TECHNICAL REPORT
R-122
Feb, 2025

Structural Causal Bandits under Markov Equivalence

Min Woo Park [1] **Andy Arditi** [2] **Elias Bareinboim** [2] **Sanghack Lee** [1]

Abstract

In a decision-making process, intelligent agents with causal knowledge can optimize action spaces to avoid unnecessary exploration. A *structural causal bandit* framework provides guidance on how to prune actions that are unable to maximize reward by leveraging prior knowledge of the underlying causal structure among actions. A key assumption of this framework is that the agent has access to a fully-specified causal diagram representing the target system. However, this assumption is often violated, making it difficult to apply in real world. In this paper, we extend the multi-armed bandits framework with structural causal models, also known as structural causal bandits, to scenarios where the agent leverages a Markov equivalent class. In such cases, the causal structure is provided to the agent in the form of maximal ancestral graphs (MAGs) or partial ancestral graphs (PAGs). We propose a generalized framework for identifying potentially optimal actions within these graph structures, thereby broadening the applicability of structural causal bandits to real-world settings.

1. Introduction

The multi-armed bandit (MAB) (Robbins, 1952; Lai and Robbins, 1985; Lattimore and Szepesvári, 2020) problem is a central topic in decision-making studies, where an agent aims to maximize cumulative rewards by repeatedly choosing actions (or pulling arms) based on observed reward, balancing the exploration-exploitation trade-off. Traditionally, MAB problems assume independence among the rewards of different arms, meaning that the reward obtained from one arm provides no information about the others e.g., KL-UCB (Cappé et al., 2013) and Thompson sampling (Thompson, 1933). Although this independence assumption simplifies the problem, it limits its applicability to real-world scenarios where dependencies among actions are common, e.g., in a movie recommendation system, the positive reaction of a user to one genre can indicate a higher likelihood of a positive reaction to similar genres.

Recent research has increasingly recognized the importance of structured dependencies among arms and reward (Li et al., 2010; Abbasi-Yadkori et al., 2011; Cesa-Bianchi and Lugosi, 2012), leading to the development of structured bandits. Concurrently, the integration of causal inference into the MAB framework has opened new avenues for modeling and solving decision problems with richer dependency structures (Bareinboim et al., 2021). Causal diagrams (Pearl, 1995) have been employed to represent causal relationships among actions, rewards, and other relevant factors. This approach enables agents to make informed decisions by considering how each action causally influences the reward through causal pathways. Existing studies (Bareinboim et al., 2015; Lattimore et al., 2016; Forney et al., 2017) have shown that causality-aware strategies can significantly outperform MAB algorithms that do not account for such underlying causal relationships. Subsequent work has explored various specialized settings by introducing additional structural assumptions, such as the availability of both observational and experimental distributions, or linear mechanisms (Zhang and Bareinboim, 2017; Lu et al., 2020; Bilodeau et al., 2022; Feng and Chen, 2023; Varici et al., 2023).

Specifically, Lee and Bareinboim (2018) formalized the *structural causal bandit* (SCM-MAB) without any parametric assumptions, where causal dependencies between arms are modeled using a structural causal model (SCM) (Pearl, 2000). They proposed a sound and complete graphical characterization to identify *minimal intervention sets* (MISs) and *possibly-optimal minimal intervention sets* (POMISs), where the former includes only the variables that affect the reward, and the latter refers to actions that could be part of an optimal strategy among MISs, thereby guiding the agent to avoid unnecessary exploration without any actual interaction. Lee and Bareinboim (2019) extended this approach to accommodate scenarios involving non-manipulable variables among all the variables in the graph. Lee and Bareinboim (2020); Everitt et al. (2021) established SCM-MAB

[1]Graduate School of Data Science, Seoul National University, Seoul, South Korea [2]Columbia University. Correspondence to: Sanghack Lee <sanghack@snu.ac.kr>, Elias Bareinboim <eb@cs.columbia.edu>.

