Research Article

AND DATA SCIENCES
BULLETIN OF COMPUTER

Interventional DN2CN: Causal Structure Learning from Mixed Observational and Experimental Data with Context-Specific Independence

Mobeen Akhter

INTI International University, Malaysia

Abstract

Learning causal Bayesian networks (CBNs) from data is challenging when the available information consists of a mixture of observational and interventional samples collected under heterogeneous experimental conditions. Standard structure learning algorithms either ignore interventions or treat them only as hard background constraints, and most of them do not exploit context-specific independence (CSI) patterns that arise in realistic high-dimensional domains. In this paper we introduce *Interventional DN2CN*, a two-stage method that extends dependency-network-to-causal-network (DN2CN) approaches to the mixed-regime setting. In the first stage, we learn a regime-aware dependency network whose local conditional distributions are represented by shallow decision trees that explicitly condition on both parent variables and an intervention indicator. This representation captures CSI and allows us to detect invariance and change of local mechanisms across different experimental regimes. In the second stage, we convert the (generally cyclic) dependency network into an acyclic causal Bayesian network by (i) removing edges that are incompatible with intervention targets and regime-specific invariance constraints, and (ii) orienting remaining undirected edges using a combination of mutual information, stability across regimes, and simple intervention-based orientation rules. We formalize the problem setting, describe the algorithm, and discuss identifiability conditions that arise from combining CSI with intervention information. We then outline an empirical evaluation on both synthetic benchmarks and a real biological network, comparing Interventional DN2CN with purely observational DN2CN and with classical constraint-based and score-based learners that incorporate intervention targets. Our results design aims to show that explicitly modeling regime-dependent context-specific structure improves both structural recovery and causal effect estimation, particularly for nodes adjacent to manipulated variables. We conclude by discussing limitations and potential extensions to dynamic settings and latent variable models.

Keywords: causal Bayesian network learning, mixed observational and interventional data, context-specific independence, dependency-to-causal network conversion, regime-aware causal discovery

1. Introduction

Causal Bayesian networks (CBNs) provide a powerful graphical language for representing cause–effect relationships among a collection of random variables and for reasoning about the effects of

Corresponding author: Mobeen Akhter (mobeen2398@gmail.com)

BCDS 2024, Vol. 5(3), 35-48. DOI: 10.71448/bcds2453-3

interventions [1–3]. When randomized experiments are expensive or ethically constrained, researchers often rely on structure learning methods that infer causal graphs from observational data alone. However, in many scientific domains — such as molecular biology, clinical research, or online platform experimentation — data are collected under a mixture of observational and interventional regimes. Some units are observed passively, while others are subject to targeted manipulations of a subset of variables.

Existing structure learning algorithms for mixed observational and experimental data typically fall into two main families. Constraint-based approaches extend conditional independence tests with background knowledge about intervention targets, while score-based approaches augment their likelihood with interventional likelihood terms [4–6]. In both cases, interventions are treated as constraints or as additional samples, but the algorithms usually operate on homogeneous parametric families (e.g., multinomial or Gaussian conditionals) and do not exploit *context-specific independence* (CSI) — the phenomenon that a variable may become independent of some of its parents in certain value configurations of other parents [7, 8].

Parallel to this literature, dependency networks (DNs) have been proposed as flexible models in which each variable is associated with a local conditional distribution given its Markov blanket, potentially leading to directed cycles [9]. When local conditionals are represented by decision trees, DNs naturally encode CSI and are easy to learn in parallel. Recent work has shown that, under suitable conditions, a DN can be converted into an acyclic causal network (CN) by removing edges within strongly connected components using mutual information (MI) and then orienting the remaining edges according to local Markov structure. This DN2CN strategy combines the flexibility of DNs with the interpretability of CBNs.

However, existing DN2CN approaches are designed for purely observational data. They ignore experimental regimes, do not use do-calculus, and treat all samples as coming from a single, stationary data-generating process. This leaves several questions open:

- How can we extend DN2CN to explicitly model multiple regimes, including interventions that manipulate specific variables?
- Can we use regime-specific CSI and invariance patterns to obtain stronger orientation rules than those available in purely observational settings?
- What are the practical benefits of combining DNs, CSI, and intervention targets for structural recovery and causal effect estimation?

In this paper we address these questions by proposing *Interventional DN2CN*, a method for causal structure learning from mixed observational and experimental data. Our main contributions are: We formalize a problem setting in which data are collected under multiple regimes, some of which correspond to atomic interventions on subsets of variables, and we introduce a regime indicator that can be incorporated into local DN conditionals; we propose a regime-aware DN learning procedure in which each variable is modeled by a decision tree conditional on its Markov blanket and on the regime indicator. This allows the model to capture CSI and to detect invariance and heterogeneity of local mechanisms across regimes; we develop a two-stage DN2CN conversion procedure that uses both mutual information and regime-specific invariance constraints to remove cycles and orient edges. In particular, we employ simple intervention-based orientation rules that leverage observed changes in local conditionals when targets are manipulated.; we sketch an empirical evaluation protocol on synthetic and real-world benchmarks, comparing Interventional DN2CN with baseline methods that either ignore CSI or do not explicitly model interventions, and discuss how structure recovery and causal effect estimation may benefit from our approach.

