

Relaxing Faithfulness with Intervention-Only Causal Discovery

Bijan H. S. Mazaheri^{1,2}

Jiaqi Zhang^{2,3}

Caroline Uhler^{2,3}

¹Thayer School of Engineering, Dartmouth College, Hanover, New Hampshire, USA

²Schmidt Center, Broad Institute of MIT and Harvard, Cambridge, Massachusetts, USA

³Laboratory for Information and Decision Systems, Massachusetts Institute of Technology, Cambridge, Massachusetts, USA

Abstract

Causal discovery algorithms learn a network that describes the causal dependencies among random variables. A common workflow involves first utilizing conditional independence properties on observational data to determine partially directed causal relationships, then applying interventions to orient the unknown causal directions. A critical assumption for the first step is faithfulness: a requirement that causally linked variables exhibit statistical dependence. Many natural systems include buffering and stabilizing pathways that cancel out to achieve systemic robustness. This cancellation of pathways violates faithfulness, leading causal discovery algorithms to incorrectly remove causal dependencies. In this paper, we argue that hard interventions contain information about the presence/absence of causal linkage that is overlooked in the first stage of structure discovery. We show that a mild assumption — called intervention-immediacy faithfulness — that allows cancellations, is sufficient to non-parametrically identify causal structures with hard interventions. These results position interventions as the primary carriers of information about causal structure, which should take precedence over conditional independence testing. To flip the paradigm, we also specify equivalence classes when the identification criteria are not met due to limitations in the scope of interventions.

(DAG) or a representation of an equivalence class of DAGs, though other targets, such as ancestral graphs with latent variables or cyclic models, are also studied (see [Vowels et al. \[2022\]](#), [Squires and Uhler \[2023\]](#) for reviews). One approach to causal discovery involves a *constraint-based* search guided by conditional independence (“CI-tests”), e.g., the PC-algorithm from [Spirtes et al. \[2000\]](#). Constraint-based approaches use the observation that variables without a direct causal link can be made independent by conditioning on variables on intermediary causal paths, known as the causal Markov condition [[Pearl, 2009](#)].

Structural causal models imply a data-generating process whereby each random variable is generated from its (causal) parents. These processes are sometimes specified using equations, leading to structural *equation* models (SEMs). Under these settings, these equations may be restricted using “parametric assumptions” on their functional form and noise. The most popular such assumption is linearity with additive Gaussian noise, yielding a multivariate Gaussian distribution on the full system.

Causal discovery algorithms therefore fall under two categories: (1) algorithms that make use of parametric assumptions on the structural equations, and (2) *non*-parametric algorithms that only make use of conditional independence and other broad distributional properties. While parametric settings, such as Linear Non-Gaussian Additive Noise Models (LinGAMs) [[Shimizu et al., 2006](#)], can sometimes fully recover a *single* DAG structure, most non-parametric algorithms for causal discovery can only narrow the structure down to what is known as a “Markov equivalence class,” uniquely specified by its conditional independence constraints [[Spirtes et al., 2000](#)]. Such a class is usually represented first by an undirected “skeleton,” denoting the *presence* of causal relationships, which is then partially oriented into a completed partial DAG (CP-DAG) [[Spirtes et al., 2000](#)].

Interventions. To resolve a Markov equivalence class, nonparametric causal discovery algorithms make *secondary*

1 INTRODUCTION

Causal Discovery. Structural Causal Models (SCMs), popularized by the works of [Pearl \[1998, 2009\]](#), graphically describe networks of causal dependencies. Causal discovery is the task of recovering the underlying causal structure from data, typically in the form of a *directed acyclic graph*

use of *interventions*. An intervention (a “perturbation” in biology) is a change to the system that *may* elicit a change in causally downstream variables, and *never* elicits a change in any other variables. Such interventions may be “soft”, such as “shift interventions” [Sani et al., 2020], which change a mean while retaining randomness and also retaining dependence on causally upstream variables). Real-world examples of soft interventions include CRISPR “gene knockdowns” [Chang et al., 2016] that partially suppress expression. Interventions may also be “hard”, such as “do interventions” [Pearl, 2009] which interrupt the data-generating process, force random variables to take a specific value, and dissociate their dependence from causally upstream variables. Real-world examples of these hard interventions are “gene knockouts” [Guan et al., 2010], which *fully suppress* their expression.

Faithfulness. Under causal sufficiency (no unobserved confounding), a data-generating process given by an SCM will exhibit the causal Markov condition, but the converse is not necessarily true: A causal path between two variables does not necessarily require that they be statistically dependent [Ramsey et al., 2012]. “Faithfulness” is the assumption that causal links imply statistical dependencies, which gives a two-way correspondence between CI-tests and graphical properties known as “d-separation” [Pearl, 2009]. This notion can be relaxed for interventions [Chevalley et al., 2025], requiring only that directed paths exhibit dependence.

Violations of both types of faithfulness are often due to “path cancellation,” i.e., dependencies whose effects cancel each other out. While often assumed away, it is important to understand that path cancellations occur by design in engineered and evolutionarily-driven systems. For example, many biochemical pathways are known to be buffered using “compensating paths” [Hartman IV et al., 2001]. Hence, while path cancellations may be unlikely in a randomly selected model for synthetic data, they are *very likely* in practice due to the formation of control systems.

Example 1.1. *Gene A may directly affect both B, C, but B may act as a compensating path to stabilize the expression of C even when A’s expression is suppressed. As such $A \rightarrow C$ and $A \rightarrow B \rightarrow C$ cancel out so that changes in A do not elicit changes in C.*

Notice that Example 1.1 exhibits path cancellation for both the conditional independence and intervention-based faithfulness definitions.

1.1 CONTRIBUTIONS

In this paper, we show that hard interventions can nonparametrically identify causal structures without assumptions that prohibit path cancellation. This result is based on a new notion of faithfulness, “intervention immediacy (II) faith-

fulness,” which we contrast with conditional independence (CI) faithfulness. II faithfulness only requires interventional dependence in a node’s “most immediate effects” — effects with only a single path and no opportunity for path cancellation (e.g., B for an intervention on A in Example 1.1).

While most notions of faithfulness are developed for non-parametric algorithms,¹ analyzing these assumptions in finite-sample regimes requires a parametric framework to quantify “closeness” to a violation. CI faithfulness has been studied under the parametric assumption of a multivariate Gaussian, where relaxation into “ ϵ -strong (CI) faithfulness” makes the assumption very strong. To compare II faithfulness and CI faithfulness, we provide a parametric study in the multivariate Gaussian setting and compare our results to those of Uhler et al. [2013]. II faithfulness is significantly milder than CI faithfulness under these metrics.

Our results position interventions as *superior* to CI testing, not only for *orienting* undirected relationships, but also for *discovering* their presence. This challenges the prevailing hierarchy that positions conditional independence as the primary source of causal information for learning the “causal skeleton” and demotes interventions as a secondary refinement.

The disruption of the CI/intervention hierarchy suggests a flipped workflow: utilize information from the available interventions first, then refine the *intervention* equivalence class with conditional independence testing that does not require assumptions beyond II faithfulness. To reorient algorithms around intervention-based causal discovery, we define these intervention equivalence classes based on limitations of the intervention cardinality.

1.2 OVERVIEW

Our algorithm is built on definitions of change sets and II faithfulness, formalized in Section 3. II faithfulness is briefly analyzed in the parametric setting of Gaussian random variables with linear dependencies in Section 4 for the sake of comparison to Uhler et al. [2013]. None of the algorithms in this paper require parametric assumptions.

In Section 5 we present two new causal discovery algorithms built on change sets and the assumption of II faithfulness. In this setting, identifiability depends on the number of variables on which we can intervene simultaneously (k) relative to the maximum local (vertex) connectivity $\kappa_{\max}$ between non-adjacent pairs of the graph (i.e., the largest, over all non-adjacent ordered pairs, of the minimum number of vertices whose removal destroys all directed paths between them).

To specify a new “intervention equivalence class,” we define the “ k -robust transitive closure” of $\mathcal{G}$. Informally, starting

¹Parametric models use more precise notions tuned to their settings, such as the matrix condition number.

from $\mathcal{G}$ and following a topological order (so that acyclicity is preserved), we add an edge $V_i \rightarrow V_j$ whenever V_i and V_j cannot be disconnected by removing fewer than k vertices; that is, whenever *every* vertex cut separating them has size at least k . Under limited cardinality of interventions, DAGs that have the same k -robust transitive closure form an equivalence class, which we cannot distinguish between using $\leq k$ -node interventions. We prove two guarantees: (1) a graph is identifiable with $\mathcal{O}(|\mathbf{V}|^2)$ hard interventions on up to $\kappa_{\max} + 1$ nodes, and (2) under a cardinality limit k it is identifiable up to its k -robust transitive closure. These equivalence classes provide insight into the information redundancies between intervention outcomes and conditional independence testing, and may be further resolved using regular CI tests as secondary sources of information.

In Section 6, we perform an empirical study on synthetic data to verify the relative robustness of intervention-only causal discovery. In the absence of real datasets with large-scale interventions, our results motivate the development of new methods for performing multi-node interventions to improve the accuracy of causal discovery.

1.3 RELATED WORKS

Interventions in Causal Discovery. Seminal work by Eberhardt and Scheines [2007] studied the integration of interventions into causal discovery, particularly on orienting Markov equivalence classes. Shanmugam et al. [2015] explored limited intervention size when orienting a Markov equivalence class, but not learning the graph. Hauser and Bühlmann [2012] defined interventional Markov equivalence classes with respect to intervention sets, then extended this to joint observational and interventional measurements in Hauser and Bühlmann [2015], which provided a critical framework for utilizing an intervention-second approach to causal discovery. Yang et al. [2018] and Kocaoglu et al. [2019] extended this characterization to soft interventions.

