32, 2019.
 - [13] Alain Hauser and Peter Bühlmann. Two optimal strategies for active learning of causal models from interventional data. *International Journal of Approximate Reasoning*, 55(4):926–939, 2014.
 - [14] Yang-Bo He and Zhi Geng. Active learning of causal networks with intervention experiments and optimal designs. *Journal of Machine Learning Research*, 9(Nov):2523–2547, 2008.
 - [15] Aapo Hyvärinen and Erkki Oja. Independent component analysis: algorithms and applications. *Neural networks*, 13(4-5):411–430, 2000.
 - [16] Mark Jerrum, Alistair Sinclair, and Eric Vigoda. A polynomial-time approximation algorithm for the permanent of a matrix with nonnegative entries. *Journal of the ACM (JACM)*, 51(4):671–697, 2004.
 - [17] Ilyes Khemakhem, Diederik Kingma, Ricardo Monti, and Aapo Hyvarinen. Variational autoencoders and nonlinear ica: A unifying framework. In *International conference on artificial intelligence and statistics*, pages 2207–2217. PMLR, 2020.
 - [18] Bohdan Kivva, Goutham Rajendran, Pradeep Ravikumar, and Bryon Aragam. Identifiability of deep generative models without auxiliary information. *Advances in Neural Information Processing Systems*, 35:15687–15701, 2022.

- [19] Murat Kocaoglu, Alex Dimakis, and Sriram Vishwanath. Cost-optimal learning of causal graphs. In *International Conference on Machine Learning*, pages 1875–1884. PMLR, 2017.
- [20] Murat Kocaoglu, Karthikeyan Shanmugam, and Elias Bareinboim. Experimental design for learning causal graphs with latent variables. *Advances in Neural Information Processing Systems*, 30, 2017.
- [21] Gustavo Lacerda, Peter L Spirtes, Joseph Ramsey, and Patrik O Hoyer. Discovering cyclic causal models by independent components analysis. In *Conference on Uncertainty in Artificial Intelligence*, pages 366–374, 2008.
- [22] Gustavo Lacerda, Peter L Spirtes, Joseph Ramsey, and Patrik O Hoyer. Discovering cyclic causal models by independent components analysis. *arXiv preprint arXiv:1206.3273*, 2012.
- [23] Ehsan Mokhtarian, Saber Salehkaleybar, AmirEmad Ghassami, and Negar Kiyavash. A unified experiment design approach for cyclic and acyclic causal models. *Journal of Machine Learning Research*, 24(354):1–31, 2023.
- [24] Joris M Mooij and Tom Claassen. Constraint-based causal discovery using partial ancestral graphs in the presence of cycles. In *Conference on Uncertainty in Artificial Intelligence*, pages 1159–1168. PMLR, 2020.
- [25] Judea Pearl. *Causality*. Cambridge university press, 2009.
- [26] Jonas Peters, Dominik Janzing, and Bernhard Schölkopf. *Elements of causal inference: foundations and learning algorithms*. The MIT Press, 2017.
- [27] Patrik Reizinger, Yash Sharma, Matthias Bethge, Bernhard Schölkopf, Ferenc Huszár, and Wieland Brendel. Jacobian-based causal discovery with nonlinear ica. *Transactions on Machine Learning Research*, 2023.
- [28] Thomas S Richardson. A discovery algorithm for directed cyclic graphs. *arXiv preprint arXiv:1302.3599*, 2013.
- [29] Thomas S Richardson. A polynomial-time algorithm for deciding markov equivalence of directed cyclic graphical models. *arXiv preprint arXiv:1302.3600*, 2013.
- [30] Karthikeyan Shanmugam, Murat Kocaoglu, Alexandros G Dimakis, and Sriram Vishwanath. Learning causal graphs with small interventions. *Advances in Neural Information Processing Systems*, 28, 2015.
- [31] Shohei Shimizu, Patrik O Hoyer, Aapo Hyvärinen, Antti Kerminen, and Michael Jordan. A linear non-gaussian acyclic model for causal discovery. *Journal of Machine Learning Research*, 7(10), 2006.
- [32] Peter Spirtes, Clark N Glymour, and Richard Scheines. *Causation, prediction, and search*. MIT press, 2000.
- [33] Chandler Squires, Sara Magliacane, Kristjan Greenewald, Dmitriy Katz, Murat Kocaoglu, and Karthikeyan Shanmugam. Active structure learning of causal dags via directed clique trees. *Advances in Neural Information Processing Systems*, 33:21500–21511, 2020.
- [34] Panagiotis Tigas, Yashas Annadani, Andrew Jesson, Bernhard Schölkopf, Yarin Gal, and Stefan Bauer. Interventions, where and how? experimental design for causal models at scale. *Advances in Neural Information Processing Systems*, 35:24130–24143, 2022.
- [35] Roman Vershynin. *High-dimensional probability: An introduction with applications in data science*, volume 47. Cambridge university press, 2018.
- [36] Marcel Wienöbst, Max Bannach, and Maciej Liskiewicz. Polynomial-time algorithms for counting and sampling markov equivalent dags. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 12198–12206, 2021.
- [37] Marcel Wienöbst, Max Bannach, and Maciej Liśkiewicz. Polynomial-time algorithms for counting and sampling markov equivalent dags with applications. *Journal of Machine Learning Research*, 24(213):1–45, 2023.

Supplementary Material

Contents

A	Distribution Equivalence: Formal Definitions, Characterization, and Proofs	14
A.1	Cycle Decomposition and Structural Characterization	14
A.2	Proofs of Section A	15
A.3	Proof of Theorem 4.4	16
A.4	Example: Matchings in Bipartite Graphs	17
B	Proofs and Examples of Section 5	17
B.1	Proof of Proposition 5.1	17
B.2	Example: How Interventions Reveal True Matching Edges	17
C	Proofs of Section 6	18
C.1	Proof of Theorem 6.6	18
D	Concentration Bounds for the Normalized Marginal Benefit	22
D.1	Bounding the Bias Term	22
D.2	Bounding the Deviation Term	23
D.3	Proof of Theorem 7.1	24
E	Implementation and Experimental Details	24
E.1	Pseudocode for the Idealized Setting	24
E.2	Experimental Details for the Idealized Setting	25
E.3	Implementation in the Finite-Sample Setting	26
F	Details on Extensions	28
F.1	Multi-Node Interventions	28
F.2	Imperfect and Noisy Interventions	29
F.3	Generalization to Non-linear Models	29

A Distribution Equivalence: Formal Definitions, Characterization, and Proofs

This appendix provides the formal definitions and results omitted from the main text for brevity. We define two key notions: distribution-entailment equivalence (for parameterized LSCMs) and distribution equivalence (for graphical structures), and show that they are equivalent under the linear non-Gaussian SCM assumption.

Definition A.1 (Distribution Set). Let G be a directed graph (without self-loops) with binary adjacency matrix $B_G \in \{0, 1\}^{n \times n}$. The *distribution set* $\mathcal{P}(G)$ is the collection of all observational distributions $P_{\mathbf{x}}$ that can arise from any linear SCM consistent with G :

$$\mathcal{P}(G) := \left\{ P_{\mathbf{x}} = (I - T_W)^{-1}_{\#}(P_{\mathbf{e}}) \mid \text{supp}(W) = \text{supp}(B_G), P_{\mathbf{e}} \in \mathcal{P}(\mathcal{E}) \right\},$$

where T_W denotes the linear operator defined by the matrix W , and $\mathcal{P}(\mathcal{E})$ is the class of exogenous distributions defined in Assumption 3.3.

Definition A.2 (Distribution Equivalence). Two directed **graphs** G_1 and G_2 are said to be *distribution equivalent*, denoted by $G_1 \equiv G_2$, if $\mathcal{P}(G_1) = \mathcal{P}(G_2)$.

Now consider two *parameterized* LSCMs,

$$\mathbf{x} = (I - W_1)^{-1} \mathbf{e}_1 \quad \text{and} \quad \mathbf{x} = (I - W_2)^{-1} \mathbf{e}_2,$$

with $P_{\mathbf{e}_1}, P_{\mathbf{e}_2} \in \mathcal{P}(\mathcal{E})$, that induce the same observational distribution. By ICA identifiability, there are some permutation matrix P and diagonal scaling matrix D such that:

$$I - W_2 = PD(I - W_1), \quad \mathbf{e}_2 := PDe_1. \quad (11)$$

In this case, the two **models** are said to be *distribution-entailment equivalent*—they have different structure matrices but yield the same distribution over observables.

In the literature, these two concepts are typically treated as distinct. However, in our linear non-Gaussian setting, we show that they coincide. To establish this, we proceed with an algebraic characterization of distribution equivalence.