The remainder of this paper is organized as follows. Section 2 reviews causal Bayesian networks, interventions, dependency networks, and context-specific independence. Section 3 defines the mixed-regime problem setting. Section 4 introduces the Interventional DN2CN algorithm. Section 5 describes the experimental design. Section 6 discusses implications and limitations, and Section 7 concludes.

2. Background

2.1. Causal Bayesian networks and interventions

Let $V = \{X_1, \ldots, X_p\}$ be a set of random variables with joint distribution $P(V)$. A causal Bayesian network (CBN) consists of a directed acyclic graph (DAG) $G = (V, E)$ together with a set of conditional probability distributions (CPDs) $\{P(X_j \mid \mathrm{Pa}_G(X_j))\}_{j=1}^{p}$, where $\mathrm{Pa}_G(X_j)$ denotes the parents of X_j in G [2, 3, 10]. The joint distribution factorizes as

$$P(V) = \prod_{j=1}^{p} P(X_j \mid \mathrm{Pa}_G(X_j)). \tag{1}$$

In a causal interpretation, each CPD encodes a local causal mechanism, and interventions correspond to modifications of specific CPDs. For an atomic intervention $\mathrm{do}(X_k = x)$, the interventional distribution is

$$P(V \mid \mathrm{do}(X_k = x)) = \delta(X_k = x) \prod_{j \neq k} P(X_j \mid \mathrm{Pa}_G(X_j)), \tag{2}$$

where $\delta(\cdot)$ is a point mass. More generally, interventions may set variables to stochastic policies or modify their mechanisms in other ways [2, 10].

In practice, we often do not know the true DAG G and must infer it from data. Classical constraint-based and score-based structure learning methods exploit conditional independence tests or decomposable scores to search over DAGs that are Markov equivalent to the underlying data-generating distribution [3, 11, 12]. Observations collected under different regimes correspond to different interventional distributions $P_r(V)$, where r indexes the regime (e.g., observational, intervention on X_k, joint interventions, etc.). Structure learning from such mixed data requires modeling both the common causal structure and the regime-specific modifications of local mechanisms and has motivated a rich literature on causal discovery under interventions and heterogeneous environments [13, 14].

2.2. Dependency networks and context-specific independence

A dependency network (DN) [9] associates each variable X_j with a local conditional distribution $P(X_j \mid \mathcal{N}(X_j))$ given a neighborhood $\mathcal{N}(X_j) \subseteq V \setminus \{X_j\}$. Unlike CBNs, DNs do not require global acyclicity; the directed graph induced by neighborhoods may contain cycles. Under mild regularity conditions, the local conditionals define a unique joint distribution that is the stationary distribution of a Markov chain that sequentially updates each variable given its neighbors [9].

When each local conditional is represented by a decision tree or decision graph, DNs can compactly represent context-specific independence (CSI) [7, 8]. Intuitively, CSI occurs when X_j is independent of some subset of its neighbors in certain contexts (value configurations) of other neighbors. For example, X_j may depend on X_k only when $X_\ell = 1$, but not when $X_\ell = 0$. In a decision tree representation, this is captured by splitting on X_ℓ and including X_k as a predictor only in the subtree corresponding to $X_\ell = 1$. CSI has been widely exploited to obtain compact and interpretable representations of local mechanisms in Bayesian networks and related models [15, 16], and the same ideas carry over directly to the DN setting.

The DN framework is attractive for structure learning because local conditionals can be learned independently using standard supervised learning methods. When the predictors are discrete and the regression model is a decision tree, learning is efficient and naturally parallelizable, and the resulting models can be combined into a coherent joint distribution without enforcing acyclicity during the local learning phase [9].

2.3. From dependency networks to causal networks

Recent work has shown that DNs can serve as an intermediate representation for learning causal networks. The basic idea is to learn a DN in which each variable is predicted from its neighbors using a decision tree, interpret the directed edges as candidate causal dependencies while allowing for cycles, remove edges within strongly connected components (SCCs) using a criterion such as mutual information (MI) between parent and child to break cycles, and then orient any remaining undirected edges according to local Markov properties, yielding an acyclic causal network (CN). This DN2CN approach leverages the flexibility of DNs and the interpretability of CBNs, and fits into a broader line of work that uses rich local predictors as a surrogate for explicit global causal structure during the early stages of learning [10, 14]. However, existing DN2CN methods are typically designed for a single, stationary regime and do not explicitly utilize information from interventions or environment shifts.

2.4. Structure learning with interventional data

Existing methods for structure learning from mixed observational and interventional data often extend Bayesian or constraint-based approaches. For example, [4] integrate interventional data into a Bayesian score, while [5] characterize the effect of interventions on Markov equivalence classes and develop greedy search procedures such as GIES. Earlier work by [17] and [6] studies the combinatorial and statistical trade-offs involved in choosing optimal interventions for causal discovery, and more recent contributions explore learning from general interventions or soft manipulations using graphical and optimization-based methods [13, 18]. In general, knowing which variables were manipulated in each regime can reduce Markov equivalence and aid in edge orientation, because descendants of intervention targets may change distribution while non-descendants remain invariant [14].