A large body of modern work integrates observational and interventional data *jointly* rather than sequentially, including GIES [Hauser and Bühlmann, 2012], IGSP and its variants [Wang et al., 2017, Squires et al., 2020], JCI [Mooij et al., 2020], and score-based continuous-optimization methods such as DCDI [Brouillard et al., 2020]. These methods typically operate *offline*, pooling a fixed observational sample with a pre-collected pool of interventional data. Our approach is instead *online* and *active*, in the tradition of adaptive experimental design for causal discovery [He and Geng, 2008, Shanmugam et al., 2015, Choo et al., 2022]: because the pruning step dynamically searches for a minimum vertex cut $\mathbf{Z}$, it designs targeted k -node interventions on the fly. It departs from that line in what the interventions are *for*: prior active methods design interventions to *orient* a skeleton that is still recovered from observational conditional-independence tests, whereas we use interven-

tions to establish edge *presence*, relying on observational CI only as an optional secondary refinement when the intervention cardinality is limited—never for the skeleton itself. This is what lets us side-step path cancellation entirely.

The work conceptually closest to ours is Chevalley et al. [2025], who introduce *intervention faithfulness* and use single-variable interventions to recover a causal order; II faithfulness weakens their condition so that it tolerates path cancellation (Lemma 4.3, Theorem 4.4), and our algorithms recover full structure rather than only an ordering. Also closely related, Zhou et al. [2025] characterizes and learns causal graphs from hard interventions in the *presence of latents* via conditional-independence-invariance machinery—complementary to our causally-sufficient but cancellation-tolerant setting. Two further methods are contrasts rather than close relatives: DCDI (above) assumes standard (λ -strong) faithfulness, which path cancellation violates, and identifies only up to the interventional Markov equivalence class, whereas our framework is nonparametric and faithfulness-relaxed; and effect identification given a known partial ancestral graph [Jaber et al., 2022] is orthogonal, addressing identification rather than structure learning.

Finite Sample Concerns for Faithfulness. Meek [1995] argued that faithfulness is a relatively mild assumption under exact statistics because it has full Lebesgue measure. However, success under finite-sample uncertainty requires a stronger notion — referred to as λ -strong faithfulness [Zhang and Spirtes, 2012] — to ensure that stochastic deviations from faithfulness are unlikely. From this perspective, Uhler et al. [2013] showed that faithfulness is a strong assumption because the manifold of violations behaves like a space-filling curve. These results are discussed in more detail in Section 2.

Causal Direction and Structure Learning. Significant recent work has shown that access to a causal ordering, even if such an ordering is searched for, drastically improves the performance of causal discovery. A notable algorithm that searches along permutations is “greedy sparsest permutation” (GSP) [Wang et al., 2017]. See Squires and Uhler [2023] for a review. These approaches are still based on CI testing, but we believe they are the preliminary foundations of a new causal discovery workflow.

A related discovery involves improving the number of CI tests utilized by causal discovery by integrating edge-orientation sooner than the original PC algorithm [Shiragur et al., 2024, Monés et al., 2026], which is still CI-based, gleanings orientations from unshielded collider structures.

Limiting Testing. Kocaoglu [2023] characterized the equivalence classes of graphs with conditional independence tests with restricted conditioning sets. We similarly characterize the equivalence classes of limited *interventions*.

Information Theoretic Causal Discovery. Intervention-based discovery shares many algorithmic similarities with information-theoretic causal discovery [Janzing and Schölkopf, 2010, Xu et al., 2025]. In particular, the algorithms in these approaches have a similar structure: they begin by isolating the transitive closure of the true graph and then removing edges. Our work diverges from this approach by focusing on interventions and the responses they elicit within the system.

2 PRELIMINARIES

2.1 NOTATION

We will use the capital Roman alphabet to denote random variables (e.g., A, B, C, V) and the lowercase Roman alphabet to denote assignments to those random variables (e.g., $A = a$ or just a). Bold will indicate a set of random variables, e.g., $\mathbf{V} = (V_1, V_2, \dots)^\top$, and $\mathbf{v}$ is an assignment to $\mathbf{V}$. Parents (**PA**), children (**CH**), ancestors (**AN**), and descendants (**DE**) in graphs will also follow these conventions, e.g., $\mathbf{PA}(V) = \mathbf{pa}^{\mathbf{v}}(v)$, where the assignments to those parents come from values specified in $\mathbf{v}$. We use subscripts to indicate the relevant graph structure, e.g., $\mathbf{PA}_{\mathcal{G}}(V)$. We will generally use the Greek alphabet (e.g., α, β) to represent parameters for structural equations and thresholds to quantify faithfulness.

Throughout the paper we assume the causal structure is a DAG and that the system is *causally sufficient* (no unobserved confounders). We follow the convention that a vertex is *not* counted among its own ancestors or descendants. The *transitive closure* $\text{TC}(\mathcal{G})$ of a DAG $\mathcal{G}$ is the graph on the same vertex set that contains a directed edge $V_i \rightarrow V_j$ whenever there is a directed path from V_i to V_j in $\mathcal{G}$ (equivalently, whenever $V_j \in \mathbf{DE}_{\mathcal{G}}(V_i)$). The *local (vertex) connectivity* of a non-adjacent ordered pair (V_i, V_j) is the minimum number of vertices whose removal destroys all directed paths from V_i to V_j ; by Menger’s theorem this equals the maximum number of internally vertex-disjoint directed $V_i \rightsquigarrow V_j$ paths. We write $\kappa_{\max}$ for the maximum of this quantity over all non-adjacent ordered pairs.

2.2 INTERVENTIONS

An intervention on a variable modifies its natural causal mechanism. In a general setting, multiple distinct interventions may exist for a single variable (e.g., setting its value, shifting its mean, etc.). For the purposes of this paper, however, we consider a simplified setting: for each variable $V \in \mathbf{V}$, there exists a single, unique **hard intervention** that can be applied. We denote the simultaneous application of this unique intervention to all variables in a set $\mathbf{A} \subseteq \mathbf{V}$ as $I(\mathbf{A})$. This operation severs the dependence between the variables in $\mathbf{A}$ and their parents, causing their value to be

determined by a new, exogenous mechanism. We denote the resulting probability distribution as $P_{I(\mathbf{A})}$, where P denotes the observational distribution (i.e., no intervention). All interventions in this paper are hard interventions in this sense. Where we informally speak of a “strong” intervention, we mean a hard intervention whose induced change is large enough (relative to the noise) to be detected by a finite-sample two-sample test; this signal magnitude is precisely what the parameter λ quantifies in Section 4.

2.3 FAITHFULNESS

D-separation provides a graphical criterion that constitutes a necessary but not sufficient condition for conditional independence under the causal Markov condition. We say that two variables V_i, V_j are d-separated if there is no active path between them in the causal graph. See Pearl [2009] for a more detailed description.

To provide intuition about path cancellation, we give an example with linear Gaussian structural equation models (SEMs). CI faithfulness with respect to a causal graph $\mathcal{G}$ corresponds to nonzero covariance for all d-connected pairs. Consider the following example with N_1, N_2, N_3 as independent, zero-mean, unit-variance Gaussians.

$$\begin{aligned} X_1 &= N_1, \\ X_2 &= \alpha_{12}X_1 + N_2, \\ X_3 &= \alpha_{13}X_1 + \alpha_{23}X_2 + N_3. \end{aligned} \tag{1}$$

The covariance between X_1 and X_3 is a function of the model coefficients:

$$\text{Cov}(X_1, X_3) = \alpha_{13} + \alpha_{23}\alpha_{12}.$$

A monomial of parameters emerges from each active path between X_1 and X_3 . When $\alpha_{13} + \alpha_{23}\alpha_{12} = 0$, we have $\text{Cov}(X_1, X_3) = 0$, a CI faithfulness violation. Such violations occur in two ways: (1) trivial coefficients (e.g., $\alpha_{ij} = 0$) and (2) “cancellation” of monomial terms from multiple paths.

λ -strong faithfulness [Uhler et al., 2013] requires $|\text{Cov}(V_i, V_j)| > \lambda\sqrt{\text{Var}(V_i)\text{Var}(V_j)}$ for all d-connected V_i, V_j . Geometrically, violations correspond to distributions close to the hypersurfaces defined by regular CI faithfulness violations. The argument for λ -strong faithfulness stems from finite-sample uncertainty, where true and empirical covariance differ.

Theorem 2.1 (informal, [Uhler et al., 2013]). *The volume of λ -strong CI faithfulness violations in a linear SEM with coefficients in $[-1, 1]$ is at least $\omega(\lambda^{\text{poly}(n)} 2^{|\mathbf{E}|})$ in the worst case.*

2.4 INTERVENTION FAITHFULNESS

Chevalley et al. [2025] introduced a notion of “ ε -strong intervention faithfulness” that formalizes the requirement

that intervening on a variable must induce a change on all of its descendants. Such an assumption is still vulnerable to a “cancellation of paths.” For example, in the model from Eq. (1), when $\alpha_{13} + \alpha_{23}\alpha_{12} = 0$, changing X_1 will elicit no change in X_3 , violating this assumption.