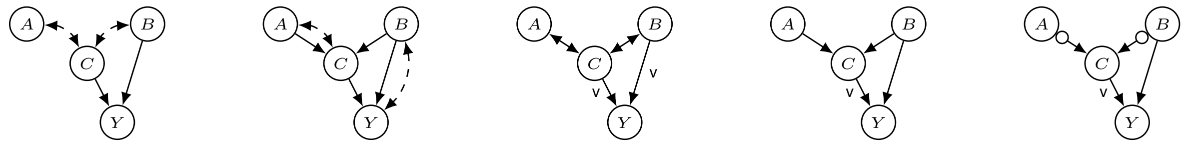

(a) $\mathcal{G}_1$ (b) $\mathcal{G}_2$ (c) $\mathcal{M}_1$ (d) $\mathcal{M}_2$ (e) $\mathcal{P}$

Figure 1: Causal diagrams (a, b) with corresponding (c, d) MAGs and (e) PAG. A visible edge is marked with v next to it.

with stochastic policies and Carey et al. (2024) studied the completeness of its graphical characterization.

While SCM-MAB has been established as a general framework, it is challenging to align with real-world scenarios, as these studies assume that the decision-making agent has perfect access to the entire causal structure. In practice, only a Markov equivalence class of the true causal diagram over observed variables can be inferred from observational data without making a substantial assumption about causal mechanisms such as causal sufficiency (Verma and Pearl, 1990; Spirtes et al., 2001; Chickering, 2002; Tsamardinos et al., 2006) or a functional assumption (Perry et al., 2022; Peters et al., 2016; Ghassami et al., 2017; Heinze-Deml et al., 2018; Huang et al., 2020; Ghassami et al., 2018; Zeng et al., 2021). A prominent representation of the equivalence class is known as partial ancestral graphs (PAGs), and any causal diagrams can be uniquely represented by a PAG (Richardson and Spirtes, 2002; Zhang, 2006; 2008a;b; Ali, 2005).

Motivation and Contributions. With observational data, we can only learn a PAG, which represents *infinite* causal diagrams over supersets of the observed variables. Therefore, considering each causal diagram consistent with the PAG is computationally exhaustive. Identifying conditions for MIS and POMIS at the level of ancestral graphs *directly* would allow one to circumvent the issue. Recognizing the gap between the theoretical advantages of SCM-MAB and their practical applicability, we study SCM-MAB in terms of ancestral graphs (i.e., MAGs and PAGs).

Our key contributions are: (1) We generalize MIS and develop its graphical criteria in ancestral graphs, enabling an agent to identify and exclude variables that have no effect on the reward. (2) We devise POMIS for ancestral graphs along with its graphical characterization, leading to an action space that is worth exploring. (3) We present an efficient algorithm to determine whether a given set can be a POMIS in the Markov equivalence class represented by a PAG.

2. Preliminaries

We introduce notation and review relevant prior work. Following conventions, we use a capital letter, such as X, to represent a variable, with its corresponding lowercase letter, x, denoting a realization of the variable. Boldface is employed to represent a set of variables or values, denoted by $\mathbf{X}$ or $\mathbf{x}$. The domain of X is indicated by $\mathfrak{X}_X$. We use calligraphic letters for graphs and models such as $\mathcal{G}$ and $\mathcal{S}$.

Graphical notations. We consider a graph $\mathcal{G}$ having vertices $\mathbf{V}$ and edges $\mathbf{E}$ composed of directed ($\rightarrow$) and bidirected edges ($\leftrightarrow$). If there is an edge between two vertices X and Y in $\mathcal{G}$, we say that the two vertices are *adjacent* in $\mathcal{G}$ denoted by $Y \in \mathtt{Adj}(X)_{\mathcal{G}}$ or $X \in \mathtt{Adj}(Y)_{\mathcal{G}}$. An ordered sequence of distinct nodes in $\mathcal{G}$ is called a *path* between X and Y in $\mathcal{G}$ if (1) the start node is X and the end node is Y, and (2) there is an edge between any two subsequent variables in the sequence. If a path consists of directed edges with the same orientation, we say the path is *directed*. A variable Z is called a *collider* on the path if the path contains two edges having arrowheads toward Z. We define a path as a *collider path* if all non-endpoint vertices along the path are colliders. A path is *uncovered (unshielded)* if, for every consecutive triple on the path, its endpoints are not adjacent.

A path is *possibly directed* from X to Y if there is no arrowhead on the path pointing towards X. If there is a (possibly) directed path from X to Y, then Y is called a *(possible) descendant* of X, and X is a *(possible) ancestor* of Y. A variable Y is referred to as a *possible child* of X, and X is a *possible parent* of Y if they are adjacent and the edge is not directed into X. We denote the ancestors, descendants, parents, and children of a given variable as $\mathtt{An}$, $\mathtt{De}$, $\mathtt{Pa}$, and $\mathtt{Ch}$, respectively. Ancestors and descendants include the variable itself. For a set of variables, we define the ancestral set as $\mathtt{An}(\mathbf{X})_{\mathcal{G}} = \bigcup_{X \in \mathbf{X}} \mathtt{An}(X)_{\mathcal{G}}$, and similarly for other relationships. We add the prefix $\mathtt{Poss}$ when referring to possible relationships, such as $\mathtt{PossAn}$.