However, most of these methods assume homogeneous parametric CPDs, such as multinomial or linear-Gaussian families, and do not exploit CSI in their local models. Our goal is to combine the strengths of DN2CN—flexible local models with context-specific structure—with the additional information provided by known interventions, thereby obtaining causal graphs that better leverage both observational dependence patterns and interventional invariance across regimes.

3. Problem setting

We consider a collection of p discrete random variables $V = \{X_1, \ldots, X_p\}$ whose joint state space is the Cartesian product $\mathcal{X}_1 \times \ldots \times \mathcal{X}_p$. These variables are observed under a finite set of data-generating regimes $\mathcal{R} = \{0, 1, \ldots, R\}$. The index $r = 0$ is reserved for the purely observational regime, in which no external manipulation is applied and the system evolves according to its natural dynamics, whereas each $r \geq 1$ corresponds to an interventional regime associated with a particular experimental condition. For a fixed regime $r \in \mathcal{R}$, we observe n_r independent and identically distributed samples

$$\mathcal{D}_r = \{x^{(i)} : i = 1, \ldots, n_r\}, \quad x^{(i)} \in \mathcal{X}_1 \times \ldots \times \mathcal{X}_p,$$

where each $x^{(i)}$ represents a complete assignment to all variables in V recorded under regime r.

To keep track of the regime that generated each sample, we introduce a regime indicator random variable R taking values in $\mathcal{R}$. Conceptually, R encodes whether a given observation comes from the baseline observational condition or from one of the experimental conditions. For each interventional regime $r \geq 1$, we assume that we know which subset of variables is directly targeted by the intervention, and we denote this set by $T_r \subseteq V$. The observational regime has no targets, so we define $T_0 = \emptyset$. Knowing the targets means that for each experimental condition r we are informed which variables are being manipulated by the experimenter (for example, forced to specific values or otherwise externally controlled), but we do not initially commit to a particular parametric form of that manipulation.

We assume that, behind all these regimes, there is a single underlying causal structure governing the relationships among the variables. This structure is represented by a directed acyclic graph (DAG) $G^\star$ on the vertex set V. Each directed edge $X_k \to X_j$ in $G^\star$ is interpreted causally, indicating that X_k is a direct cause of X_j. The parents of a variable X_j in the true graph are denoted by $\mathrm{Pa}_{G^\star}(X_j)$, and the corresponding conditional distribution $P(X_j \mid \mathrm{Pa}_{G^\star}(X_j))$ describes the local causal mechanism by which X_j responds to its parents. Crucially, we assume that this causal graph $G^\star$ is invariant across all regimes: interventions may change the local mechanisms associated with some nodes, but they do not create or destroy causal edges between variables.

Formally, for each regime $r \in \mathcal{R}$ we assume that the joint distribution over V factorizes according to the same DAG $G^\star$:

$$P_r(V) = \prod_{j=1}^{p} P_r(X_j \mid \mathrm{Pa}_{G^\star}(X_j)). \tag{3}$$

The subscript r on the conditional distributions emphasizes that the local mechanism of each variable is allowed to depend on the regime. However, we impose a mechanism invariance assumption for non-target variables: if a variable X_j is not directly manipulated in regime r, i.e., $j \notin T_r$, then its conditional distribution coincides with its observational conditional. In symbols,

$$j \notin T_r \quad \Longrightarrow \quad P_r(X_j \mid \mathrm{Pa}_{G^\star}(X_j)) = P_0(X_j \mid \mathrm{Pa}_{G^\star}(X_j)). \tag{4}$$

Intuitively, this expresses the idea that interventions act locally on the mechanisms of their direct targets while leaving the causal laws of non-target variables unchanged. For variables that are targets, $j \in T_r$, we do not restrict how their mechanisms may change: the conditional distributions $P_r(X_j \mid \mathrm{Pa}_{G^\star}(X_j))$ can differ arbitrarily from $P_0(X_j \mid \mathrm{Pa}_{G^\star}(X_j))$, thereby covering both hard interventions that fix X_j to specific values and softer or stochastic interventions that modify its dependence on its parents.

The complete dataset available to the learner consists of the union of all regime-specific datasets, enriched with their regime labels,

$$\mathcal{D} = \bigcup_{r \in \mathcal{R}} \{(x^{(i)}, r) : x^{(i)} \in \mathcal{D}_r\}, \tag{5}$$

as well as the collection of target sets $\{T_r : r \in \mathcal{R}\}$. Thus, for each individual observation we know both the realized values of all variables in V and the regime in which it was generated, and for each regime we know which variables were directly intervened upon. We deliberately do not assume any additional knowledge about the functional form or parametric structure of the interventions themselves; our only structural assumption is that interventions may arbitrarily change the local mechanisms of their targets while leaving non-target mechanisms invariant and preserving the underlying causal DAG.