3 ASSUMPTIONS AND FRAMEWORK

To address the problem of path cancellation, we develop intervention immediacy faithfulness (II faithfulness) as a milder and more robust assumption for causal discovery. We also develop intervention locality as a complementary interventional analog to the causal Markov property.

Definition 3.1 (Most Immediate Child (MIC)). We say that $B \in \mathbf{V}$ is a “most immediate child” (MIC) of $A \in \mathbf{V}$ in $\mathcal{G} = (\mathbf{V}, \mathbf{E})$ if $B \in \text{CH}(A)$ and there are *no other directed paths* from A to B .

Here “no other directed paths” refers to *all* directed paths in $\mathcal{G}$ (of any length), not only paths of length one; equivalently, B is a MIC of A if and only if the edge $A \rightarrow B$ is the *unique* directed path from A to B .

Lemma 3.2. *In a DAG, for any vertex $V \in \mathbf{V}$, if $\text{CH}(V) \neq \emptyset$, then there is always at least one MIC.*

To see the intuition behind Lemma 3.2, consider the fact that any child with more than one path from V must have one path that is longer than length one, which means there is another child earlier in the topological ordering. Repeat this until you find the “most immediate child.” The proof is deferred to Appendix A.

II faithfulness is based on the notion of a “change” between different interventional distributions.

Definition 3.3 (Change Set). The change set $\text{CHG}(I(\mathbf{A}))$ is the set of all variables whose marginal distributions change under an intervention $I(\mathbf{A})$.

$$\text{CHG}(I(\mathbf{A})) := \{V \in \mathbf{V} \mid P_{I(\mathbf{A})}(v) \neq P(v)\}.$$

Definition 3.4 (Conditional Change Set). Let $I(\mathbf{A})$ and $I(\mathbf{B})$ be interventions on disjoint sets $\mathbf{A}$ and $\mathbf{B}$. The conditional change set of $I(\mathbf{A})$ given $I(\mathbf{B})$ is the set of variables whose distributions change when adding the intervention $I(\mathbf{A})$ to the existing intervention $I(\mathbf{B})$.

$$\text{CHG}(I(\mathbf{A}) \mid I(\mathbf{B})) := \{V \in \mathbf{V} \mid P_{I(\mathbf{A} \cup \mathbf{B})}(v) \neq P_{I(\mathbf{B})}(v)\}.$$

Conditional change sets arise naturally in sequential experimentation: a canonical example is sequential gene knockouts, where a first knockout $I(\mathbf{B})$ removes upstream regulatory effects so that the effect of a second knockout $I(\mathbf{A})$ can be attributed to the paths that remain. This mechanism is

what step 2 of our algorithms exploits. A change set may include descendants of the intervened variables, as well as the intervened variables themselves. However, it cannot include non-descendants.

Definition 3.5 (Intervention Locality). $I(\mathbf{A})$ satisfies intervention locality if its change set is restricted to the intervened variables and their descendants in $\mathcal{G}$.

$$\text{CHG}(I(\mathbf{A})) \subseteq \text{DE}_{\mathcal{G}}(\mathbf{A}) \cup \mathbf{A}.$$

Intervention locality is satisfied by the hard interventions considered throughout this paper (and, more generally, by any intervention that only modifies the mechanisms of the intervened variables); as noted below, it is implied by the SCM data-generating process. A critical property of hard interventions is the severance of relationships between the intervened variables and their causes. This breaks down potentially canceling paths, giving rise to the notion of “conditional” MICs. Throughout, a *conditioning intervention* $I(\mathbf{B})$ is a hard intervention on a disjoint set $\mathbf{B}$ that is held fixed while we probe the effect of a second intervention.

Definition 3.6 (Conditional MIC). Let $\mathcal{G} = (\mathbf{V}, \mathbf{E})$ be a graph and $I(\mathbf{B})$ be a conditioning intervention on the set $\mathbf{B} \subset \mathbf{V}$. The residual graph $\mathcal{G}_{I(\mathbf{B})}$ is the subgraph formed by removing the vertices in $\mathbf{B}$ and their incident edges. We say C is a “conditional MIC” of A given $I(\mathbf{B})$ if C is a MIC of A in the residual graph $\mathcal{G}_{I(\mathbf{B})}$.

Definition 3.7 (II Faithfulness). A model is II faithful if for any pair of disjoint intervention sets $\mathbf{A}$ and $\mathbf{B}$, all conditional MICs of any $A \in \mathbf{A}$ given $I(\mathbf{B})$ are in the corresponding conditional change set $\text{CHG}(I(\mathbf{A}) \mid I(\mathbf{B}))$.

Intervention locality and II faithfulness will make up the two main assumptions of our intervention-only causal discovery framework. Like the causal Markov condition, intervention locality is implied by the SCM data-generating process: a hard intervention on $\mathbf{A}$ replaces only the structural assignments of the variables in $\mathbf{A}$ and leaves every other mechanism intact, so a variable’s distribution can change only if it is reachable from $\mathbf{A}$ along directed paths—that is, only if it lies in $\mathbf{A} \cup \text{DE}_{\mathcal{G}}(\mathbf{A})$. II faithfulness is similar to CI faithfulness in that it must be assumed. However, this assumption is significantly weaker, as we will explain by studying it in a parametric setting.

4 PARAMETRIC II FAITHFULNESS

To analyze the parametric implications of II faithfulness, we return to the setting of linear SEMs where each variable is a linear function of its parents plus an independent noise term ε_i with $\mathbb{E}[\varepsilon_i] = 0$ and $\text{Var}(\varepsilon_i) = 1$. We can quantify II faithfulness with a λ -strong version.

Definition 4.1 (λ -strong Mean II faithfulness). A linear SEM is “ λ -strong mean II faithful” if the following holds for every variable A , every conditioning intervention $I(\mathbf{B})$ with $A \notin \mathbf{B}$, and every conditional MIC C of A given $I(\mathbf{B})$: there exists an intervention on A that shifts its mean by some non-zero amount δ such that the induced change in the mean of C (relative to the distribution under $I(\mathbf{B})$ alone) satisfies $|\mathbb{E}_{I(\{A\})|I(\mathbf{B})}[C] - \mathbb{E}_{I(\mathbf{B})}[C]| \geq \lambda|\delta|$. In particular, the intervention (and hence δ) may depend on C and $\mathbf{B}$; we do *not* require a single intervention to work uniformly across all conditioning sets.

Distributional vs. mean faithfulness. Definition 3.7 (II faithfulness) requires a change in the full *distribution* of each conditional MIC, whereas Definition 4.1 (λ -strong mean II faithfulness) constrains the *mean*. These are not equivalent, and mean II faithfulness is the strictly stronger requirement: an intervention can change a conditional MIC’s distribution while leaving its mean fixed. Our parametric analysis and our algorithmic implementation (Section 6) target mean shifts, and thus certify only mean II faithfulness. Replacing the difference-in-means test with a nonparametric two-sample test (e.g., a Kolmogorov–Smirnov or maximum-mean-discrepancy test) recovers the full distributional notion of Definition 3.7 empirically.

A similar notion can be defined for the variance, but we stick with means for simplicity, noting that this stronger assumption provides upper bounds on the “strength” of also including variance II faithfulness. λ -strong mean II faithfulness can be characterized parametrically in this setting.

Lemma 4.2. *In a linear SEM with unit-variance noise, a model is λ -strong mean II faithful if and only if every path coefficient in the model is bounded away from zero, i.e., $|\beta_{ij}| \geq \lambda$ for all $(i, j) \in \mathbf{E}$.*

This proof is deferred to Appendix A. This lemma reveals that λ -strong mean II faithfulness is parametrically equivalent to preventing any single causal link from being trivially weak. In the example from Eq. (1), this means $|\alpha_{12}| \geq \lambda$, $|\alpha_{23}| \geq \lambda$, and $|\alpha_{13}| \geq \lambda$.

A local characterization. Lemma 4.2 is inherently *local*: mean II faithfulness constrains only individual edge coefficients β_{ij} , and never sums or products of coefficients along competing paths (as in CI faithfulness). Equivalently, it suffices to require that the single coefficient β_{AC} be bounded away from zero for each node A and each conditional MIC C of A . Every edge $(A, C) \in \mathbf{E}$ is realized as such a pair, by conditioning on one internal vertex of every other $A \rightsquigarrow C$ path. Hence, this local condition coincides with the edge-wise bound $|\beta_{ij}| \geq \lambda$ for all $(i, j) \in \mathbf{E}$.

Lemma 4.3. *In a linear SEM with unit-variance noise, the set of parameterizations that violate λ -strong mean II faithfulness is a strict subset of those that violate λ -strong CI faithfulness.*

The intuition behind this proof comes from Lemma 4.2, since trivial coefficients must also result in some trivial correlations (relative to the unit variance). The full proof is deferred to Appendix A.

Finally, we can characterize the volume of II faithfulness violations.

Theorem 4.4. *The volume of λ -strong mean II faithfulness violations within the parameter space where all edge coefficients β_{ij} , $(i, j) \in \mathbf{E}$, lie in $[-1, 1]$, and for $\lambda \in (0, 1)$, is $O(\lambda|\mathbf{E}|)$.*

This linear growth in λ , when compared to Theorem 2.1, illustrates that II faithfulness is a much milder assumption than CI faithfulness.

5 ALGORITHM

In Section 3, we laid the foundation for intervention-only causal discovery using change sets and introduced two key assumptions: II faithfulness and intervention locality. We will now present an algorithm that uses this non-parametric framework to learn causal structures, identifying the conditions for complete recovery and establishing equivalence classes otherwise. For this entire section, $\mathcal{G}$ will denote the true graph with n vertices, and we will assume the model satisfies II faithfulness and intervention locality.