Proposition A.3. *Two directed graphs (without self-loops) G_1 and G_2 are distribution equivalent, that is, $G_1 \equiv G_2$, if and only if there exists a permutation matrix $P \in \{0, 1\}^{n \times n}$ such that*

$$I + B_{G_2} = P(I + B_{G_1}). \quad (12)$$

This result implies that if two LSCMs are distribution-entailment equivalent, then their associated graphs are distribution equivalent (see the proof A.2.1). Conversely, all models consistent with the same observational distribution must correspond to graphs in the same distribution equivalence class. For linear SCMs with non-Gaussian noise, the two notions of equivalence are therefore interchangeable.

To understand how graphs within the same equivalence class differ, we introduce the cycle decomposition of permutations and relate it to structural changes in the graph. While this idea is discussed in [22], we include it here for completeness—especially since their treatment is limited to distribution-entailment equivalent LSCMs, rather than distribution-equivalent graphs.

A.1 Cycle Decomposition and Structural Characterization

Let $\sigma : [n] \rightarrow [n]$ be a permutation. A *cycle decomposition* of σ expresses it as a product of disjoint cycles. A cycle $(i_1, i_2, \dots, i_\ell)$ maps $i_j \mapsto i_{j+1}$ for $1 \leq j < \ell$, and $i_\ell \mapsto i_1$. Two cycles are disjoint if they share no elements. Every permutation has a unique cycle decomposition (up to cycle order).

Based on Proposition A.3, if $G_1 \equiv G_2$, then there exists a permutation P satisfying $I + B_{G_2} = P(I + B_{G_1})$. The next proposition relates this to cycles in the graph.

Proposition A.4. *Let $G_1 \equiv G_2$. Let P_π be the permutation matrix such that $I + B_{G_2} = P_\pi(I + B_{G_1})$. Suppose $\pi = c_1 c_2 \dots c_k$ is its cycle decomposition, where $c_j = (i_1^j, i_2^j, \dots, i_{n_j}^j)$ and $n_j \geq 2$. Then for each $j \in [k]$,*

- $(x_{i_1^j} \rightarrow x_{i_2^j} \rightarrow \cdots \rightarrow x_{i_{n_j}^j} \rightarrow x_{i_1^j})$ forms a cycle in G_1 , and
- $(x_{i_1^j} \leftarrow x_{i_2^j} \leftarrow \cdots \leftarrow x_{i_{n_j}^j} \leftarrow x_{i_1^j})$ forms a cycle in G_2 .

Remark. Proposition A.4 highlights that each cycle in the permutation π corresponds to a cycle in one graph whose direction is reversed in the other. However, this characterization does not fully capture the structural relationship between distribution-equivalent graphs. In the main text, we define the *Cycle Reversion* (CR) operation 4.1, which provides a more complete description. Unlike the statement of the proposition, the CR operation explicitly specifies how incoming edges from nodes outside the cycle are redirected after reversion. This reveals that the permutation π encodes not only reversed cycles, but also a deeper structural transformation of the graph that preserves distribution equivalence.

A.2 Proofs of Section A

A.2.1 Proof of Proposition A.3

We first show that $G_1 \equiv G_2$ implies the existence of such a permutation matrix P . Suppose $\mathcal{P}(G_1) = \mathcal{P}(G_2) = \mathcal{P}$, i.e., both graphs admit the same set of observational distributions. Take an arbitrary $P_{\mathbf{x}} \in \mathcal{P}$, then by definition:

$\mathbf{x} = (I - W_1)^{-1} \mathbf{e}_1$ with $\text{supp}(W_1) = \text{supp}(B_{G_1})$, $P_{\mathbf{e}_1} \in \mathcal{P}(\mathcal{E})$, and also, $\mathbf{x} = (I - W_2)^{-1} \mathbf{e}_2$ with $\text{supp}(W_2) = \text{supp}(B_{G_2})$, $P_{\mathbf{e}_2} \in \mathcal{P}(\mathcal{E})$.

From the ICA identifiability result, since \mathbf{x} admits two LSCM representations with non-Gaussian independent noises, there must exist a permutation matrix P and a diagonal scaling matrix D $I - W_2 = PD(I - W_1)$ and $\mathbf{e}_2 = PDe_1$.

Now, since both graphs lack self-loops, the diagonal entries of $I - W_2$ must be 1. We expand this constraint. Let w_i^1 be the i -th row of W_1 , and $u_j \in \mathbb{R}^n$ be the standard basis row vector with 1 in position j . Then:

$$P = \begin{bmatrix} u_{\pi(1)} \\ u_{\pi(2)} \\ \vdots \\ u_{\pi(n)} \end{bmatrix}, \quad D = \text{diag}(d_1, d_2, \dots, d_n),$$

where π is the permutation induced by P . Therefore, the matrix product becomes:

$$PD(I - W_1) = \begin{bmatrix} d_{\pi(1)}(u_{\pi(1)} - w_{\pi(1)}^1) \\ d_{\pi(2)}(u_{\pi(2)} - w_{\pi(2)}^1) \\ \vdots \\ d_{\pi(n)}(u_{\pi(n)} - w_{\pi(n)}^1) \end{bmatrix}. \quad (13)$$

To ensure the diagonal entries of $I - W_2 = PD(I - W_1)$ are 1, we must have, for all $i \in [n]$,

$$d_i = \frac{1}{u_{i, \pi^{-1}(i)} - w_{i, \pi^{-1}(i)}^1} = \begin{cases} 1 & \text{if } \pi(i) = i, \\ \frac{-1}{w_{i, \pi^{-1}(i)}^1} & \text{if } \pi(i) \neq i. \end{cases} \quad (14)$$

Hence, D is uniquely determined by P and W_1 . Since $P_{\mathbf{x}}$ is arbitrary in \mathcal{P} , we can select a distribution generated by $W_1 = -B_{G_1}$. With this choice we get $D = I$, yielding $I - W_2 = P(I + B_{G_1})$. Since $\text{supp}(W_2) = \text{supp}(B_{G_2})$, we conclude that $W_2 = -B_{G_2}$.

Now, we will prove the other direction. For every $P_{\mathbf{x}} \in \mathcal{P}(G_1)$, there is some W_1 and $P_{\mathbf{e}_1} \in \mathcal{P}(\mathcal{E})$ such that $\mathbf{x} = (I - W_1)^{-1} \mathbf{e}_1$ and $\text{supp}(W_1) = \text{supp}(B_{G_1})$. Define the scaling matrix D as defined in (14) based on the given P in the assumption and W_1 . Then define $W_2 = I - PD(I - W_1)$ and $\mathbf{e}_2 = PDe_1$. Applying a scaling and a permutation matrix to an arbitrary matrix does not add or remove zero entries. Therefore, it can be shown that $\text{supp}(W_2) = \text{supp}(B_{G_2})$. This results in $P_{\mathbf{x}} \in \mathcal{P}(G_2)$ and then $\mathcal{P}(G_1) \subseteq \mathcal{P}(G_2)$. With the same arguments and substituting P by P^T , we have $\mathcal{P}(G_2) \subseteq \mathcal{P}(G_1)$ and therefore, $G_1 \equiv G_2$. \square

A.2.2 Proof of Proposition A.4

Let b_i denote the i -th row of B_{G_1} , and let u_j be the standard basis row vector in \mathbb{R}^n with 1 at position j and 0 elsewhere. As in the proof of Proposition A.3, the fact that all diagonal entries of $P_\pi(I + B_{G_1})$ equal 1 implies that $u_{\pi(i),i} + b_{\pi(i),i} = 1$ for all $i \in [n]$.

If i is a fixed point, then $\pi(i) = i$ and the condition becomes $u_{i,i} + b_{i,i} = 1 + 0 = 1$, since G_1 has no self-loops. Now, let i belong to a cycle $c_j = (i_1^j, i_2^j, \dots, i_{n_j}^j)$ of length $n_j \geq 2$. Then, applying the condition to each i_r^j gives $b_{i_{r+1}^j, i_r^j} = 1$ for all $r = 1, \dots, n_j$ (with indices modulo n_j). Thus, the sequence of edges $(x_{i_1^j} \rightarrow x_{i_2^j} \rightarrow \dots \rightarrow x_{i_{n_j}^j} \rightarrow x_{i_1^j})$ forms a cycle in G_1 .

To obtain the analogous result in G_2 , observe that $I + B_{G_1} = P_{\pi^{-1}}(I + B_{G_2})$. Repeating the same argument for G_2 using $P_{\pi^{-1}}$ implies that $(x_{i_1^j} \leftarrow x_{i_2^j} \leftarrow \dots \leftarrow x_{i_{n_j}^j} \leftarrow x_{i_1^j})$ forms a cycle in G_2 . \square

A.3 Proof of Theorem 4.4

To prove this, we show that the CR operation does not move any vertex from one SCC to another. Since any two distribution-equivalent graphs can be transformed into each other through a sequence of CR operations, this implies that SCCs are preserved.