Under this problem setting, the learning task is to recover the true causal structure $G^\star$ as faithfully as possible from the mixed observational and interventional data. In practice, owing to Markov equivalence and finite-sample limitations, this means identifying the Markov equivalence class (or, when interventions provide enough information, a refined equivalence class) of DAGs that are compatible with the observed conditional independence structure and with the assumed intervention targets across regimes. The central question addressed by our method is how to use the combined information in the regime indicator R, the target sets $\{T_r\}$, and the observed distributions $P_r(V)$ to infer a causal graph that best represents the invariant and regime-specific aspects of the data-generating process.

4. Interventional DN2CN

This section introduces the Interventional DN2CN algorithm. The method proceeds in three stages. First, we learn a regime-aware dependency network using decision trees that condition on both neighbors and the regime indicator R. Second, we remove edges to obtain an acyclic skeleton by combining mutual information with regime-specific invariance constraints and intervention targets. Third, we orient remaining undirected edges using local Markov structure, intervention-based rules, and context-specific independence (CSI) patterns.

4.1. Stage 1: Learning a regime-aware dependency network

For each variable X_j we define a neighborhood $\mathcal{N}(X_j) \subseteq V \setminus \{X_j\}$, initially taken to be $V \setminus \{X_j\}$ or obtained from an initial screening step (for example, mutual information thresholding). We then learn a predictive model of the form

$$P(X_j \mid \mathcal{N}(X_j), R), \tag{6}$$

using a decision tree classifier trained on the pooled data across all regimes:

$$\mathcal{D}^{(j)} = \bigcup_{r \in \mathcal{R}} \{(x^{(i)}_{\mathcal{N}(X_j)}, r, x^{(i)}_j) : x^{(i)} \in \mathcal{D}_r\}.$$

The tree is grown using standard impurity-based splits, such as information gain or Gini reduction, and at internal nodes we allow splits both on predictors in $\mathcal{N}(X_j)$ and on the regime indicator R. Splits on R partition the data into subtrees corresponding to different regimes, allowing the model to learn distinct local mechanisms in interventional regimes. Splits on predictors within each regime-specific subtree capture CSI within that regime, because some predictors may only be relevant in particular contexts determined by other variables or by the regime. Leaves then correspond to contexts $(x_{\mathcal{N}(X_j)}, r)$ in which X_j has a specific categorical distribution.

From the learned tree for X_j we extract a set of candidate parents $\widehat{\mathrm{Pa}}(X_j)$ by collecting all variables (excluding R) that appear in internal splits. This defines a directed edge set

$$E^{\mathrm{DN}} = \{(X_k, X_j) : X_k \in \widehat{\mathrm{Pa}}(X_j)\}. \tag{7}$$

The resulting directed graph $G^{\mathrm{DN}} = (V, E^{\mathrm{DN}})$ is a dependency network and may contain cycles.

4.2. Stage 2: Removing edges using MI and invariance constraints

To obtain an acyclic skeleton suitable for a causal Bayesian network, we must remove edges that participate in directed cycles while preserving the most informative and causally plausible dependencies. We build on the DN2CN idea of using mutual information (MI), but augment it with regime-specific invariance constraints.

For each directed edge $(X_k, X_j) \in E^{\mathrm{DN}}$, we estimate the conditional mutual information

$$I(X_k; X_j \mid \widehat{\mathrm{Pa}}(X_j) \setminus \{X_k\}, R), \tag{8}$$

using the empirical joint distribution over $(X_j, X_k, \widehat{\mathrm{Pa}}(X_j) \setminus \{X_k\}, R)$, smoothed as needed. Intuitively, this quantity measures the residual dependence between X_k and X_j after conditioning on other parents and on the regime, and thus captures how strongly X_k contributes unique predictive information about X_j.

To incorporate intervention information, we define an invariance score that quantifies how stable the local conditional $P(X_j \mid \widehat{\mathrm{Pa}}(X_j), R)$ is across regimes in which X_j is not directly manipulated. Let $\mathcal{R}_j^{\mathrm{non}} = \{r \in \mathcal{R} : X_j \notin T_r\}$ be the set of regimes where X_j is not targeted. For each such regime r, we can estimate the conditional $P_r(X_j \mid \widehat{\mathrm{Pa}}(X_j))$ by restricting the data to $\mathcal{D}_r$. We then define an invariance score for edge (X_k, X_j) as

$$\mathrm{Inv}(X_k \to X_j) = - \sum_{r, r' \in \mathcal{R}_j^{\mathrm{non}}} w_{r,r'} \, D_{\mathrm{KL}} \left(P_r(X_j \mid \widehat{\mathrm{Pa}}(X_j)) \| P_{r'}(X_j \mid \widehat{\mathrm{Pa}}(X_j)) \right), \tag{9}$$

where $w_{r,r'}$ are non-negative weights, for example normalized by sample size, and D_{KL} is the Kullback–Leibler divergence. High (that is, less negative) values of Inv indicate that the local mechanism for X_j is approximately invariant across non-targeting regimes, which is consistent with X_j being a non-intervention target and with the corresponding parents representing stable causes. The score can be refined to focus more directly on the role of X_k by comparing models with and without X_k included as a parent, or by computing invariance only on subtrees of the decision tree that involve X_k.