The identifiability of causal discovery from interventions depends on the number of variables that can be simultaneously intervened upon. We refer to this as the cardinality of the intervention, denoted $|\mathbf{A}|$ for an intervention $I(\mathbf{A})$. We will first present an algorithm for the unrestricted case and then provide versions for limited intervention cardinality.

5.1 UNRESTRICTED CARDINALITY

When the number of simultaneous interventions is not limited, the Unrestricted Intervention Cardinality (UIC) algorithm (Algorithm 1) recovers the true graph $\mathcal{G} = (\mathbf{V}, \mathbf{E})$. Step 1 performs every single-node intervention and takes the transitive closure of the observed change sets to obtain a supergraph $\mathcal{G}'$ of $\mathcal{G}$; step 2 then walks $\mathcal{G}'$ in topological order and removes each edge $V_i \rightarrow V_j$ whose endpoint is not a conditional MIC of V_i given the remaining parents.

To prove correctness, we first characterize the graph $\mathcal{G}'$ learned in step 1.

Lemma 5.1. *The graph $\mathcal{G}'$ from step 1 is the transitive closure of the true graph $\mathcal{G}$.*

The two inclusions are shown separately: $\text{TC}(\mathcal{G}_{\text{temp}}) \subseteq \text{TC}(\mathcal{G})$ follows from intervention locality (every edge added to $\mathcal{G}_{\text{temp}}$ marks a true descendant relation), and $\text{TC}(\mathcal{G}) \subseteq \text{TC}(\mathcal{G}_{\text{temp}})$ by induction in reverse topological order, using

Algorithm 1 Unrestricted Intervention Cardinality (UIC)

Require: change-set oracle $\text{CHG}(\cdot)$ obtained from hard interventions; vertex set $\mathbf{V}$.

Ensure: estimated causal DAG $\mathcal{G}''$.

```
1:  $\mathcal{G}_{temp} \leftarrow (\mathbf{V}, \emptyset)$ 
2: for each  $V_i \in \mathbf{V}$  do  $\triangleright$  Step 1: recover the transitive closure
3:   perform  $I(\{V_i\})$ 
4:   for each  $V_j \in \text{CHG}(I(\{V_i\})) \setminus \{V_i\}$  do
5:     add edge  $V_i \rightarrow V_j$  to  $\mathcal{G}_{temp}$ 
6:   end for
7: end for
8:  $\mathcal{G}' \leftarrow \text{TC}(\mathcal{G}_{temp})$ ;  $\mathcal{G}'' \leftarrow \mathcal{G}'$ 
9: for each  $V_j$  in a topological order of  $\mathcal{G}'$  do  $\triangleright$  Step 2: prune spurious edges
10:   for each  $V_i \in \text{PA}_{\mathcal{G}'}(V_j)$  do
11:      $\mathbf{B} \leftarrow \text{PA}_{\mathcal{G}''}(V_j) \setminus \{V_i\}$ 
12:     if  $V_j \notin \text{CHG}(I(\{V_i\}) \mid I(\mathbf{B}))$  then
13:       remove edge  $V_i \rightarrow V_j$  from  $\mathcal{G}''$ 
14:     end if
15:   end for
16: end for
17: return  $\mathcal{G}''$ 
```

If faithfulness to realize each true edge $V_m \rightarrow V_j$ as a path through a MIC of V_m . The full proof is given in Appendix A.

We now show that step 2 correctly prunes all spurious edges from the transitive closure without removing any true edges.

Lemma 5.2. *An edge $V_i \rightarrow V_j$ is in the true graph $\mathbf{E}$ if and only if $V_j \in \text{CHG}(I(\{V_i\}) \mid I(\text{PA}_{\mathcal{G}''}(V_j) \setminus \{V_i\}))$.*

This proof relies on the fact that a superset of the parent set of V_j is easily obtained from the transitive closure that we recovered in the previous step, which can block all paths from V_i to V_j . The full proof is given in Appendix A.

Theorem 5.3. *Assuming If faithfulness and intervention locality, the UIC algorithm learns $\mathcal{G}$ using $\mathcal{O}(n^2)$ interventions of cardinality up to $n - 1$.*

Proof. Step 1 performs n interventions. Step 2 performs at most one intervention for each edge in the transitive closure $\mathcal{G}'$, which is at most $\mathcal{O}(n^2)$. The cardinality of interventions in step 2 can be up to $n - 1$. Correctness follows from Lemmas 5.1 and 5.2. $\square$

5.2 RESTRICTED CARDINALITY

The UIC algorithm may require interventions on $n - 1$ nodes simultaneously, which is often unrealistic. We now consider the case where the cardinality of any intervention is limited, i.e., $|\mathbf{A}| \leq k$. The following k -Restricted Intervention Cardinality (k -RIC) algorithm shows that identifiability is

Algorithm 2 k -Restricted Intervention Cardinality (k -RIC)

Require: change-set oracle $\text{CHG}(\cdot)$; vertex set $\mathbf{V}$; cardinality budget k .

Ensure: estimated graph $\mathcal{G}''$ (equal to $\mathcal{G}$ when $k > \kappa_{\max}$).

```
1: run Step 1 of Algorithm 1 to obtain  $\mathcal{G}'$ ;  $\mathcal{G}'' \leftarrow \mathcal{G}'$ 
2: for each  $V_j$  in a topological order of  $\mathcal{G}'$  do
3:   for each  $V_i \in \text{PA}_{\mathcal{G}'}(V_j)$  do
4:      $\mathcal{G}^- \leftarrow \mathcal{G}''$  with edge  $V_i \rightarrow V_j$  removed
5:      $\mathbf{Z} \leftarrow$  minimum vertex cut in  $\mathcal{G}^-$  separating  $V_i$  from  $V_j$ 
6:     if  $|\mathbf{Z}| < k$  and  $V_j \notin \text{CHG}(I(\{V_i\}) \mid I(\mathbf{Z}))$  then
7:       remove edge  $V_i \rightarrow V_j$  from  $\mathcal{G}''$ 
8:     end if
9:   end for
10: end for
11: return  $\mathcal{G}''$ 
```

still possible if k is larger than the maximum local (vertex) connectivity $\kappa_{\max}$ between non-adjacent vertex pairs, i.e., the largest number of vertices needed to block all directed paths between any non-adjacent pair. Throughout, a set $\mathbf{Z}$ “separates” an ordered pair (V_i, V_j) if removing $\mathbf{Z}$ (as a hard intervention does to its targets’ incoming edges) leaves no directed path from V_i to V_j . Step 1 is identical to UIC; step 2 (Algorithm 2) differs only in that, before testing an edge $V_i \rightarrow V_j$, it finds a *minimum vertex cut* $\mathbf{Z}$ separating V_i from V_j in $\mathcal{G}''$ and conditions on $I(\mathbf{Z})$, pruning the edge when $|\mathbf{Z}| < k$ and $V_j \notin \text{CHG}(I(\{V_i\}) \mid I(\mathbf{Z}))$.

Theorem 5.4. *Assume If faithfulness and intervention locality, and let $\kappa_{\max}$ denote the maximum, over all non-adjacent ordered pairs (V_i, V_j) , of the local connectivity of the pair (the minimum number of vertices whose removal destroys all directed paths from V_i to V_j in $\mathcal{G}$). Then the k -RIC algorithm with $k > \kappa_{\max}$ learns $\mathcal{G}$ using $\mathcal{O}(n^2)$ interventions.*

The proof hinges on the topological order traversal, which guarantees that $\mathcal{G}''$ has the correct structure on the ancestors of vertex V_i when V_j is considered. Hence, we can find a vertex cut separating paths to V_j that matches the minimum vertex cut in the true graph $\mathcal{G}$.

The full proof (by case analysis on true edges and non-edges, with the non-edge case an induction in topological order) is given in Appendix A.

5.3 PARTIAL IDENTIFICATION

When the intervention cardinality k is less than or equal to the local connectivity required to separate a non-adjacent pair, we lose identifiability. For example, consider two graphs: (1) $A \rightarrow B \rightarrow C$ and (2) $A \rightarrow B \rightarrow C, A \rightarrow C$. With only single-node interventions ($k = 1$), we observe that $B, C \in \text{CHG}(I(\{A\}))$, and $C \in \text{CHG}(I(\{B\}))$. This

information produces the same transitive closure for both graphs, and we cannot distinguish them without a cardinality 2 intervention ($I(\{B\})$ compared to $I(\{A, B\})$) or a CI test ($A \perp\!\!\!\perp C \mid B$).

Definition 5.5. The k -robust transitive closure of a graph $\mathcal{G}$ is a supergraph of $\mathcal{G}$ containing an edge $V_i \rightarrow V_j$ (for i, j in topological order) if the minimum vertex cut separating V_i and V_j in $\mathcal{G}$ has cardinality k or greater.

Theorem 5.6. Assuming II faithfulness, the k -RIC algorithm recovers $\mathcal{G}$ up to its k -robust transitive closure.