Let $C = (x_0, x_1, \dots, x_{m-1})$ be a directed cycle in a graph G , and let \mathcal{S} denote the strongly connected component (SCC) that contains all nodes of C . We apply the *Cycle Reversion (CR)* operation to C and denote the resulting graph by G' . We aim to prove that SCCs remain unchanged under this operation; specifically, no vertex in \mathcal{S} exits the component, and no external vertex enters it.

SCCs do not shrink. Take any two nodes $u, v \in \mathcal{S}$. Since \mathcal{S} is strongly connected in G , there exist directed paths $u \rightsquigarrow v$ and $v \rightsquigarrow u$ in G . If these paths do not pass through nodes in C , they remain unaffected in G' .

Now, suppose that the path from u to v passes through nodes in C . Without loss of generality, assume the path takes the form:

$$u \rightarrow p_1 \rightarrow \underbrace{x_i \rightarrow x_{i+1} \rightarrow \dots \rightarrow x_{i+k}}_{\text{subpath in the cycle } C} \rightarrow p_2 \rightarrow v,$$

where p_1 is a subpath from u to node $x_i \in C$, and p_2 is a subpath from $x_{i+k} \in C$ to v (with indices modulo m). Before cycle reversion, the edges within C follow the original cycle direction $x_0 \rightarrow x_1 \rightarrow \dots \rightarrow x_{m-1} \rightarrow x_0$. After applying CR, this direction is reversed: $x_{m-1} \rightarrow x_{m-2} \rightarrow \dots \rightarrow x_0 \rightarrow x_{m-1}$ and any edge from a node $y \notin C$ into a node $x_j \in C$ is redirected to point to x_{j-1} . Thus, p_1 , which originally led into x_i , will now lead to its predecessor x_{i-1} in the reversed cycle. So we can construct a modified path p'_1 from u to x_{i-1} . Moreover, within the reversed cycle, there exists a path from $x_{i-1} \rightarrow x_{i-2} \rightarrow \dots \rightarrow x_{i+k}$, corresponding to the reversed traversal of the original path segment in C . Since outgoing edges from nodes in C to nodes outside are preserved, p_2 remains valid from x_{i+k} to v .


Hence, the new path in G' is:

$$u \rightarrow p'_1 \rightarrow x_{i-1} \rightarrow x_{i-2} \rightarrow \dots \rightarrow x_{i+k} \rightarrow p_2 \rightarrow v.$$


By symmetric reasoning, we can similarly construct a path from v to u in G' , showing that u and v remain strongly connected. Therefore, SCCs do not shrink under CR.

SCCs do not expand. Assume for contradiction that a node $z \notin \mathcal{S}$ becomes strongly connected with all nodes in \mathcal{S} after applying CR, i.e., $z \in \mathcal{S}'$, where \mathcal{S}' is the SCC containing \mathcal{S} in G' . This would mean there exist paths $z \rightsquigarrow u$ and $u \rightsquigarrow z$ in G' for each $u \in \mathcal{S}$. Now, reapplying CR to reverse the cycle C again (returning to graph G), the same argument as above shows that SCCs cannot shrink. Thus, $z \in \mathcal{S}$ in G as well, contradicting the assumption $z \notin \mathcal{S}$.

Conclusion. Since no node leaves or enters \mathcal{S} under the CR operation, the strongly connected components remain unchanged. Figure 1b provides a visual illustration of this argument. \square



(a) Permutation as a matching.



(b) Matching after intervention.

Figure 3: (a) Each row permutation corresponds to a matching between sources and observed variables. (b) An intervention on a specific variable reveals its correct row assignment, narrowing the equivalence class.

A.4 Example: Matchings in Bipartite Graphs

As discussed in Section 4.2, each graph in a distribution equivalence class corresponds to a perfect matching in a bipartite graph defined from the matrix $I - W_{\text{ICA}}$. The following example illustrates how a specific valid row permutation corresponds to a certain perfect matching.

$$I - W_{\text{ICA}} = \begin{pmatrix} \circ & \times & \times & \times & \circ \\ \times & \circ & \circ & \times & \times \\ \circ & \times & \circ & \times & \times \\ \times & \times & \times & \circ & \circ \\ \times & \circ & \times & \circ & \times \end{pmatrix}$$

Here, cross marks indicate non-zero entries in the matrix, and the red entries mark a specific row permutation π that selects a unique non-zero element in each row—corresponding to a perfect matching. This selection defines a valid permutation matrix P_π such that $P_\pi(I - W_{\text{ICA}})$ is a valid causal model in the equivalence class with non-zero elements on the diagonal. Figure 3a further illustrates how such permutations correspond to bipartite matchings.

B Proofs and Examples of Section 5

B.1 Proof of Proposition 5.1

Applying ICA to both observational and interventional data yields $PD(I - W)$ and $P_i D_i(I - W^{(i)})$, where P, P_i are permutation matrices and D, D_i are diagonal scaling matrices.

Since W and $W^{(i)}$ differ only in the i -th row, and all rows are linearly independent (due to invertibility), their normalized forms will differ in exactly one row. This unique row in $PD(I - W)$, absent in $P_i D_i(I - W^{(i)})$, corresponds to the i -th row of W , revealing its support and weights. Thus, the valid row permutation must place this row at index i .

B.2 Example: How Interventions Reveal True Matching Edges

As discussed in Proposition 5.1, interventions provide critical information about the true matching in the bipartite representation. Recall that each perfect matching in the bipartite graph $G_b = (\{r_1, \dots, r_n\}, [n], E)$, where an edge $(r_i, j) \in E$ exists iff $W_{\text{ICA}}[i, j] \neq 0$, corresponds to a valid row permutation matrix P_π defining a distribution-equivalent graph.

When we intervene on variable x_i , the corresponding row in $I - W^{(i)}$ becomes identifiable. In the bipartite graph, this uniquely reveals the edge $(r_{\pi^{-1}(i)}, i)$, which is part of the true matching. This discovery not only pins down the correct match for one node pair but also restricts the set of consistent permutations, narrowing the equivalence class.

Example. Suppose ICA applied to observational and interventional data (intervention on x_4) yields the following recovered matrices:

$$P_1(I - W) = \begin{pmatrix} r_1^T \\ r_2^T \\ r_3^T \\ r_4^T \\ r_5^T \end{pmatrix} \quad P_2(I - W^{(4)}) = \begin{pmatrix} r_3^T \\ r_5^T \\ r_1^T \\ u_4^T \\ r_4^T \end{pmatrix}$$

Here, row r_2^T appears only in the observational matrix, while $u_4^T = [0, 0, 0, 1, 0]$ appears only in the interventional one. This mismatch allows us to identify r_2^T as the true row of x_4 , revealing the matching edge $(r_2, 4)$ in the bipartite graph. As this process is repeated across multiple variables, more edges in the true matching are discovered, eventually resolving the full permutation π . Figure 3b illustrates this idea of how such interventional data reveals some information about the true hidden perfect matching corresponding to the true causal model.

C Proofs of Section 6

C.1 Proof of Theorem 6.6

In this appendix, we provide the full proof that the reward function

$$f(\mathcal{I}, \phi) = |\Omega| - |\Omega^{(\mathcal{I}, \phi)}|$$

is both adaptive monotone and adaptive submodular.

Adaptive Monotonicity

By construction, performing an additional intervention can only eliminate more graphs (never re-introduce any). Hence, for any partial realization ψ and any $v \notin \text{dom}(\psi)$,

$$f(\text{dom}(\psi) \cup \{v\}, \phi) \geq f(\text{dom}(\psi), \phi) \quad \text{for every realization } \phi.$$

Taking expectation conditional on $\Phi \sim \psi$ preserves the inequality, establishing adaptive monotonicity (inequality (6)).

Adaptive Submodularity

Fix two partial realizations $\psi \subseteq \psi'$ with

$$\mathcal{I}_1 = \text{dom}(\psi), \quad \mathcal{I}_2 = \text{dom}(\psi'), \quad \mathcal{I}_1 \subseteq \mathcal{I}_2,$$

and let ϕ be any realization consistent with ψ' . Define

$$\Omega_1 = \Omega^{(\mathcal{I}_1, \phi)}, \quad \Omega_2 = \Omega^{(\mathcal{I}_2, \phi)},$$

and denote their sizes by

$$N_1 = |\Omega_1|, \quad N_2 = |\Omega_2|,$$

and $N_2 \leq N_1$. Let $v \notin \mathcal{I}_2$ be the next intervention candidate. In the bipartite-matching view, assume that v has k possible edges across Ω_1 . Concretely, recall that each candidate graph in our current equivalence class Ω_1 corresponds to a perfect matching in the bipartite graph G_b whose left nodes are rows r_1, \dots, r_n and whose right nodes are column positions $1, \dots, n$. Before intervening on variable v , the right-hand node v may be matched to different rows r_i across the $|\Omega_1| = N_1$ graphs. Call these distinct row-indices $i = z_1, z_2, \dots, z_k$, so that the set of “candidate edges” for node v is

$$\{(r_i, v) : i = z_1, z_2, \dots, z_k\},$$

and each graph $g \in \Omega_1$ picks exactly one of these edges in its matching (Figure 4).