For each edge we define a combined score

$$S(X_k \to X_j) = \alpha \, I(X_k; X_j \mid \widehat{\mathrm{Pa}}(X_j) \setminus \{X_k\}, R) + (1 - \alpha) \, \mathrm{Inv}(X_k \to X_j), \tag{10}$$

where $\alpha \in [0, 1]$ balances raw predictive strength and mechanism invariance. Edges with higher values of S are considered more important and more consistent with the underlying causal structure, because they both carry strong conditional dependence and participate in mechanisms that remain stable across non-target regimes.

To break cycles, we identify strongly connected components (SCCs) in G^{DN} and then process each SCC that contains at least one cycle. Within such a component, we iteratively remove the edge with the lowest combined score $S(X_k \to X_j)$, update the SCC decomposition, and repeat until the component becomes acyclic. This procedure yields an acyclic directed graph $G^{\mathrm{acy}} = (V, E^{\mathrm{acy}})$ that defines the skeleton for the final causal Bayesian network.

4.3. Stage 3: Orienting edges using interventions and CSI

After cycle removal, some edges may effectively be undirected in the sense that they appear as mutual parents in local trees or that their orientation is not fully determined by the dependency network alone. At this stage we exploit intervention targets and regime-specific changes in local mechanisms to orient the remaining edges.

One guiding principle is that intervention targets behave as exogenous variables within regimes in which they are manipulated. If a variable X_k is directly manipulated in regime r, so that $X_k \in T_r$, then in that regime its distribution is determined by the intervention rather than by its usual parents, and thus it should not have incoming causal edges from variables whose mechanisms remain unchanged. In practice, we enforce a weaker rule: for any variable X_k that is targeted in at least one regime and any edge (X_j, X_k) with low mutual information and a low invariance score, we favor orienting the edge as $X_k \to X_j$ or removing it entirely, rather than maintaining the orientation $X_j \to X_k$.

A second principle is based on the effect of interventions on conditional distributions. Suppose we intervene on X_k in regime r and compare the conditional distributions of another variable X_j across regimes,

$$P_r(X_j \mid \mathbf{Z}) \quad \text{vs.} \quad P_{r'}(X_j \mid \mathbf{Z}), \tag{11}$$

for some subset $\mathbf{Z}$ of other variables and for a non-interventional regime r'. If the distribution of X_j changes significantly when X_k is manipulated, while the distribution of X_k remains essentially unaffected by manipulations of X_j, this asymmetry supports the orientation $X_k \to X_j$. Operationally, we approximate this reasoning by measuring regime-specific shifts in the decision tree for X_j along paths that include splits on X_k, and by comparing them to corresponding shifts in the tree for X_k when X_j is targeted.

CSI in the learned trees provides additional orientation cues. For instance, if the tree for X_j contains a split on X_k only in the context where $X_\ell = x_\ell$, and we also observe a split on X_ℓ in the tree for X_k, this pattern may indicate a context-specific v-structure $X_k \to X_j \leftarrow X_\ell$ in that region of the state space. While CSI alone does not fully determine the presence or direction of v-structures, combining CSI patterns with intervention-based invariance information can strengthen orientation decisions and help distinguish genuine converging causal paths from spurious associations.

We implement the orientation stage via local rules applied iteratively. We begin by initializing the graph with edges from E^{acy} and a partial orientation based on the directions induced by the dependency network. For each pair of adjacent nodes (X_k, X_j) whose orientation remains ambiguous, we compute statistics that summarize regime-specific shifts in their local decision trees and apply the intervention-based orientation principles described above. For triples (X_k, X_j, X_ℓ) that form an undirected chain, we examine CSI patterns consistent with v-structure formation and orient the edges accordingly whenever the data support such a configuration. Throughout this process we enforce acyclicity by discarding candidate orientations that would reintroduce directed cycles, preferentially keeping those orientations that are supported by higher combined scores S and stronger invariance evidence. The result of this stage is a DAG $\widehat{G}$ and a set of local conditionals $\{\widehat{P}(X_j \mid \mathrm{Pa}_{\widehat{G}}(X_j), R)\}$ that together define an interventional causal Bayesian network model.

Algorithm 1 Interventional DN2CN

Require: Mixed dataset $\mathcal{D} = \bigcup_{r\in\mathcal{R}} \mathcal{D}_r$, regime indicator R, intervention targets $\{T_r\}$
Ensure: Estimated causal DAG $\widehat{G}$