Proof. The algorithm prunes a spurious edge $V_i \rightarrow V_j$ if and only if it can find a separating vertex cut $\mathbf{Z}$ of size less than k . If the minimum separating set for a non-edge has size $\geq k$, the spurious edge remains. $\square$

6 EMPIRICAL VALIDATION

Setup. We instantiate intervention-based causal discovery with a two-sample Student- t test (`scipy.stats`, significance level 0.05, Bonferroni-corrected for the number of tests) to detect a change in mean, and run the k -RIC algorithm for cardinality budgets $k \in \{1, 2, 3, 10\}$.² On 10-node graphs, $k = 10$ is unrestricted and coincides with UIC. We compare against two conditional-independence-reliant baselines. The first is the PC algorithm [Spirtes et al., 2000] (implementation of Zhang et al., 2021) run on the observational data initially, with the skeleton edges then oriented by the interventions in the usual way (“PC→Interv”). The second is IGSP [Wang et al., 2017], a joint interventional method, given the *same* single-node interventional datasets that k -RIC uses in its first step; unlike k -RIC, its permutation search still relies on observational CI tests.

We evaluate on 30 random DAGs on 10 vertices with 25 expected edges, with additive standard Gaussian noise; the edge-weight range is regime-specific and stated below. To emulate a gene knockout, an intervention severs a node from its parents and resets it to a draw from a Gaussian centered at δ with standard deviation 0.5—a hard (perfect) intervention producing a mean shift of δ . Because the exogenous noise has unit variance, δ is expressed in units of the noise standard deviation (e.g., $\delta = 5$ is a five-standard-deviation shift). Keeping a non-negligible intervention variance leaves the intervened variable well-conditioned for the Gaussian CI and invariance tests that IGSP relies on, since a near-deterministic intervention would destabilize the partial correlations it estimates. We draw 1000 observational samples and 1000 samples per intervention. All methods receive the same observational sample and the same single-node interventional datasets; k -RIC’s pruning is additionally allowed to query targeted interventions on the separating sets

it searches over (adaptive, and used only for $k \geq 2$), which is intrinsic to the method, whereas PC→Interv and IGSP are non-adaptive baselines that consume only the shared pool. We report the structural Hamming distance (SHD) between the true and recovered structures.

CI refinement. k -RIC returns the k -robust transitive closure, a supergraph of $\mathcal{G}$. We utilize a refinement step that removes the remaining spurious edges with observational CI tests: fixing a topological order of the recovered graph, we test each surviving edge $V_i \rightarrow V_j$ by conditioning on all other predecessors of V_j in that order, deleting the edge when $V_i \perp\!\!\!\perp V_j$ given that set (a partial-correlation test at level α). Conditioning on the full predecessor set blocks every directed path into V_j , so a true edge’s partial correlation stays proportional to its direct coefficient and is not eliminated by marginal cancellation; an edge is dropped only when that coefficient is too weak for the observational test to detect—the low-signal regime where interventions hold no advantage either. We apply this step in the coefficient-strength sweep below and the appendix analysis; the violin comparisons report pure k -RIC without this refinement, to isolate the cardinality effect.

Cardinality (standard regime, weights $[-1, 1]$). We first isolate the effect of the intervention cardinality k , the quantity our theory turns on. At $k = 1$, k -RIC performs no multi-node pruning and carries the spurious ancestral edges of the transitive closure. As the budget grows, multi-node interventions prune these edges and SHD falls—but the size of this effect is governed by δ . When $\delta = 1\sigma$ the induced shifts are too small for the pruning tests to fire reliably, and all methods sit at comparable SHD; as δ increases, k -RIC improves monotonically with k and drops clearly below both baselines (by $k = 2$ for $\delta \geq 5\sigma$, and by $k = 3$ at $\delta = 2\sigma$), reaching the lowest error of any method at $k = 10$ (SHD ≈ 4 at $\delta = 10\sigma$; Figure 1). The two baselines are independent of k , and both are overtaken once the interventions carry an appreciable signal. That an intervention-only method, using no CI tests, matches or beats methods that additionally exploit CI is direct evidence of the informational primacy of interventions; the advantage becomes overwhelming under path cancellation, where the CI baselines collapse (SHD ≈ 15 –19) while k -RIC recovers most of the structure (Appendix B, Figure A1).

Coefficient-strength sweep (weights $[-\omega, \omega]$). Figure 2 compares the three methods across the coefficient range $[-\omega, \omega]$ and the intervention strength δ , reporting average SHD over 20 DAGs per cell. The intervention-first pipeline— k -RIC at $k = 2$ followed by a CI-refinement step—is best across the great majority of the plane, and its margin widens with the signal: for $\omega \geq 0.5$ and $\delta \geq 5$ it drives SHD down to roughly 4–6, about half of either baseline. The conditional-independence baselines (PC→Interv and IGSP)

²Code to reproduce all experiments is available on [GitHub](#).

Baselines (PC→Interv, IGSP) vs k -RIC across cardinality $k=1, 2, 3, 10$ (Nodes=10, Edges=25, weights in $(-1, 1)$)




Figure 1: Intervention-only k -RIC across the cardinality budget $k \in \{1, 2, 3, 10\}$ versus two CI-reliant baselines, at several intervention strengths δ (weights in $[-1, 1]$; 30 DAGs, 10 nodes, 25 edges; SHD, lower is better). **PC→Interv** (grey) and **IGSP** (orange) are independent of k ; the remaining violins are k -RIC at each k (on 10 nodes, $k = 10$ is unrestricted). As δ grows, multi-node interventions ($k > 1$) prune the spurious ancestral edges that single-node interventions ($k = 1$) leave behind, and k -RIC drops below both baselines.




Figure 2: Average SHD (lower is better) of three methods across the coefficient bound ω and intervention strength δ (in noise SDs), on 10-node graphs with 25 edges (20 DAGs per cell). **Left:** the intervention-first pipeline (k -RIC at $k = 2$, then a CI-refinement step). **Middle:** IGSP. **Right:** PC→Interv. A green box marks each cell where a method is best, or within one standard error of the best. The intervention-first pipeline is best across most of the plane, ceding only the weakest-signal corner—the smallest δ or ω —to PC→Interv.

trail the pipeline across the bulk of the plane and catch up only in the weakest-signal corner—the smallest δ and smallest ω —where the induced mean shift ($\propto \omega\delta$) is too small for the interventions to resolve structure and observational CI is all that remains. The two baselines degrade differently as ω weakens—PC→Interv stays nearly flat (it orients from the interventions) while IGSP falls behind (its search leans on observational CI throughout)—as detailed in Appendix C.

7 DISCUSSION

Differing geometries of faithfulness assumptions drive different robustness to finite-sample noise. We show that causal structure remains recoverable under a much milder notion of intervention faithfulness, positioning interventions as primary rather than supporting carriers of causal information. Empirically, intervention-only k -RIC outperforms the CI-based baselines PC and IGSP, and a secondary CI-refinement step yields a complete pipeline that returns a fully

resolved DAG. Since a permutation is the weakest (0-robust) intervention equivalence class, generalizing permutation-based methods to k -robust transitive closures is a natural next step toward the intervention-first workflow.

Limitations. Our framework does not address cancellation due to latent confounders, so extending II faithfulness to ancestral graphs (ADMGs/MAGs) is important future work. Our approach also presumes that hard interventions are available. This is natural in domains such as cell biology (gene knockouts), but interventions are infeasible or ethically impermissible in many settings (e.g., much of environmental and social science), where CI-based methods remain the appropriate tool. Finally, the algorithms as implemented test for mean shifts and therefore certify only mean II faithfulness; a nonparametric two-sample test recovers the distributional notion. Soft interventions would require strengthening II faithfulness, since they do not sever a node from its parents and thus do not break canceling paths.

Acknowledgements

We thank the anonymous reviewers for their helpful feedbacks. B.M. is partially supported by the Advanced Research Concepts (ARC) COMPASS program, sponsored by the Defense Advanced Research Projects Agency (DARPA) under agreement number HR001-25-3-0212. J.Z. was partially supported by the Eric and Wendy Schmidt Center. C.U. was partially supported by NCCIH/NIH (1DP2AT012345), ONR (N00014-24-1-2687), DOE-ASCR (DE-SC0023187), and the Eric and Wendy Schmidt Center at the Broad Institute.