Let

$$n_i = |\{g \in \Omega_1 : g \text{ matches } v \text{ to } r_{z_i}\}|, \quad \sum_{i=1}^k n_i = N_1.$$

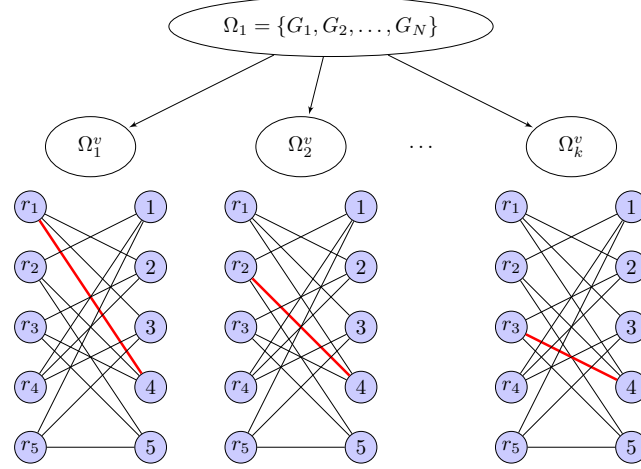


Figure 4: Illustration of equivalence class partitioning induced by a candidate intervention on variable $v = 4$. The top node represents the current equivalence class $\Omega_1 = \{G_1, \dots, G_N\}$, and each child Ω_i^v consists of graphs whose perfect matching assigns variable v to row r_{z_i} , i.e., includes the edge (r_{z_i}, v) . The bipartite graphs below depict three such subsets, each highlighting a different candidate edge (r_{z_i}, v) in red. This illustrates how an intervention on v resolves ambiguity by eliminating all but one of these subsets, reducing the equivalence class.

Hence prior to the intervention, the probability that v is matched to row r_i is $p_i = n_i/N_1$. Once we intervene on v and recover its true structural row, we immediately identify which one of the k edges (r_i, v) is correct, and eliminate all graphs in Ω_1 that did not use that edge. If edge i is observed, exactly n_i graphs survive, so the number of eliminated graphs is $N_1 - n_i$. Hence, the conditional expected marginal benefit under ψ is

$$\Delta(v|\psi) = \sum_{i=1}^k p_i (N_1 - n_i) = \sum_{i=1}^k \frac{n_i}{N_1} (N_1 - n_i) = N_1 \sum_{i=1}^k p_i (1 - p_i). \quad (15)$$

Similarly, under the realization ψ' , only the subset $\Omega_2 \subseteq \Omega_1$ remains. Let

$$m_i = |\{g \in \Omega_2 : g \text{ matches } v \text{ to } r_{z_i}\}|, \quad \sum_{i=1}^k m_i = N_2,$$

and set $q_i = m_i/N_2$. Then

$$\Delta(v|\psi') = N_2 \sum_{i=1}^k q_i (1 - q_i). \quad (16)$$

We now show

$$\Delta(v|\psi) - \Delta(v|\psi') = N_1 \sum_{i=1}^k p_i (1 - p_i) - N_2 \sum_{i=1}^k q_i (1 - q_i) \geq 0.$$

To see this, note that transition from Ω_1 to Ω_2 can be viewed as **removing** one graph in $\Omega_1 \setminus \Omega_2$ at a time. We will verify that removing a single graph from a collection of size N can only **decrease** $N \sum_i p_i (1 - p_i)$.

Removal of a single graph. Suppose we remove one graph of “type” j from a collection of size N_1 . Then

$$n'_j = n_j - 1, \quad n'_i = n_i \quad (i \neq j), \quad N' = N_1 - 1,$$

and the new probabilities are

$$p'_j = \frac{n_j - 1}{N'}, \quad p'_i = \frac{n_i}{N'} \quad (i \neq j).$$

We compare

$$M(\Omega_1) = N_1 \sum_i p_i(1 - p_i) \quad \text{and} \quad M(\Omega') = N' \sum_i p'_i(1 - p'_i).$$

Let $S = \sum_i n_i^2$. Then

$$M(\Omega_1) = N_1 \left(1 - \sum_i p_i^2\right) = N_1 - \frac{S}{N_1},$$

and

$$M(\Omega') = N' - \sum_i \frac{(n'_i)^2}{N'} = (N_1 - 1) - \frac{(S - n_j^2 + (n_j - 1)^2)}{N_1 - 1} = (N_1 - 1) - \frac{S - 2n_j + 1}{N_1 - 1}.$$

Hence

$$\begin{aligned} M(\Omega_1) - M(\Omega') &= \left(N_1 - \frac{S}{N_1}\right) - \left((N_1 - 1) - \frac{S - 2n_j + 1}{N_1 - 1}\right) \\ &= 1 - \frac{S}{N_1} + \frac{S - 2n_j + 1}{N_1 - 1} \\ &= \frac{N_1(N_1 - 1) - S(N_1 - 1) + N_1(S - 2n_j + 1)}{N_1(N_1 - 1)} \\ &= \frac{N_1^2 - 2n_j N_1 + S + (n_j^2 - n_j^2)}{N_1(N_1 - 1)} \\ &= \frac{(N_1 - n_j)^2 + \sum_{i \neq j} n_i^2}{N_1(N_1 - 1)} \geq 0. \end{aligned}$$

Thus, removing a single graph cannot increase M .

Induction to multiple removals. To reach Ω_2 of size N_2 , we remove $N_1 - N_2$ graphs one at a time. Each removal step leaves M unchanged or decreases it. Hence

$$N_1 \sum_i p_i(1 - p_i) \geq N_2 \sum_i q_i(1 - q_i) \implies \Delta(v|\psi) \geq \Delta(v|\psi'),$$

establishing adaptive submodularity (7).

By the result of Golovin and Krause [11], adaptive monotonicity and submodularity guarantee that the adaptive greedy policy achieves at least a $(1 - 1/e)$ -approximation to the optimal policy. \square

Alternative α -Parameterization Proof

An equivalent way to see that conditioning (going from ψ to ψ') only *decreases* the marginal benefit is via the following re-parameterization. We present this alternative α -parameterization proof because it is more general, as it does not rely on the discreteness of variables. Unlike previous approaches, this proof remains valid even when the variables are continuous, thereby broadening the scope of applicability to a wider range of models.

Fix a partial realization ψ so that $\Omega_1 = \Omega^{(\text{dom}(\psi), \phi)}$ has size N and assign to the next intervention v the matching edge (r_{z_j}, v) with probability

$$p_j = \frac{n_j}{N},$$

where n_j of the N surviving graphs in Ω_1 match v to row r_j . After further conditioning under ψ' , only $\Omega_2 \subseteq \Omega_1$ of size $N' \leq N$ remains, and among those $m_j \leq n_j$ graphs still use edge j . Define

$$\alpha_j = \frac{m_j}{n_j} \in [0, 1], \quad Z = \sum_{i=1}^k p_i \alpha_i.$$

Conditioned on ψ' , the probability that we observe edge j is

$$\frac{m_j}{\sum_i m_i} = \frac{\alpha_j n_j}{\sum_i \alpha_i n_i} = \frac{N \alpha_j p_j}{N \sum_i \alpha_i p_i} = \frac{\alpha_j p_j}{Z},$$

and upon observing it we eliminate $m_j = \alpha_j n_j$ graphs from $\sum_{i=1}^k m_i = ZN$ graphs. Hence the conditional expected marginal benefit under ψ' is

$$\begin{aligned}\Delta(v|\psi') &= \sum_{j=1}^k \frac{p_j \alpha_j}{Z} (NZ - n_j \alpha_j) = N \sum_{j=1}^k p_j \alpha_j \left(1 - \frac{p_j \alpha_j}{Z}\right) \\ &= N \sum_{j=1}^k \frac{p_j \alpha_j (Z - p_j \alpha_j)}{Z} = \frac{N}{Z} \left(Z^2 - \sum_{j=1}^k p_j^2 \alpha_j^2\right) \\ &= \frac{N}{Z} \sum_{i \neq j} p_i \alpha_i p_j \alpha_j.\end{aligned}\tag{17}$$

Equivalently, in matrix form let $\mathbf{p} = (p_1, \dots, p_k)^T$ and $\boldsymbol{\alpha} = (\alpha_1, \dots, \alpha_k)^T$. Then $\mathbf{p}^T(\boldsymbol{\alpha}\boldsymbol{\alpha}^T - \text{diag}(\boldsymbol{\alpha}^2))\mathbf{p} = \sum_{i \neq j} p_i \alpha_i p_j \alpha_j$, and $Z = \sum_{j=1}^k p_j \alpha_j = \mathbf{p}^T \boldsymbol{\alpha}$. So we may define

$$f(\boldsymbol{\alpha}) = \Delta(v|\psi') = N \frac{\mathbf{p}^T(\boldsymbol{\alpha}\boldsymbol{\alpha}^T - \text{diag}(\boldsymbol{\alpha}^2))\mathbf{p}}{\mathbf{p}^T \boldsymbol{\alpha}}.$$

To prove adaptive submodularity, it suffices to show f is coordinate-wise non-decreasing on $[0, 1]^k$, since then conditioning from $\boldsymbol{\alpha} = \mathbf{1}$ (smaller partial observation) to any $\boldsymbol{\alpha} \in [0, 1]^k$ (bigger partial observation) can only decrease f (the marginal benefit of adding v to the intervention set).