Initialize empty parent sets $\widehat{\mathrm{Pa}}(X_j)$ for all $X_j \in V$.
for each variable $X_j \in V$ **do**
    Learn a decision tree model for $P(X_j \mid \mathcal{N}(X_j), R)$ using the pooled data $\mathcal{D}^{(j)} = \bigcup_{r\in\mathcal{R}}\{(x^{(i)}_{\mathcal{N}(X_j)}, r, x^{(i)}_j) : x^{(i)} \in \mathcal{D}_r\}$.
    Extract candidate parents $\widehat{\mathrm{Pa}}(X_j)$ as the set of variables (excluding R) that appear in internal splits of the tree.
Construct the dependency network $G^{\mathrm{DN}} = (V, E^{\mathrm{DN}})$ with edges $(X_k, X_j) \in E^{\mathrm{DN}}$ whenever $X_k \in \widehat{\mathrm{Pa}}(X_j)$.
for each edge $(X_k, X_j) \in E^{\mathrm{DN}}$ **do**
    Compute conditional mutual information $I(X_k; X_j \mid \widehat{\mathrm{Pa}}(X_j) \setminus \{X_k\}, R)$.
    Compute invariance score $\mathrm{Inv}(X_k \to X_j)$ based on regime-specific conditionals $P_r(X_j \mid \widehat{\mathrm{Pa}}(X_j))$ for $r \in \mathcal{R}^{\mathrm{non}}_j$.
    Combine them into an edge score $S(X_k \to X_j) = \alpha\, I(X_k; X_j \mid \widehat{\mathrm{Pa}}(X_j) \setminus \{X_k\}, R) + (1-\alpha)\,\mathrm{Inv}(X_k \to X_j)$.
Identify strongly connected components (SCCs) of G^{DN}.
for each SCC that contains at least one directed cycle **do**
    while the SCC contains a directed cycle **do**
        Find the edge (X_k, X_j) in the SCC with the smallest score $S(X_k \to X_j)$.
        Remove (X_k, X_j) from E^{DN}.
        Update the SCC decomposition of G^{DN}.
Let $G^{\mathrm{acy}} = (V, E^{\mathrm{acy}})$ be the resulting acyclic graph.
Apply intervention-based and CSI-based orientation rules to ambiguous edges in G^{acy}, enforcing acyclicity, to obtain the final DAG $\widehat{G}$.
Optionally refit local conditional distributions $P(X_j \mid \mathrm{Pa}_{\widehat{G}}(X_j), R)$ for all $X_j \in V$ under the learned parent sets.

4.4. *Algorithm summary*

Algorithm 1 summarizes the overall procedure. At a high level, the algorithm takes as input the mixed dataset $\mathcal{D} = \bigcup_r \mathcal{D}_r$, the regime indicator R, and the collection of intervention targets $\{T_r\}$, and produces as output an estimated causal DAG $\widehat{G}$. For each variable X_j, it first learns a decision tree approximating the conditional distribution $P(X_j \mid \mathcal{N}(X_j), R)$ and extracts candidate parents $\widehat{\mathrm{Pa}}(X_j)$ from the variables that appear in internal splits. These parent sets are used to construct the initial dependency network G^{DN} with edges (X_k, X_j) whenever $X_k \in \widehat{\mathrm{Pa}}(X_j)$. For each such edge, the algorithm then computes the conditional mutual information $I(X_k; X_j \mid \widehat{\mathrm{Pa}}(X_j) \setminus \{X_k\}, R)$ and the invariance score $\mathrm{Inv}(X_k \to X_j)$, and combines them into the edge score $S(X_k \to X_j)$. Strongly connected components in G^{DN} are identified, and edges with the lowest scores are iteratively removed until all components are acyclic, yielding the acyclic graph G^{acy}. Finally, intervention- and CSI-based orientation rules are applied to ambiguous edges in G^{acy}, while maintaining acyclicity, to obtain the final DAG $\widehat{G}$. In a last optional step, local conditional distributions $P(X_j \mid \mathrm{Pa}_{\widehat{G}}(X_j), R)$ can be refitted under the learned parent sets in order to improve the quantitative accuracy of the model.

5. Experimental design

This section outlines an empirical evaluation protocol for Interventional DN2CN. We describe synthetic benchmarks, a real-world dataset, baselines, and evaluation metrics. Concrete numerical results should be added once experiments have been carried out.

5.1. Synthetic benchmarks

To systematically assess structural recovery and causal effect estimation, we consider a synthetic data generation scheme in which we first generate random DAGs over $p \in \{10, 20, 50\}$ variables using Erdos–Renyi or scale-free graph models with a specified expected in-degree. For each such DAG, we then parameterize discrete CPDs either by sampling probability tables from Dirichlet distributions or by constructing decision-tree-based mechanisms that explicitly incorporate context-specific independence. On top of these observational models, we define a collection of interventional regimes by selecting subsets of variables T_r and modifying their local mechanisms according to hard or soft interventions; examples include fixing the value distributions of targets or altering selected rows of their CPDs. From these models we sample a mixture of observational and interventional data with regime-specific sample sizes n_r, chosen to reflect realistic imbalance between observational and experimental conditions.

Within this synthetic framework, we systematically vary several factors to probe the robustness of Interventional DN2CN and competing methods. In particular, we vary the number and size of the intervention target sets, ranging from single-node interventions to larger joint manipulations; we alter the proportion of interventional samples relative to purely observational samples, thereby simulating settings with abundant or scarce experimental data; and we adjust the strength and prevalence of context-specific independence in the generating mechanisms, moving from nearly context-free models to highly context-dependent ones. These variations allow us to examine how performance scales with graph size, intervention design, and the degree to which CSI is present in the underlying causal processes.