References

- Philippe Brouillard, Sébastien Lachapelle, Alexandre Lacoste, Simon Lacoste-Julien, and Alexandre Drouin. Differentiable Causal Discovery from Interventional Data. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 21865–21877. Curran Associates, Inc., 2020. URL https://proceedings.neurips.cc/paper_files/paper/2020/file/f8b7aa3a0d349d9562b424160ad18612-Paper.pdf.
- Hong Chang, Bin Yi, Ruixia Ma, Xiaoguo Zhang, Hongyou Zhao, and Yaguang Xi. CRISPR/cas9, a novel genomic tool to knock down microRNA in vitro and in vivo. *Scientific reports*, 6(1):22312, 2016.
- Mathieu Chevalley, Patrick Schwab, and Arash Mehrjou. Deriving causal order from single-variable interventions: Guarantees & algorithm. In *The Thirteenth International Conference on Learning Representations*, 2025. URL <https://openreview.net/forum?id=u630VngeSp>.
- Davin Choo, Kirankumar Shiragur, and Arnab Bhattacharyya. Verification and search algorithms for causal dags. *Advances in Neural Information Processing Systems*, 35:12787–12799, 2022.
- Frederick Eberhardt and Richard Scheines. Interventions and causal inference. *Philosophy of science*, 74(5):981–995, 2007.
- Chunmei Guan, Chao Ye, Xiaomei Yang, and Jiangang Gao. A review of current large-scale mouse knockout efforts. *genesis*, 48(2):73–85, 2010.
- John L Hartman IV, Barbara Garvik, and Lee Hartwell. Principles for the buffering of genetic variation. *Science*, 291(5506):1001–1004, 2001.
- Alain Hauser and Peter Bühlmann. Characterization and greedy learning of interventional markov equivalence classes of directed acyclic graphs. *The Journal of Machine Learning Research*, 13(1):2409–2464, 2012.
- Alain Hauser and Peter Bühlmann. Jointly interventional and observational data: estimation of interventional markov equivalence classes of directed acyclic graphs. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 77(1):291–318, 2015.
- Yang-Bo He and Zhi Geng. Active Learning of Causal Networks with Intervention Experiments and Optimal Designs. *Journal of Machine Learning Research*, 9(84):2523–2547, 2008. URL <http://jmlr.org/papers/v9/he08a.html>.
- Amin Jaber, Adele Ribeiro, Jiji Zhang, and Elias Bareinboim. Causal Identification under Markov equivalence: Calculus, Algorithm, and Completeness. In S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh, editors, *Advances in Neural Information Processing Systems*, volume 35, pages 3679–3690. Curran Associates, Inc., 2022.
- Dominik Janzing and Bernhard Schölkopf. Causal inference using the algorithmic markov condition. *IEEE Transactions on Information Theory*, 56(10):5168–5194, 2010.
- Murat Kocaoglu. Characterization and learning of causal graphs with small conditioning sets. *Advances in Neural Information Processing Systems*, 36, 2023.
- Murat Kocaoglu, Amin Jaber, Karthikeyan Shanmugam, and Elias Bareinboim. Characterization and Learning of Causal Graphs with Latent Variables from Soft Interventions. In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.
- Christopher Meek. Strong completeness and faithfulness in Bayesian networks. In *Proceedings of the Eleventh conference on Uncertainty in artificial intelligence*, UAI’95, pages 411–418, San Francisco, CA, USA, August 1995. Morgan Kaufmann Publishers Inc. ISBN 978-1-55860-385-1. URL <https://dl.acm.org/doi/10.5555/2074158.2074205>.
- Marc Franquesa Monés, Jiaqi Zhang, and Caroline Uhler. On the number of conditional independence tests in constraint-based causal discovery. *arXiv preprint arXiv:2603.21844*, 2026.
- Joris M. Mooij, Sara Magliacane, and Tom Claassen. Joint causal inference from multiple contexts. *The Journal of Machine Learning Research*, 21(1):99:3919–99:4026, January 2020. ISSN 1532-4435. URL <https://dl.acm.org/doi/10.5555/3455716.3455815>.
- Judea Pearl. Graphs, causality, and structural equation models. *Sociological Methods & Research*, 27(2):226–284, 1998.

- Judea Pearl. *Causality*. Cambridge university press, 2009.
- Joseph Ramsey, Jiji Zhang, and Peter L Spirtes. Adjacency-faithfulness and conservative causal inference. *arXiv preprint arXiv:1206.6843*, 2012.
- Numair Sani, Jaron Lee, and Ilya Shpitser. Identification and Estimation of Causal Effects Defined by Shift Interventions. In *Proceedings of the 36th Conference on Uncertainty in Artificial Intelligence (UAI)*, pages 949–958. PMLR, August 2020. URL <https://proceedings.mlr.press/v124/sani20a.html>.
- Karthikeyan Shanmugam, Murat Kocaoglu, Alexandros G Dimakis, and Sriram Vishwanath. Learning causal graphs with small interventions. *Advances in Neural Information Processing Systems*, 28, 2015.
- Shohei Shimizu, Patrik O. Hoyer, Aapo Hyvärinen, and Antti Kerminen. A Linear Non-Gaussian Acyclic Model for Causal Discovery. *Journal of Machine Learning Research*, 7(72):2003–2030, 2006. ISSN 1533-7928. URL <http://jmlr.org/papers/v7/shimizu06a.html>.
- Kirankumar Shiragur, Jiaqi Zhang, and Caroline Uhler. Causal discovery with fewer conditional independence tests. *arXiv preprint arXiv:2406.01823*, 2024.
- Peter Spirtes, Clark N Glymour, and Richard Scheines. *Causation, prediction, and search*. MIT press, 2000.
- Chandler Squires and Caroline Uhler. Causal structure learning: A combinatorial perspective. *Foundations of Computational Mathematics*, 23(5):1781–1815, 2023.
- Chandler Squires, Yuhao Wang, and Caroline Uhler. Permutation-based causal structure learning with unknown intervention targets. In *Conference on Uncertainty in Artificial Intelligence*, pages 1039–1048. PMLR, 2020.
- Caroline Uhler, Garvesh Raskutti, Peter Bühlmann, and Bin Yu. Geometry of the faithfulness assumption in causal inference. *The Annals of Statistics*, pages 436–463, 2013.
- Matthew J Vowels, Necati Cihan Camgoz, and Richard Bowden. D’ya like dags? a survey on structure learning and causal discovery. *ACM Computing Surveys*, 55(4):1–36, 2022.
- Yuhao Wang, Liam Solus, Karren Yang, and Caroline Uhler. Permutation-based causal inference algorithms with interventions. *Advances in Neural Information Processing Systems*, 30, 2017.
- Sascha Xu, Sarah Mameche, and Jilles Vreeken. Information-theoretic causal discovery in topological order. In *The 28th International Conference on Artificial Intelligence and Statistics*, 2025. URL <https://openreview.net/forum?id=9pjJXQWYXc>.
- Karren Yang, Abigail Katcoff, and Caroline Uhler. Characterizing and learning equivalence classes of causal dags under interventions. In *International Conference on Machine Learning*, pages 5541–5550. PMLR, 2018.
- Jiji Zhang and Peter L Spirtes. Strong faithfulness and uniform consistency in causal inference. *arXiv preprint arXiv:1212.2506*, 2012.
- Keli Zhang, Shengyu Zhu, Marcus Kalander, Ignavier Ng, Junjian Ye, Zhitang Chen, and Lujia Pan. gcastle: A python toolbox for causal discovery, 2021.
- Zihan Zhou, Muhammad Qasim Elahi, and Murat Kocaoglu. Characterization and Learning of Causal Graphs from Hard Interventions. In D. Belgrave, C. Zhang, H. Lin, R. Pascanu, P. Koniusz, M. Ghassemi, and N. Chen, editors, *Advances in Neural Information Processing Systems*, volume 38, pages 77620–77663. Curran Associates, Inc., 2025.

Relaxing Faithfulness with Intervention-Only Causal Discovery (Supplementary Material)

Bijan H. S. Mazaheri^{1,2}

Jiaqi Zhang^{2,3}

Caroline Uhler^{2,3}

¹Thayer School of Engineering, Dartmouth College, Hanover, New Hampshire, USA

²Schmidt Center, Broad Institute of MIT and Harvard, Cambridge, Massachusetts, USA

³Laboratory for Information and Decision Systems, Massachusetts Institute of Technology, Cambridge, Massachusetts, USA

A DEFERRED PROOFS

A.1 PROOF OF LEMMA 3.2

For any vertex $V \in \mathbf{V}$, if $\mathbf{CH}(V) \neq \emptyset$, then there is always at least one MIC.

We assume $\mathcal{G}$ is a Directed Acyclic Graph (DAG). As $\mathcal{G}$ is a DAG, there exists a topological ordering of its vertices. Let's denote the position of a vertex X in such an ordering by $\text{ord}(X)$. A key property of a topological ordering is that if there is a directed path from X to Y , then $\text{ord}(X) < \text{ord}(Y)$.

Let V be a vertex such that its set of children, $\mathbf{CH}(V)$, is non-empty. Consider the set of topological ordering positions for all children of V :

$$S = \{\text{ord}(C) \mid C \in \mathbf{CH}(V)\}$$

Since $\mathbf{CH}(V)$ is a finite, non-empty set, S is a finite, non-empty set of integers. Therefore, S must contain a minimum element.

Let $B \in \mathbf{CH}(V)$ be a child of V that achieves this minimum, i.e., $\text{ord}(B) = \min(S)$. We claim that B is a most immediate child of V .

Suppose for the sake of contradiction that B is *not* a most immediate child. By definition, this means there must exist another directed path from V to B besides the direct edge $V \rightarrow B$. Any such path must be of length at least 2 and must pass through another child of V . Let this path be:

$$V \rightarrow C \rightarrow \dots \rightarrow B$$

where $C \in \mathbf{CH}(V)$ and $C \neq B$.

The existence of a directed path from C to B implies that $\text{ord}(C) < \text{ord}(B)$. However, C is a child of V , so $\text{ord}(C) \in S$. We have found an element in S that is strictly smaller than $\min(S)$, which is a contradiction.

Therefore, our assumption that B is not a most immediate child must be false. It follows that B is a most immediate child, and such a child must always exist if $\mathbf{CH}(V)$ is non-empty.

A.2 PROOF OF LEMMA 4.2

In a linear SEM with unit-variance noise, a model is λ -strong mean II faithful if and only if every path coefficient in the model is bounded away from zero, i.e., $|\beta_{ij}| \geq \lambda$ for all $(i, j) \in \mathbf{E}$.

($\Rightarrow$) Assume the model is λ -strong mean II faithful. Consider an arbitrary edge $(A, C) \in \mathbf{E}$ with coefficient β_{AC} . Let $\mathcal{P}_{AC}$ be the set of all other directed paths from A to C . For each path $p \in \mathcal{P}_{AC}$, select one intermediate vertex V_p along that path. Let the set of all such selected vertices be $\mathbf{B} = \{V_p \mid p \in \mathcal{P}_{AC}\}$.