Now, we can compute the partial derivatives as follows

$$\frac{\partial f}{\partial \alpha_i} = N \frac{\left(\frac{\partial}{\partial \alpha_i} \mathbf{p}^T(\boldsymbol{\alpha}\boldsymbol{\alpha}^T - \text{diag}(\boldsymbol{\alpha}^2))\mathbf{p}\right) Z - \mathbf{p}^T(\boldsymbol{\alpha}\boldsymbol{\alpha}^T - \text{diag}(\boldsymbol{\alpha}^2))\mathbf{p} \frac{\partial Z}{\partial \alpha_i}}{Z^2}.\tag{18}$$

One checks easily

$$\frac{\partial Z}{\partial \alpha_i} = p_i, \quad \frac{\partial}{\partial \alpha_i} [\mathbf{p}^T(\boldsymbol{\alpha}\boldsymbol{\alpha}^T - \text{diag}(\boldsymbol{\alpha}^2))\mathbf{p}] = 2p_i \sum_{j \neq i} p_j \alpha_j.$$

Substituting into (18) gives

$$\begin{aligned}\frac{\partial f}{\partial \alpha_i} &= N \frac{2p_i(\sum_{j \neq i} p_j \alpha_j) Z - (\sum_{r \neq s} p_r \alpha_r p_s \alpha_s) p_i}{Z^2} \\ &= N p_i \frac{2Z(Z - p_i \alpha_i) - (Z^2 - \sum_{j=1}^k (p_j \alpha_j)^2)}{Z^2} \\ &= \frac{N p_i}{Z^2} (Z^2 - 2p_i \alpha_i Z + \sum_{j=1}^k (p_j \alpha_j)^2)\end{aligned}$$

Thus, to show that $\frac{\partial f}{\partial \alpha_i} \geq 0$, one needs to show $Z^2 - 2p_i \alpha_i Z + \sum_{j=1}^k (p_j \alpha_j)^2 \geq 0$. It can be shown by some simple algebraic manipulation. By defining $Z_{-i} = \sum_{j \neq i} p_j \alpha_j$, we have

$$\begin{aligned}Z^2 - 2p_i \alpha_i Z + \sum_{j=1}^k (p_j \alpha_j)^2 &= (Z_{-i} + p_i \alpha_i)^2 - 2p_i \alpha_i (Z_{-i} + p_i \alpha_i) + \sum_{j=1}^k (p_j \alpha_j)^2 \\ &= Z_{-i}^2 + 2p_i \alpha_i Z_{-i} + p_i^2 \alpha_i^2 - 2p_i \alpha_i Z_{-i} - 2p_i^2 \alpha_i^2 + \sum_{j=1}^k (p_j \alpha_j)^2 \\ &= Z_{-i}^2 + \sum_{j \neq i} (p_j \alpha_j)^2 \geq 0\end{aligned}$$

Hence f is coordinatewise non-decreasing on $[0, 1]^k$. In particular, passing from $\boldsymbol{\alpha} = \mathbf{1}$ to any $\boldsymbol{\alpha} \in [0, 1]^k$ can only decrease f , which completes the proof of adaptive submodularity via α -parameterization. \square

D Concentration Bounds for the Normalized Marginal Benefit

This section establishes a concentration bound for the empirical normalized marginal benefit

$$L(\hat{\mathbf{p}}) = \sum_{i=1}^k \hat{p}_i(1 - \hat{p}_i),$$

computed from M i.i.d. samples of a categorical distribution $\mathbf{p} \in \Delta_k$. The quantity $L(\hat{\mathbf{p}})$ is used in our algorithm as a heuristic to estimate the benefit of intervening on a given variable. We aim to bound the estimation error

$$|L(\mathbf{p}) - L(\hat{\mathbf{p}})|.$$

Using the triangle inequality, we write:

$$|L(\mathbf{p}) - L(\hat{\mathbf{p}})| \leq |L(\mathbf{p}) - \mathbb{E}[L(\hat{\mathbf{p}})]| + |\mathbb{E}[L(\hat{\mathbf{p}})] - L(\hat{\mathbf{p}})|.$$

This gives rise to a bias term and a deviation term, which we control separately below.

D.1 Bounding the Bias Term

Lemma D.1. *Let $\hat{\mathbf{p}}$ denote the empirical distribution over M i.i.d. samples from a categorical distribution $\mathbf{p} \in \Delta_k$. Then the bias of the estimator satisfies*

$$|L(\mathbf{p}) - \mathbb{E}[L(\hat{\mathbf{p}})]| \leq \frac{L(\mathbf{p})}{M} \leq \frac{1}{M}.$$

Proof. We compute the bias directly:

$$\begin{aligned} \mathbb{E}[L(\hat{\mathbf{p}})] &= \mathbb{E} \left[\sum_{i=1}^k \hat{p}_i(1 - \hat{p}_i) \right] \\ &= \sum_{i=1}^k \mathbb{E}[\hat{p}_i] - \sum_{i=1}^k \mathbb{E}[\hat{p}_i^2]. \end{aligned}$$

Note that $\mathbb{E}[\hat{p}_i] = p_i$, and:

$$\begin{aligned} \mathbb{E}[\hat{p}_i^2] &= \mathbb{E} \left[\left(\frac{1}{M} \sum_{j=1}^M \mathbf{1}\{X_j = i\} \right)^2 \right] \\ &= \frac{1}{M^2} \sum_{j=1}^M \sum_{\ell=1}^M \mathbb{E}[\mathbf{1}\{X_j = i\} \cdot \mathbf{1}\{X_\ell = i\}] \\ &= \frac{1}{M^2} \sum_{j=1}^M \mathbb{P}[X_j = i] + \frac{1}{M^2} \sum_{j \neq \ell} \mathbb{P}[X_j = i, X_\ell = i] \\ &= \frac{1}{M^2} (Mp_i + M(M-1)p_i^2) \\ &= \frac{p_i}{M} + \left(1 - \frac{1}{M}\right) p_i^2. \end{aligned}$$

Therefore,

$$\begin{aligned} \mathbb{E}[L(\hat{\mathbf{p}})] &= \sum_{i=1}^k p_i - \sum_{i=1}^k \left(\frac{p_i}{M} + \left(1 - \frac{1}{M}\right) p_i^2 \right) \\ &= \left(1 - \frac{1}{M}\right) \left(\sum_{i=1}^k p_i - \sum_{i=1}^k p_i^2 \right) \\ &= \left(1 - \frac{1}{M}\right) L(\mathbf{p}), \end{aligned}$$

which yields

$$\begin{aligned}
|L(\mathbf{p}) - \mathbb{E}[L(\hat{\mathbf{p}})]| &= \frac{1}{M} L(\mathbf{p}) \\
&= \frac{1}{M} \sum_{i=1}^k p_i (1 - p_i) \\
&\stackrel{*}{\leq} \left(1 - \frac{1}{k}\right) \frac{1}{M} \\
&\leq \frac{1}{M}.
\end{aligned}$$

Justification for *. The bound follows from the identity

$$\sum_{i=1}^k p_i (1 - p_i) = \sum_{i=1}^k p_i - \sum_{i=1}^k p_i^2 = 1 - \sum_{i=1}^k p_i^2,$$

using the fact that $\sum_{i=1}^k p_i = 1$ for any $p \in \Delta_k$. To maximize this expression, we minimize $\sum_{i=1}^k p_i^2$ under the simplex constraint. By the Cauchy-Schwarz inequality,

$$\left(\sum_{i=1}^k p_i^2\right) \left(\sum_{i=1}^k 1^2\right) \geq \left(\sum_{i=1}^k p_i\right)^2 = 1^2,$$

so $\sum_{i=1}^k p_i^2 \geq \frac{1}{k}$, with equality if and only if $p_i = \frac{1}{k}$ for all i . Thus, the maximum of the original expression is $1 - \frac{1}{k}$, attained uniquely by the uniform distribution. \square

D.2 Bounding the Deviation Term

To bound the deviation term $|\mathbb{E}[L(\hat{\mathbf{p}})] - L(\hat{\mathbf{p}})|$, we invoke the bounded differences inequality. We begin with the following definition:

Definition D.2. A function $f : \mathcal{X}^M \rightarrow \mathbb{R}$ satisfies the bounded difference property with parameters c_1, \dots, c_M if changing the i -th coordinate of the input changes the function value by at most c_i . That is, for all $i \in [M]$ and any input vectors $x_1, \dots, x_M \in \mathcal{X}$, and $x'_i \in \mathcal{X}$,

$$|f(x_1, \dots, x_i, \dots, x_M) - f(x_1, \dots, x'_i, \dots, x_M)| \leq c_i.$$

A well-known result [35] establishes concentration of measure for bounded difference functions when their inputs are independent random variables.