5.2. Real-world dataset

For real data, a natural choice is a biological signaling network dataset with known interventions, such as the flow cytometry measurements of protein signaling in human T cells under various biochemical perturbations [19]. This dataset includes both observational and interventional conditions and is accompanied by a widely used reference network, making it a standard benchmark for causal discovery methods. In this setting, different perturbation conditions play the role of regimes, with documented intervention targets, and Interventional DN2CN can be applied directly by treating the condition label as the regime indicator and using the known targets as T_r.

5.3. Baselines

We compare Interventional DN2CN with several baseline methods that represent different ways of combining observational and interventional information. One baseline is DN2CN applied only to the observational subset of the data, ignoring interventional samples; this variant evaluates how much is gained by explicitly modeling interventions. A second baseline is DN2CN applied to the pooled data without any regime modeling, effectively treating interventional observations as if they were ordinary samples from a single stationary distribution. We also include a constraint-based interventional learner, such as an extension of PC or FCI that incorporates knowledge of intervention targets as background constraints on edge orientation. In addition, we consider a score-based interventional learner, for example a Bayesian or BIC-scored search algorithm that augments its likelihood with interventional likelihood terms, such as GIES [5]. Where computationally feasible, we further include a non-linear additive model approach that combines generalized additive models with ideas from invariant causal prediction to exploit interventional invariance at the level of conditional distributions.

5.4. Evaluation metrics

We evaluate the methods along three main dimensions: structural recovery, causal effect estimation, and predictive performance. For structural recovery, and in settings where ground-truth graphs are available (either synthetic or real with a reference network), we compute the Structural Hamming Distance (SHD) between the estimated DAG and the true DAG, thereby quantifying the total number of edge additions, deletions, and reversals required to reconcile the learned structure with the ground truth. We also report precision, recall, and F1 score for directed edges to capture the trade-off between false positives and false negatives in edge identification. In addition, we measure the proportion of correctly oriented edges adjacent to intervention targets, since these edges are often the most directly affected by interventional information and thus provide a focused view of orientation quality.

For causal effect estimation, we select a subset of ordered pairs (X_k, X_j) and estimate average causal effects (ACEs) under interventions using the learned DAG and its associated CPDs. These estimated effects are then compared to ground truth effects, either computed analytically from the generative models in the synthetic experiments or approximated using strong interventional regimes in the real dataset. We summarize performance using metrics such as mean absolute error (MAE) and root mean squared error (RMSE) of ACE estimates, and, in Bayesian variants, we also examine the coverage properties of credible intervals for ACEs to assess the calibration of uncertainty quantification.

As a final diagnostic, we examine predictive performance. Specifically, we evaluate the predictive log-likelihood or related scores on held-out data, stratified by regime, in order to verify that the learned models provide an adequate fit to the joint distributions observed in both observational and interventional settings. While predictive performance is not a direct measure of causal correctness, it helps identify gross model misspecification and ensures that structural advantages are not achieved at the expense of a poor overall data fit.

6. Discussion

Interventional DN2CN combines three sources of information that are often treated separately in causal discovery: observational conditional dependence, interventional perturbations, and context-specific independence. By modeling local mechanisms with regime-aware decision trees, the method captures fine-grained patterns of invariance and change across regimes, which can be exploited to orient edges and resolve ambiguities in Markov equivalence classes.

Conceptually, the approach is closely aligned with the invariance-based view of causality, according to which causal relations correspond to stable mechanisms that remain robust across a range of interventions and environments [13]. In our setting, invariance is assessed at the level of local conditionals extracted from decision trees, with interventions encoded via splits on the regime indicator. This perspective suggests several avenues for theoretical development, such as characterizing conditions under which the combination of CSI and regime-specific invariance identifies the true DAG or at least narrows down the equivalence class beyond what is possible with observational data alone.

From a practical standpoint, the dependency network stage offers computational benefits. Local decision trees can be learned independently and in parallel, which is attractive for high-dimensional problems. The main additional cost introduced by Interventional DN2CN arises from computing regime-specific invariance scores and from applying more elaborate orientation rules. These costs are still manageable for moderate numbers of variables and regimes, especially when trees are shallow and regimes are reasonably balanced.

Several limitations should be noted. First, our method currently assumes discrete variables and relies on decision trees. Extending the approach to continuous or mixed-type data would require using other local models (e.g., Gaussian mixtures or generalized additive models) and more sophisticated estimators of MI and KL divergence. Second, we assume that intervention targets are known and that interventions act locally on CPDs; violations of these assumptions (e.g., unknown off-target effects, soft interventions) may degrade performance. Third, we do not explicitly model latent confounding; in the presence of unobserved common causes, the learned DAG may spuriously attribute direct causal effects where only associations exist.

Finally, while we have focused on static causal networks, the same ideas could be extended to dynamic settings. Learning a regime-aware dependency network over time-indexed variables and converting it into a dynamic Bayesian network (DBN) could enable Interventional DN2CN to model feedback loops that unfold across time, thus enriching its applicability to temporal systems.