Now consider the intervention $I(\mathbf{B})$. In the residual graph $\mathcal{G}_{I(\mathbf{B})}$, all alternative directed paths from A to C are removed, leaving only the direct edge. Therefore, C is a conditional MIC of A given $I(\mathbf{B})$. By the definition of λ -strong mean II faithfulness, an intervention on A must produce a detectable change on C . In a linear model, this change is governed by the total effect in the residual graph, which is simply β_{AC} . The faithfulness condition thus requires $|\beta_{AC}| \geq \lambda$. Since (A, C) was an arbitrary edge, this holds for all edges in $\mathbf{E}$.

($\Leftarrow$) Assume $|\beta_{ij}| \geq \lambda$ for all edges $(i, j) \in \mathbf{E}$. Consider any disjoint intervention sets $I(\mathbf{A})$ and $I(\mathbf{B})$, and let C be a conditional MIC of some $A \in \mathbf{A}$ given $I(\mathbf{B})$. By definition, this means that in the residual graph $\mathcal{G}_{I(\mathbf{B})}$, the only directed path from A to C is the direct edge (A, C) .

In a linear model, the change in the mean of C from an intervention on A (that shifts its mean by δ) is determined by the total effect of A on C in the active graph, multiplied by δ . Here, the total effect is simply the path coefficient of the single connecting path, β_{AC} . The change in mean is thus $\beta_{AC}\delta$. The condition for λ -strong mean II faithfulness requires $|\beta_{AC}\delta| \geq \lambda|\delta|$, which simplifies to $|\beta_{AC}| \geq \lambda$. By our initial assumption, this condition is met. Since this holds for any valid choice of $A, C, \mathbf{B}$, the model is λ -strong mean II faithful.

A.3 PROOF OF LEMMA 4.3

In a linear SEM with unit-variance noise, the set of parameterizations that violate λ -strong mean II faithfulness is a strict subset of those that violate λ -strong CI faithfulness.

By Lemma 4.2, a violation of λ -strong mean II faithfulness is equivalent to $|\beta_{ij}| < \lambda$ for some edge $(i, j) \in \mathbf{E}$. Consider the adjacent pair (V_i, V_j) and condition on $\mathbf{S} = \mathbf{PA}_{\mathcal{G}}(V_j) \setminus \{V_i\}$. In a linear Gaussian SEM, the population regression coefficient of V_i in the regression of V_j on $\{V_i\} \cup \mathbf{S}$ equals the structural coefficient β_{ij} (since $V_j = \sum_{P \in \mathbf{PA}_{\mathcal{G}}(V_j)} \beta_{Pj}P + \varepsilon_j$ with ε_j independent of the parents), and the partial correlation between V_i and V_j given $\mathbf{S}$ is

$$\rho(V_i, V_j \mid \mathbf{S}) = \beta_{ij} \sqrt{\frac{\text{Var}(V_i \mid \mathbf{S})}{\text{Var}(V_j \mid \mathbf{S})}}.$$

Because all conditional variances are finite and strictly positive in a non-degenerate model, the factor multiplying β_{ij} is bounded and nonzero; hence $|\rho(V_i, V_j \mid \mathbf{S})|$ is proportional to $|\beta_{ij}|$ and is likewise small whenever $|\beta_{ij}|$ is small. Since V_i and V_j are adjacent (hence d-connected given $\mathbf{S}$), a sufficiently small $|\beta_{ij}|$ drives this partial correlation below the corresponding λ -strong CI faithfulness threshold, i.e., a λ -strong CI faithfulness violation. Thus every λ -strong mean II faithfulness violation is also a λ -strong CI faithfulness violation. Crucially, the converse fails: a model can violate CI faithfulness through cancellation (e.g., $\alpha_{13} + \alpha_{12}\alpha_{23} = 0$ in Eq. (1)) while every edge coefficient remains large, hence satisfying mean II faithfulness. Therefore the set of mean-II-unfaithful models is *strictly* contained in the set of CI-unfaithful models.

A.4 PROOF OF THEOREM 4.4

Proof. Let the parameter space be the hypercube $[-1, 1]^{|\mathbf{E}|}$. By Lemma 4.2, a λ -strong mean II faithfulness violation occurs if and only if $|\beta_{ij}| < \lambda$ for at least one edge $(i, j) \in \mathbf{E}$. For a single coefficient, the volume of violating parameters is the region where it lies in $(-\lambda, \lambda)$, which has a relative volume of λ . By applying a union bound over all $|\mathbf{E}|$ edges, the relative volume of the violation set is at most $\mathcal{O}(\lambda|\mathbf{E}|)$. $\square$

A.5 PROOF OF LEMMA 5.1

The graph $\mathcal{G}'$ from step 1 is the transitive closure of the true graph $\mathcal{G}$.

Proof. Let $\mathcal{G}_{temp}$ be the graph constructed from the single-node interventions before taking the transitive closure. We need to prove that $\text{TC}(\mathcal{G}_{temp}) = \text{TC}(\mathcal{G})$.

First, we show $\text{TC}(\mathcal{G}_{temp}) \subseteq \text{TC}(\mathcal{G})$. An edge $V_i \rightarrow V_j$ is added to $\mathcal{G}_{temp}$ only if $V_j \in \text{CHG}(I(\{V_i\}))$. By intervention locality, this implies V_j is a descendant of V_i in $\mathcal{G}$. Therefore, every edge in $\mathcal{G}_{temp}$ corresponds to a path in $\mathcal{G}$, which means $\mathcal{G}_{temp} \subseteq \text{TC}(\mathcal{G})$. The transitive closure of a subgraph is a subgraph of the original's transitive closure, so $\text{TC}(\mathcal{G}_{temp}) \subseteq \text{TC}(\mathcal{G})$.

Second, we show $\text{TC}(\mathcal{G}) \subseteq \text{TC}(\mathcal{G}_{temp})$. This is equivalent to showing that $\mathcal{G} \subseteq \text{TC}(\mathcal{G}_{temp})$. We proceed by induction over the source vertices in reverse topological order, showing that for the vertex currently under consideration, each of its outgoing edges $V_i \rightarrow V_j$ in $\mathcal{G}$ is realized as a directed path from V_i to V_j in $\mathcal{G}_{temp}$.

Let the vertices be ordered $V_1, \dots, V_n$ according to a reverse topological sort of $\mathcal{G}$.

Base case: For the first vertex V_1 (a sink in $\mathcal{G}$), the claim is vacuously true as it has no outgoing edges.

Inductive step: Assume for all $\ell < m$, the claim holds for V_ℓ . Now consider V_m . Let V_j be any child of V_m in $\mathcal{G}$. If V_j is a MIC of V_m , then by II faithfulness, $V_j \in \text{CHG}(I(\{V_m\}))$, so the edge $V_m \rightarrow V_j$ is in $\mathcal{G}_{temp}$. If V_j is not a MIC of V_m , then there exists another path from V_m to V_j in $\mathcal{G}$, which must pass through another child of V_m . Let this path be $V_m \rightarrow U_1 \rightarrow \dots \rightarrow V_j$. Since U_1 appears after V_m in the reverse topological sort, our inductive hypothesis applies to it and all subsequent nodes on the path to V_j . This means there is a path $U_1 \rightsquigarrow V_j$ in $\mathcal{G}_{temp}$. This reduces our problem to showing there is a path $V_m \rightsquigarrow U_1$ in $\mathcal{G}_{temp}$. This is the same problem we started with for V_j . However, we can construct a sequence of children $V_j, U_1, U_2, \dots$ in which each child is required for a non-MIC path from V_m to the preceding child. Since the number of children is finite, this sequence must terminate, and it must terminate at a child that is a MIC of V_m . Let this MIC be M . By II faithfulness, the edge $V_m \rightarrow M$ is in $\mathcal{G}_{temp}$. The path from M to V_j exists in $\mathcal{G}$, and all nodes on it are covered by the inductive hypothesis, so there is a path $M \rightsquigarrow V_j$ in $\mathcal{G}_{temp}$. Concatenating these, we get a path $V_m \rightarrow M \rightsquigarrow V_j$ in $\mathcal{G}_{temp}$. Thus, for any edge $V_m \rightarrow V_j$ in $\mathcal{G}$, there is a path in $\mathcal{G}_{temp}$. This means $\mathcal{G} \subseteq \text{TC}(\mathcal{G}_{temp})$, which proves $\text{TC}(\mathcal{G}) \subseteq \text{TC}(\mathcal{G}_{temp})$. $\square$

A.6 PROOF OF LEMMA 5.2

An edge $V_i \rightarrow V_j$ is in the true graph $\mathbf{E}$ if and only if $V_j \in \text{CHG}(I(\{V_i\}) \mid I(\mathbf{PA}_{\mathcal{G}''}(V_j) \setminus \{V_i\}))$.

($\Rightarrow$) Assume $V_i \rightarrow V_j$ is in $\mathbf{E}$. Let $\mathbf{B} = \mathbf{PA}_{\mathcal{G}''}(V_j) \setminus \{V_i\}$. When we intervene on $\mathbf{B}$, all paths into V_j from its parents (other than V_i) are blocked. Because $\mathcal{G}'$ is the transitive closure, $\mathbf{PA}_{\mathcal{G}'}(V_j)$ contains all ancestors of V_j , so intervening on $\mathbf{B}$ blocks all paths from V_i to V_j that do not begin with the edge $V_i \rightarrow V_j$. Thus, V_j becomes a conditional MIC of V_i . By II faithfulness, $V_j \in \text{CHG}(I(\{V_i\}) \mid I(\mathbf{B}))$, so the edge is correctly kept.