Theorem D.3 (Bounded Differences Inequality, Theorem 2.9.1 in [35]). *Let $X = (X_1, \dots, X_M)$ be a vector of independent random variables, and let $f : \mathcal{X}^M \rightarrow \mathbb{R}$ satisfy the bounded difference property with parameters c_i . Then for any $t > 0$,*

$$\mathbb{P}[f(X) - \mathbb{E}f(X) \geq t] \leq \exp\left(-\frac{2t^2}{\sum_{i=1}^M c_i^2}\right).$$

Lemma D.4. *Let $\hat{\mathbf{p}}$ denote the empirical distribution over M i.i.d. samples $X_1, \dots, X_M \in [k]$, and define $L(X_1, \dots, X_M) = L(\hat{\mathbf{p}}) = \sum_{i=1}^k \hat{p}_i (1 - \hat{p}_i)$. Then the function $L(X_1, \dots, X_M) : [k]^M \rightarrow \mathbb{R}$ satisfies the bounded difference property with parameters $c_i = \frac{2}{M}$ for all $i \in [M]$. Based on theorem D.3, it implies the following concentration inequality:*

$$\mathbb{P}[|L(\hat{\mathbf{p}}) - \mathbb{E}[L(\hat{\mathbf{p}})]| \geq t] \leq 2 \exp\left(-\frac{Mt^2}{2}\right).$$

Proof. We prove the bounded difference property with respect to the first argument; the proofs for the remaining coordinates follow analogously.

Consider two datasets that differ only in their first element:

$$\mathbf{X} = (X_1 = a, X_2, X_3, \dots, X_N) \quad \text{and} \quad \mathbf{X}' = (X'_1 = b, X_2, X_3, \dots, X_N),$$

where $a \neq b$. Let $\hat{\mathbf{p}}$ denote the empirical distribution based on \mathbf{X} , and $\hat{\mathbf{q}}$ the empirical distribution based on \mathbf{X}' . Since the two datasets differ only at position 1, the empirical distributions $\hat{\mathbf{p}}$ and $\hat{\mathbf{q}}$ will differ only in the entries corresponding to a and b . We can therefore compute the difference in the function values based on the differences in these two entries as follows:

$$\begin{aligned}
|L(\mathbf{X}) - L(\mathbf{X}')| &= |L(\hat{\mathbf{p}}) - L(\hat{\mathbf{q}})| \\
&= |\hat{p}_a(1 - \hat{p}_a) + \hat{p}_b(1 - \hat{p}_b) - (\hat{q}_a(1 - \hat{q}_a) + \hat{q}_b(1 - \hat{q}_b))| \\
&\leq |\hat{p}_a(1 - \hat{p}_a) - \hat{q}_a(1 - \hat{q}_a)| + |\hat{p}_b(1 - \hat{p}_b) - \hat{q}_b(1 - \hat{q}_b)| \\
&= |\hat{p}_a - \hat{q}_a - (\hat{p}_a^2 - \hat{q}_a^2)| + |\hat{p}_b - \hat{q}_b - (\hat{p}_b^2 - \hat{q}_b^2)| \\
&\leq |\hat{p}_a - \hat{q}_a| \cdot \underbrace{|1 - (\hat{p}_a + \hat{q}_a)|}_{\leq 1} + |\hat{p}_b - \hat{q}_b| \cdot \underbrace{|1 - (\hat{p}_b + \hat{q}_b)|}_{\leq 1} \\
&\stackrel{(1)}{\leq} \underbrace{|\hat{p}_a - \hat{q}_a|}_{\leq \frac{1}{M}} + \underbrace{|\hat{p}_b - \hat{q}_b|}_{\leq \frac{1}{M}} \\
&\stackrel{(2)}{\leq} \frac{2}{M}
\end{aligned}$$

Since $\hat{p}_a, \hat{q}_a \in [0, 1]$, it follows that $|1 - (\hat{p}_a + \hat{q}_a)| \leq 1$. The same argument applies to the term involving b . This yields inequality (1).

Inequality (2) uses the fact that changing one data point in the dataset affects at most two entries in the empirical distribution: one entry increases by $1/M$, and another decreases by $1/M$. Therefore,

$$|\hat{p}_a - \hat{q}_a| \leq \frac{1}{M}, \quad |\hat{p}_b - \hat{q}_b| \leq \frac{1}{M}.$$

□

D.3 Proof of Theorem 7.1

To bound the estimation, we can combine the previous results as follows

$$\begin{aligned}
\mathbb{P}[|L(\mathbf{p}) - L(\hat{\mathbf{p}})| \geq t] &\leq \mathbb{P}[|L(\hat{\mathbf{p}}) - \mathbb{E}[L(\hat{\mathbf{p}})]| + \underbrace{|L(\mathbf{p}) - \mathbb{E}[L(\hat{\mathbf{p}})]|}_{\stackrel{\text{lemma D.1}}{\leq} \frac{1}{M}} \geq t] \\
&\leq \mathbb{P}[|L(\hat{\mathbf{p}}) - \mathbb{E}[L(\hat{\mathbf{p}})]| \geq t - \frac{1}{M}] \\
&\stackrel{\text{lemma D.4}}{\leq} 2 \exp\left(-\frac{M}{2} \left(t - \frac{1}{M}\right)^2\right)
\end{aligned}$$

Setting $t = \frac{1}{M} + \sqrt{\frac{2}{M} \log(\frac{2}{\epsilon})}$ implies

$$\mathbb{P}\left[|L(\mathbf{p}) - L(\hat{\mathbf{p}})| \geq \frac{1}{M} + \sqrt{\frac{2}{M} \log(\frac{2}{\epsilon})}\right] \leq \epsilon. \quad (19)$$

□

E Implementation and Experimental Details

E.1 Pseudocode for the Idealized Setting

The pseudocode for our adaptive intervention selection strategy, assuming an ideal ICA oracle, is shown in Algorithm 1. The method begins by extracting an observational estimate of the causal matrix, which is assumed to be a perfectly permuted and scaled version of the true matrix. It then builds a

bipartite graph representing the equivalence class of candidate graphs. At each iteration, the algorithm greedily selects the intervention expected to eliminate the largest number of inconsistent graphs by sampling perfect matchings. Once an intervention is performed, the corresponding matching edge is identified, and the process continues until the causal graph is uniquely identified or the budget is exhausted.

Algorithm 1 Adaptive Experiment Design for Cyclic LSCMs (Idealized Setting)

- 1: **Input:** Observational distribution, budget K
 - 2: **Output:** Reduced equivalence class after K adaptive interventions
 - 3: Run **ideal ICA** on observational data to estimate $I - W_{\text{ICA}}$
 - 4: Construct the bipartite graph from nonzero entries of $I - W_{\text{ICA}}$
 - 5: **for** $t = 1$ to K **do**
 - 6: Check if the bipartite graph admits a **unique perfect matching**; if so, **terminate**
 - 7: Sample M_t perfect matchings from the bipartite graph using a sampler oracle
 - 8: Estimate the normalized marginal benefit $\hat{B}(j) = \sum_i \hat{p}_{ij}(1 - \hat{p}_{ij})$ for each variable j
 - 9: Select variable j^* with maximum estimated benefit and intervene on it
 - 10: Run **ideal ICA** on the interventional distribution to estimate $I - W_{\text{ICA}}^{(j^*)}$
 - 11: Recover the true matching edge $(r_{\pi^{-1}(j^*)}, j^*)$ by comparing pre- and post-intervention estimates.
 - 12: Remove $r_{\pi^{-1}(j^*)}, j^*$, and their connecting edge from the bipartite graph
 - 13: **end for**
-

E.2 Experimental Details for the Idealized Setting


This section provides additional implementation details for the simulations presented under the ideal ICA assumption in the main text. The intervention sampler, averaging methodology, and results on dense graphs are described below.