7. Conclusion

We have proposed Interventional DN2CN, a method for learning causal Bayesian networks from mixed observational and experimental data by leveraging regime-aware dependency networks with context-specific independence. By treating the regime indicator as an explicit predictor in local decision trees, the method captures how interventions modify local mechanisms, and by combining mutual information with invariance-based scoring, it removes cycles and orients edges in a way that respects both statistical dependence and causal constraints.

The proposed algorithm provides a flexible and interpretable framework that can serve as a bridge between pure observational DN2CN approaches and classical interventional structure learners. Future work includes extending the method to continuous and mixed data, incorporating latent variable structure, and developing formal identifiability guarantees under realistic assumptions about CSI and intervention design. Empirically validating the method on a range of domains with rich interventional structure will further clarify its strengths and limitations.

References

[1] Ramanan, N., & Natarajan, S. (2020). Causal learning from predictive modeling for observational data. *Frontiers in big Data, 3*, 535976.

[2] Pearl, J. (2009).Neuberg, L. G. (2003). Causality: models, reasoning, and inference, by judea pearl, cambridge university press, 2000. *Econometric Theory, 19*(4), 675-685.

[3] Spirtes, P., Glymour, C. N., & Scheines, R. (2000). *Causation, Prediction, and Search*. MIT press.

[4] Eaton, D., & Murphy, K. (2007, March). Exact Bayesian structure learning from uncertain interventions. In *Artificial Intelligence and Statistics* (pp. 107-114). PMLR.

[5] Hauser, A., & Bühlmann, P. (2012). Characterization and greedy learning of interventional Markov equivalence classes of directed acyclic graphs. *The Journal of Machine Learning Research, 13*(1), 2409-2464.

[6] He, Y. B., & Geng, Z. (2008). Active learning of causal networks with intervention experiments and optimal designs. *Journal of Machine Learning Research, 9*(11), 2523–2547.

[7] Boutilier, C., Friedman, N., Goldszmidt, M., & Koller, D. (2013). Context-specific independence in Bayesian networks. arXiv preprint arXiv:1302.3562. Morgan Kaufmann.

[8] Poole, D., & Zhang, N. L. (2003). Exploiting contextual independence in probabilistic inference. *Journal of Artificial Intelligence Research, 18*, 263-313.

[9] Heckerman, D., Chickering, D. M., Meek, C., Rounthwaite, R., & Kadie, C. (2000). Dependency networks for inference, collaborative filtering, and data visualization. *Journal of Machine Learning Research, 1*(Oct), 49-75.

[10] Shanmugam, R. (2018). *Elements of Causal Inference: Foundations and Learning Algorithms.*

[11] Chickering, D. M. (2002). Optimal structure identification with greedy search. *Journal of Machine Learning Research, 3*(Nov), 507-554.

[12] Tsamardinos, I., Brown, L. E., & Aliferis, C. F. (2006). The max-min hill-climbing Bayesian network structure learning algorithm. *Machine Learning, 65*(1), 31-78.

[13] Peters, J., Bühlmann, P., & Meinshausen, N. (2016). Causal inference by using invariant prediction: identification and confidence intervals. *Journal of the Royal Statistical Society Series B: Statistical Methodology, 78*(5), 947-1012.

[14] Mooij, J. M., Peters, J., Janzing, D., Zscheischler, J., & Schölkopf, B. (2016). Distinguishing cause from effect using observational data: methods and benchmarks. *Journal of Machine Learning Research, 17*(32), 1-102.

[15] Gao, T., Fadnis, K., & Campbell, M. (2017, July). Local-to-global Bayesian network structure learning. In *International Conference on Machine Learning* (pp. 1193-1202). PMLR.

[16] Chickering, D. M., Heckerman, D., & Meek, C. (2013). A Bayesian approach to learning Bayesian networks with local structure. arXiv preprint arXiv:1302.1528. Morgan Kaufmann.

[17] Eberhardt, F., & Scheines, R. (2007). Interventions and causal inference. *Philosophy of Science, 74*(5), 981-995.

[18] Hyttinen, A., Hoyer, P. O., Eberhardt, F., & Jarvisalo, M. (2013). Discovering cyclic causal models with latent variables: A general SAT-based procedure. arXiv preprint arXiv:1309.6836.

[19] Sachs, K., Perez, O., Pe'er, D., Lauffenburger, D. A., & Nolan, G. P. (2005). Causal protein-signaling networks derived from multiparameter single-cell data. *Science, 308*(5721), 523-529.

How to cite this article: Mobeen Akhter (2024). Interventional DN2CN: Causal Structure Learning from Mixed Observational and Experimental Data with Context-Specific Independence. *Bulletin of Computer and Data Sciences*, 5(3), 35-48. DOI: 10.71448/bcds2453-3

Received: 06/05/2024 **Revised:** 19/07/2024 **Accepted:** 22/08/2024 **Publish:** 30/09/2024

BULLETIN OF COMPUTER AND DATA SCIENCES *Bulletin of Computer and Data Sciences* is a peer-reviewed open access journal.