($\Leftarrow$) Assume $V_i \rightarrow V_j$ is not in $\mathbf{E}$. Let $\mathbf{B} = \mathbf{PA}_{\mathcal{G}''}(V_j) \setminus \{V_i\}$. Since V_i is not a true parent of V_j , the set of true parents $\mathbf{PA}_{\mathcal{G}}(V_j)$ is a subset of $\mathbf{B}$. Any directed path from V_i to V_j in $\mathcal{G}$ must pass through at least one node in $\mathbf{PA}_{\mathcal{G}}(V_j)$. By intervening on $\mathbf{B}$, all such paths are blocked. In the residual graph, V_j is not a descendant of V_i . By intervention locality, $V_j \notin \text{CHG}(I(\{V_i\}) \mid I(\mathbf{B}))$, so the edge is correctly removed.

A.7 PROOF OF THEOREM 5.4

Proof. The algorithm's correctness hinges on Step 2 correctly deciding whether to keep or remove each candidate edge $V_i \rightarrow V_j$ from the transitive closure graph $\mathcal{G}'$. We analyze this for true edges and non-edges separately.

Case 1: $V_i \rightarrow V_j$ is a true edge in $\mathcal{G}$. The algorithm considers the edge $V_i \rightarrow V_j$ and hypothetically removes it to form $\mathcal{G}^-$. It then searches for a minimum vertex cut $\mathbf{Z}$ in $\mathcal{G}^-$ that separates V_i from V_j . Two outcomes are possible:

1. No separating set $\mathbf{Z}$ with $|\mathbf{Z}| < k$ is found. The algorithm correctly keeps the edge.
2. A separating set $\mathbf{Z}$ with $|\mathbf{Z}| < k$ is found. The algorithm then tests if $V_j \in \text{CHG}(I(\{V_i\}) \mid I(\mathbf{Z}))$. By definition, the set $\mathbf{Z}$ blocks all paths from V_i to V_j in $\mathcal{G}$ *except for the direct edge* $V_i \rightarrow V_j$. In the context of the conditioning intervention $I(\mathbf{Z})$, the only remaining path from V_i to V_j is the true edge itself. This makes V_j a conditional MIC of V_i given $I(\mathbf{Z})$. By the II faithfulness assumption, V_j must be in the change set. The condition $V_j \in \text{CHG}(\dots)$ is met, and the algorithm correctly keeps the edge.

In either outcome, a true edge is never removed.

Case 2: $V_i \rightarrow V_j$ is not a true edge in $\mathcal{G}$. We need to show that the algorithm will find a separating set $\mathbf{Z}$ with $|\mathbf{Z}| < k$ and subsequently remove the edge. We prove this by induction on the position of V_j in the topological order used by the algorithm. Without loss of generality, relabel the vertices $V_1, \dots, V_n$ so that this topological order is $V_1, \dots, V_n$.

Inductive Hypothesis (IH): Assume that for every vertex V_m with $m < j$, all incoming edges to V_m in $\mathcal{G}''$ have been correctly identified; that is, the subgraph of $\mathcal{G}''$ induced by $\{V_1, \dots, V_{j-1}\}$ is identical to the subgraph of $\mathcal{G}$ induced by the same set of vertices.

Inductive Step: Consider the vertex V_j and a candidate edge $V_i \rightarrow V_j$ that is in $\mathcal{G}'$ but not in $\mathcal{G}$ (so $i < j$). Since $V_i \rightarrow V_j$ is not a true edge, (V_i, V_j) is a non-adjacent pair, and there exists a minimum vertex cut $\mathbf{Z}_{\text{true}}$ in $\mathcal{G}$ that destroys all directed paths from V_i to V_j . By Menger’s theorem, the size of this cut equals the *local* connectivity of the pair (V_i, V_j) —the maximum number of internally vertex-disjoint directed $V_i \rightsquigarrow V_j$ paths. This local quantity is at most κ_{max} by definition, and by the theorem’s premise $k > \kappa_{\text{max}}$, so $|\mathbf{Z}_{\text{true}}| \leq \kappa_{\text{max}} < k$.

Now, we must show that $\mathbf{Z}_{\text{true}}$ also separates V_i and V_j in the current state of $\mathcal{G}''$ (with the edge $V_i \rightarrow V_j$ hypothetically removed). Assume for contradiction that it does not. This means there must be a path P'' from V_i to V_j in $\mathcal{G}''$ that avoids $\mathbf{Z}_{\text{true}}$. Since $\mathbf{Z}_{\text{true}}$ separates them in $\mathcal{G}$, this path P'' must contain at least one spurious edge (an edge in $\mathcal{G}''$ but not in $\mathcal{G}$). Let (U, W) be the first such edge on path P'' .

The vertices U and W must both precede V_j in the topological ordering. Therefore, $U, W \in \{V_1, \dots, V_{j-1}\}$. By our Inductive Hypothesis, the subgraph of $\mathcal{G}''$ induced by these vertices is already identical to that of $\mathcal{G}$. This means the spurious edge (U, W) cannot exist, a contradiction.

Therefore, $\mathbf{Z}_{\text{true}}$ is a valid vertex cut for (V_i, V_j) in the current graph $\mathcal{G}''$. The algorithm searches for a *minimum* vertex cut, $\mathbf{Z}_{\text{alg}}$, whose size must be less than or equal to the size of our valid cut: $|\mathbf{Z}_{\text{alg}}| \leq |\mathbf{Z}_{\text{true}}| < k$.

The condition $|\mathbf{Z}_{\text{alg}}| < k$ is met. The algorithm tests if $V_j \in \text{CHG}(I(\{V_i\}) \mid I(\mathbf{Z}_{\text{alg}}))$. Since $\mathbf{Z}_{\text{alg}}$ is a vertex cut separating V_i and V_j in the true graph $\mathcal{G}$, by intervention locality, V_j will not be in the change set. The edge $V_i \rightarrow V_j$ is correctly removed.

By induction, this holds for all vertices and all spurious edges are removed. $\square$

B PATH-CANCELLATION COMPARISON

Figure A1 repeats the cardinality comparison of Figure 1 in the path-cancellation regime (edge weights in $[-0.1, 0.1]$), the stress test for CI-based methods.

Baselines (PC→Interv, IGSP) vs k -RIC across cardinality $k=1, 2, 3, 10$ (Nodes=10, Edges=25, weights in $(-0.1, 0.1)$)




Figure A1: The cardinality comparison under significant risk of path cancellation (weights in $[-0.1, 0.1]$; 30 DAGs, 10 nodes, 25 edges). The small, sign-varied coefficients make partial correlations cancel, leaving PC→Interv and IGSP near-uninformative at every δ , whereas intervention-only k -RIC, which uses no independence tests, recovers most of the structure once the intervention strength is sufficient ($\delta \gtrsim 5\sigma$).

Small, sign-varied coefficients make partial correlations cancel, so the conditional-independence baselines PC→Interv and IGSP stay near-uninformative ($\text{SHD} \approx 15\text{--}19$) at every δ . k -RIC never tests independencies and is immune to this cancellation; it is limited only by whether the (now small) intervention shifts are detectable. Once $\delta \gtrsim 2\sigma$ makes them visible, even the crudest intervention-only estimate already beats both baselines, and by $\delta = 5\sigma$ k -RIC recovers most of the structure ($\text{SHD} \approx 6$, roughly a third of the baselines’ error), reaching $\text{SHD} \approx 4\text{--}5$ at $\delta = 10\sigma$. Here the cardinality budget k barely matters—the tiny coefficients leave the multi-node pruning tests underpowered, so extra cardinality buys little—which is exactly why we isolate the cardinality effect in the more benign standard regime of Figure 1, where the signal is strong enough for pruning to bite.

The δ -dependence of the k -RIC violins in Figure A1 reflects the power of the underlying two-sample tests. When δ is small relative to the noise, the tests are underpowered: they miss true MIC shifts (causing missed edges) and fail to certify

redundant edges as prunable (causing under-pruning), which both raise and widen the SHD distribution and account for the residual error even at large k . As δ grows, both failure modes recede and the distributions concentrate at low SHD. This is precisely the λ -strong mean II-faithfulness regime of Section 4: the detectable mean shift at a MIC scales with the product of the intervention strength δ and the edge coefficient, so shrinking δ shrinks the effective λ that is satisfied (Figure 2).

C COEFFICIENT-STRENGTH DEPENDENCE OF THE BASELINES

Figure 2 reports mean SHD across the (ω, δ) plane; Figure A2 isolates the coefficient-strength dependence by fixing $\delta = 5\sigma$ and plotting SHD against ω for the three methods. PC $\rightarrow$ Interv is nearly flat in ω : it recovers the skeleton from observational conditional independence but orients every edge directly from the interventional shifts, which stay strong at $\delta = 5\sigma$ regardless of ω . IGSP, whose permutation search relies on observational CI tests throughout, degrades steadily as ω shrinks and those tests lose power. The intervention-first pipeline is lowest across the range. This is the same mechanism, now at the level of whole algorithms, that motivates preferring interventions to conditional independence: the more a method’s decisions rest on CI tests, the more it suffers as the parametric signal weakens.




Figure A2: Average SHD versus the coefficient bound ω at fixed intervention strength $\delta = 5\sigma$ (10-node graphs, 25 edges; shaded bands are ± 1 standard error). PC $\rightarrow$ Interv stays roughly flat because it orients edges from the interventions; IGSP degrades as ω weakens because its permutation search depends on observational CI tests; the intervention-first pipeline (k -RIC + CI, $k = 2$) is lowest throughout.