Intervention Sampler. To estimate the marginal benefit of each candidate intervention, we rely on random samples of perfect matchings in the bipartite graph. When exact enumeration is computationally infeasible, we use a greedy heuristic to generate random (but not uniform) perfect matchings. Starting from an empty matching, the sampler proceeds by iteratively selecting an unmatched right-node v that has the minimum degree, choosing an adjacent unmatched left-node u uniformly at random, and adding the pair (u, v) to the matching. Then we remove v, u , and all their incident edges from the bipartite graph. This process continues until a perfect matching is formed. One can show that this method always finds a random perfect matching if one exists. While this sampling scheme is non-uniform, it is scalable and yields accurate marginal benefit estimates in practice.


Averaging Across Trials. For each fixed number of nodes n , we repeat the simulation over 60 random graph instances to mitigate randomness in graph generation and intervention ordering. We report the *average number of interventions* required for full identification and visualize variability using *standard deviation bars* around the mean in our plots in Figure 5. These error bars demonstrate the stability of each strategy across different random seeds.

Experiments on Dense Graphs. In the main experiments, we used Erdős-Rényi random graphs with edge probability $p = c/n$, ensuring that the expected degree remains constant as n increases. In this regime, the feedback vertex set (FVS) can be computed exactly, and serves as a theoretical lower bound. To further stress-test the methods, we also evaluate performance on *dense* graphs where each edge is included independently with a fixed probability (e.g., $p = 0.2$). These dense graphs often contain an exponential number of cycles, making exact FVS computation infeasible.

Nonetheless, we compare the performance of our adaptive method against Random and Max-Degree baselines in these settings. The results, shown in Figure 5b, reveal that our algorithm continues to outperform these heuristics even in dense regimes where the graph structure is significantly more complex. For full reproducibility, including code and data used in all experiments, we refer to the repository available at <https://github.com/EhsanSharifian/exp-design-cyclic-models>.



(a) Sample mode on sparse graphs with a maximum of 40 nodes.



(b) Sample mode on dense graphs with a maximum of 60 nodes.

Figure 5: Comparison of intervention strategies under the ideal ICA assumption with a sample-based matching sampler. Error bars indicate standard deviation over trials.

E.3 Implementation in the Finite-Sample Setting

This section provides the implementation details for the finite-sample setting, where we use the FastICA algorithm on finite samples instead of having access to an ideal oracle. This introduces estimation noise, necessitating a more robust algorithm, presented in Algorithm 2.

The algorithm incorporates two key modifications to handle noise. First, after running ICA, it performs *adaptive thresholding* (Line 4): it begins with strict thresholds to remove noise and progressively relaxes them only until the resulting bipartite graph admits at least one perfect matching. Second, after an intervention, it uses a *safe matching* procedure to identify the correct row (Lines 13-20). Instead of picking the single best match directly, it ranks candidate rows by similarity and iterates through them. It accepts the first candidate that, when matched, leaves the remaining graph in a state where a perfect matching is still possible. This conservative strategy ensures robustness against estimation errors.

Algorithm 2 ICA-based Adaptive Experiment Design for Cyclic LSCMs

```

1: Input: Observational data, intervention budget  $K$ 
2: Output: Recovered causal matrix  $\widehat{W}$ 
3: Run FastICA on the observational data to estimate  $I - W_{ICA}$ 
4: Normalize and threshold  $I - W_{ICA}$  adaptively to ensure matchability
5: Construct a bipartite graph from the support of the thresholded matrix
6: for  $t = 1$  to  $K$  do
7:   Check for a unique perfect matching; if found, terminate
8:   Sample  $M_t$  perfect matchings from the bipartite graph
9:   Compute normalized marginal benefit  $\hat{B}(j) = \sum_i \hat{p}_{ij}(1 - \hat{p}_{ij})$  for each variable  $j$ 
10:  Select the variable  $j^*$  with maximum estimated benefit and intervene on it
11:  Run FastICA on the interventional distribution to estimate  $I - W_{ICA}^{(j^*)}$ 
12:  Rank candidate matching edges  $(r_i, j^*)$  by comparing rows of pre- and post-intervention
    matrices
13:  for each candidate edge  $(r_i, j^*)$  in ranked order do
14:    Temporarily remove  $r_i$  and  $j^*$  from the bipartite graph
15:    if the remaining graph admits a perfect matching then
16:      Accept  $(r_i, j^*)$  as the matching edge
17:      Remove  $r_i, j^*$ , and their connecting edge from the actual graph
18:      Break
19:    end if
20:  end for
21:  if no valid edge found then
22:    Terminate with failure
23:  end if
24: end for
  
```




Figure 6: Comparison of intervention strategies under finite-sample ICA with a sample-based matching sampler. Error bars represent standard deviation over trials.

The performance of this robust algorithm is detailed in Figure 6. The results show that our adaptive strategy continues to outperform the Random and Max Degree baselines, demonstrating its practical effectiveness. Furthermore, Table 1 quantifies the accuracy of the final recovered matrix \widehat{W} against the ground truth W^* , demonstrating high-fidelity recovery.

Table 1: Distribution of relative errors ($\varepsilon_{\text{rel}} = \frac{\|\widehat{W} - W^*\|_F}{\|W^*\|_F}$) for the recovered matrix \widehat{W} across all experimental runs in the finite-sample setting.

Relative Error Range	Number of Instances
< 0.05	860
0.05 – 0.15	275
0.15 – 0.25	95
> 0.25	470

F Details on Extensions

F.1 Multi-Node Interventions

The adaptive framework presented in the main text focuses on single-node interventions but can be naturally extended to handle simultaneous interventions on multiple nodes. This extension allows for the parallelization of experiments, which can significantly improve the efficiency of the causal discovery process. Here, we provide a detailed analysis of this generalization.

Formalizing the Multi-Node Intervention. Suppose an experiment is conducted where hard interventions are simultaneously applied to a set of t variables, denoted by $E = \{i_1, i_2, \dots, i_t\}$. A hard intervention on a variable effectively removes all its incoming causal links. The post-intervention weight matrix, $W^{(E)}$, is therefore obtained by setting the rows corresponding to the indices in E to zero. Formally, for a row index k :

$$(W^{(E)})_{k,\cdot} = \begin{cases} \mathbf{0} & \text{if } k \in E \\ W_{k,\cdot} & \text{if } k \notin E \end{cases}$$

Applying ICA to the observational and interventional data yields the permuted and scaled matrices $PD(I - W)$ and $P'D'(I - W^{(E)})$, respectively. The core idea is to identify structural information by comparing the set of rows from these two matrices.

Localizing the Permutation Ambiguity. The key difference between $(I - W)$ and $(I - W^{(E)})$ lies precisely in the t rows corresponding to the interventions. When we compare the set of n rows recovered from the observational data with the n rows from the interventional data, we will find that:

- $n - t$ rows from the observational set have a corresponding match in the interventional set (up to scaling). These are the rows associated with the non-intervened variables $V \setminus E$.
- t rows from the observational set have no match. These "unmatched" rows must correspond to the variables in the intervened set E .

This single multi-node intervention effectively partitions the global row-permutation problem. It breaks the cycle reversion ambiguity that might exist between the nodes in E and the nodes in $V \setminus E$. We now know with certainty that the t unmatched observational rows must be assigned to the t true row indices corresponding to the variables in E .

The global challenge of finding one permutation of size n is thus reduced to two smaller, independent subproblems: one for the set E and one for the set $V \setminus E$. The specific assignment of the t identified rows to the t variables within E remains to be solved, a subproblem that is addressed by the acyclicity condition discussed next.

Condition for Full Identification within the Intervened Set. While a multi-node intervention is always informative for localizing ambiguity, we can achieve an even stronger result (full identification of all intervened nodes) if the intervention set E is chosen carefully. Specifically, if the subgraph induced by the variables in E is *acyclic*, the internal permutation ambiguity within E can be fully resolved.

A crucial subtlety is that we cannot verify the acyclicity of an *arbitrary* set of variables directly from the observational matrix $I - W_{\text{ICA}}$, due to the row-permutation ambiguity we aim to solve. However,

a **verifiable sufficient condition** can be considered herein. The **Strongly Connected Components (SCCs)** of the causal graph are identifiable from the observational data, as the component structure is invariant across all graphs in the permutation-equivalence class.

Therefore, a verifiable strategy is to select an intervention set E such that each variable $i_k \in E$ belongs to a different SCC. By definition, nodes residing in distinct SCCs cannot form a cycle among themselves. This choice guarantees that the subgraph induced by E is acyclic.

Consequently, when we intervene on such a set and identify the t unmatched observational rows, there will be only one way to assign these rows to the variables in E that is consistent with an acyclic structure. This uniquely reveals the true row for each intervened variable and fully resolves the local ambiguity.

Implications for Experimental Design. This extension enables a more efficient, parallelized experimental design strategy. Instead of selecting one node at a time, the algorithm could identify maximal subsets of variables that belong to different SCCs, and intervene on them simultaneously. This would remarkably reduce the number of experiments required for full graph identification.

F.2 Imperfect and Noisy Interventions

The core framework described in this paper assumes “perfect” or “hard” interventions, where all incoming causal mechanisms for a target variable are removed (i.e., the corresponding row in the weight matrix W is set to zero). However, the method’s validity is not restricted to this idealized case. It can be extended to a broad class of imperfect or “soft” interventions, which are often more representative of real-world experiments.

The General Identifiability Condition. The fundamental mechanism of our approach is the identification of the intervened variable by detecting a change in its corresponding structural equation. This is achieved by comparing the set of rows recovered from observational data, $R_{obs} = \{\text{rows of } (I - W)\}$, with the set of rows recovered from interventional data, R_{int} . An intervention on a variable X_k is considered informative as long as it sufficiently perturbs the k -th row of the causal matrix.

Let the imperfect intervention on X_k change its corresponding row in the weight matrix from $W_{k,\cdot}$ to a new row $W'_{k,\cdot}$. The new post-intervention matrix is $(I - W')$. The intervention is guaranteed to be identifiable if the new, perturbed row $(I - W')_{k,\cdot}$ is not a scaled version of any of the *other* original rows from the observational matrix. That is, for all $j \neq k$:

$$(I - W')_{k,\cdot} \neq c \cdot (I - W)_{j,\cdot} \quad \text{for any scaling factor } c.$$

If this condition holds, the original row $(I - W)_{k,\cdot}$ will be the only row in R_{obs} that has no corresponding match in R_{int} , uniquely revealing k as the target of the intervention.

This identifiability condition holds with high probability for interventions that alter the causal mechanism of a variable. The chances of a random modification creating a row that perfectly aligns with an existing one are negligible.

F.3 Generalization to Non-linear Models

While this work focuses on linear models for their tractability and widespread use, the core principles of our adaptive design framework (representing ambiguity via a bipartite graph and greedily selecting interventions to reduce it) can be extended to non-linear Structural Causal Models (SCMs). This generalization relies on recent advances in non-linear Independent Component Analysis (ICA) and the structural information revealed by the Jacobian of the causal mechanism.

Causal Discovery via the Jacobian. Consider a general non-linear SCM defined by the structural equations $X_i = f_i(\text{Pa}_i, N_i)$, where Pa_i are the parents of node i and N_i are independent noise terms. Let \mathbf{f} be the vector-valued function that maps the exogenous noises \mathbf{N} to the observed variables \mathbf{X} . The inverse of this map, \mathbf{f}^{-1} , recovers the noises from the variables, such that $\mathbf{N} = \mathbf{f}^{-1}(\mathbf{X})$.

This formulation is a direct generalization of the linear model. In our linear case, the relationship is $\mathbf{N} = (I - W)\mathbf{X}$, meaning the matrix $(I - W)$ itself is the (linear) inverse map. For the non-linear case, a key result from [27] for acyclic graphs shows that the *Jacobian of the inverse map*, $\nabla \mathbf{f}^{-1}(\mathbf{X})$,

plays the analogous role. Specifically, its support (the pattern of non-zero entries) is equivalent to the support of $I - A$, where A is the graph’s adjacency matrix. We conjecture that this result can be extended to cyclic models under similar identifiability conditions.

Role of Non-linear ICA. The theoretical connection between the Jacobian and the graph structure can be operationalized using modern non-linear ICA algorithms [17, 18]. These methods can learn the inverse map \mathbf{f}^{-1} from observational data alone, up to known indeterminacies (such as permutation, scaling, and element-wise monotonic transformations). By computing the Jacobian of the learned function $\hat{\mathbf{f}}^{-1}$, we can obtain an estimate of the matrix $I - A$ up to a row permutation and scaling. Let us denote this recovered matrix by $(I - A)_{\text{ICA}}$.

Adapting the Experimental Design Framework. The recovered matrix $(I - A)_{\text{ICA}}$ serves the exact same purpose in the non-linear setting as $(I - W_{\text{ICA}})$ does in our linear framework. Its support can be used to construct a bipartite graph that perfectly characterizes the permutation ambiguity of the observational equivalence class. From this representation, we can identify the SCCs and apply the same adaptive intervention strategy to resolve the ambiguity. An intervention on a node X_k would modify its function f_k , which in turn would alter the k -th row of the Jacobian, allowing for its identification.

Current Limitation. The primary challenge for this non-linear extension lies in the identifiability constraints. The main structural requirement is that *each node in the causal graph must have a unique parent set* to ensure the rows of the Jacobian have unique sparsity patterns to enable correct matching in the bipartite graph. Consequently, based on current results in nonlinear ICA, it is needed that each node has a unique parent set in order to carry out the experiment design step. We hope that this assumption can be relaxed with future advances in nonlinear ICA.

NeurIPS Paper Checklist

The checklist is designed to encourage best practices for responsible machine learning research, addressing issues of reproducibility, transparency, research ethics, and societal impact. Do not remove the checklist: **The papers not including the checklist will be desk rejected.** The checklist should follow the references and follow the (optional) supplemental material. The checklist does NOT count towards the page limit.

Please read the checklist guidelines carefully for information on how to answer these questions. For each question in the checklist:

- You should answer [Yes], [No], or [NA].
- [NA] means either that the question is Not Applicable for that particular paper or the relevant information is Not Available.
- Please provide a short (1–2 sentence) justification right after your answer (even for NA).

The checklist answers are an integral part of your paper submission. They are visible to the reviewers, area chairs, senior area chairs, and ethics reviewers. You will be asked to also include it (after eventual revisions) with the final version of your paper, and its final version will be published with the paper.

The reviewers of your paper will be asked to use the checklist as one of the factors in their evaluation. While "[Yes]" is generally preferable to "[No]", it is perfectly acceptable to answer "[No]" provided a proper justification is given (e.g., "error bars are not reported because it would be too computationally expensive" or "we were unable to find the license for the dataset we used"). In general, answering "[No]" or "[NA]" is not grounds for rejection. While the questions are phrased in a binary way, we acknowledge that the true answer is often more nuanced, so please just use your best judgment and write a justification to elaborate. All supporting evidence can appear either in the main paper or the supplemental material, provided in appendix. If you answer [Yes] to a question, in the justification please point to the section(s) where related material for the question can be found.

IMPORTANT, please:

- **Delete this instruction block, but keep the section heading “NeurIPS Paper Checklist”,**
- **Keep the checklist subsection headings, questions/answers and guidelines below.**
- **Do not modify the questions and only use the provided macros for your answers.**

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper’s contributions and scope?

Answer: [Yes]

Justification: The main contribution of the paper is to introduce near-optimal experiment design in linear non-Gaussian cyclic model, as mentioned in the the abstract and introduction.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: The assumptions of the work are clearly mentioned in Assumption 3.1-3.3.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory assumptions and proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [\[Yes\]](#)

Justification: We present proofs of all theorems in the appendix, for improved readability and space limitations.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental result reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [\[Yes\]](#)

Justification: The codes are available as a supplementary material and experimental setups are explained in the appendix.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
 - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: The codes are available as a supplementary material to faithfully reproduce the results.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.

- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental setting/details

Question: Does the paper specify all the training and test details (e.g., data splits, hyper-parameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [\[Yes\]](#)

Justification: The detail of experiments are discussed in Section 8 and in the appendix.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment statistical significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [\[Yes\]](#)

Justification: The confidence intervals are reported in the appendix.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments compute resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [\[Yes\]](#)

Justification: Compute resources are reported in the appendix.

Guidelines:

- The answer NA means that the paper does not include experiments.

- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code of ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

Answer: [Yes]

Justification: -

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [NA]

Justification: This paper presents work whose goal is to advance the field of machine learning. There are general potential societal consequences as a result of improvement in alignment techniques, none of which we feel must be specifically highlighted.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: -

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: All code-bases used in the experiments are cited.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. New assets

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification: -

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. Crowdsourcing and research with human subjects

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: -

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. Institutional review board (IRB) approvals or equivalent for research with human subjects

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: -

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. Declaration of LLM usage

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does not impact the core methodology, scientific rigor, or originality of the research, declaration is not required.

Answer: [NA]

Justification: -

Guidelines:

- The answer NA means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy (<https://neurips.cc/Conferences/2025/LLM>) for what should or should not be described.