An *inducing path* relative to $\mathbf{L}$ is defined as a path where every vertex not in $\mathbf{L}$ is a collider on the path, and every collider is an ancestor of an endpoint of the path. A directed edge $X \rightarrow Y$ is *visible* if there exists no causal diagram in the corresponding equivalence class where there is an inducing path between X and Y that is into X. We refer to any edge that is not visible as *invisible*. The $\mathbf{X}$-lower-manipulation of $\mathcal{G}$ deletes all those visible edges and are out of variables in $\mathbf{X}$, and replaces all those edges that are out of variables in $\mathbf{X}$ but are invisible in $\mathcal{G}$ with bidirected edges denoted as $\mathcal{G}_{\underline{\mathbf{X}}}$. The $\mathbf{X}$-upper-manipulation of $\mathcal{G}$ deletes all those edges in $\mathcal{G}$ that are into variables in $\mathbf{X}$ denoted as $\mathcal{G}_{\overline{\mathbf{X}}}$.

We denote the set of variables in $\mathcal{G}$ by $\mathbf{V}(\mathcal{G})$. A subgraph $\mathcal{G}[\mathbf{V}']$, where $\mathbf{V}' \subseteq \mathbf{V}(\mathcal{G})$ is defined as a vertex-induced subgraph in which all edges among the vertices in $\mathbf{V}'$ are preserved. We define $\mathcal{G}\backslash\mathbf{X}$ as $\mathcal{G}[\mathbf{V}(\mathcal{G}) \setminus \mathbf{X}]$ for $\mathbf{X} \subseteq \mathbf{V}(\mathcal{G})$.

Structural Causal Model. We use structural causal model (SCM) (Pearl, 2000) as the semantical framework to represent the underlying environment a decision maker is deployed. An SCM $\mathcal{S}$ is a quadruple $\langle \mathbf{U}, \mathbf{V}, \mathbf{F}, P(\mathbf{U})\rangle$, where $\mathbf{U}$ is a set of exogenous variables determined by factors outside the model following a joint distribution $P(\mathbf{U})$, and $\mathbf{V}$ is a set of endogenous variables whose values are determined following a collection of functions $\mathbf{F} = \{f_i\}_{V_i \in \mathbf{V}}$ such that $V_i \leftarrow f_i(\mathbf{pa}_i, \mathbf{u}_i)$ where $\mathbf{PA}_i \subseteq \mathbf{V} \setminus \{V_i\}$ and $\mathbf{U}_i \subseteq \mathbf{U}$. The observational probability $P(\mathbf{v})$ is defined as $\sum_{\mathbf{u}} \prod_{V_i \in \mathbf{V}} P(v_i \mid \mathbf{pa}_i, \mathbf{u}_i) P(\mathbf{u})$. Every SCM $\mathcal{S}$ is associated with a *causal diagram* $\mathcal{G} = \langle \mathbf{V}, \mathbf{E}\rangle$ where a directed edge $V_i \rightarrow V_j \in \mathbf{E}$ if $V_i \in \mathbf{PA}_j$, and a bidirected edge between V_i and V_j if $\mathbf{U}_i$ and $\mathbf{U}_j$ are correlated. The probability of $\mathbf{V} = \mathbf{v}$ when $\mathbf{X}$ is intervened upon to take the value $\mathbf{x}$ is denoted by $P(\mathbf{v} \setminus \mathbf{x} \mid do(\mathbf{x}))$.

Ancestral graphical structures. Ancestral graphs are designed to capture graph structures without explicitly modeling latent variables. While directed edges between vertices in a causal diagram imply a direct causal effect between them, in ancestral graphs, directed edges instead represent ancestral relationships. Similar to the absence of directed cycles in causal diagrams, ancestral graphs do not permit *almost directed cycle*, which occurs when $X \leftrightarrow Y$ is present while X is an ancestor of Y.

A mixed graph is called a *maximal ancestral graph* (MAG) if (i) it does not contain any directed or almost directed cycles (i.e., ancestral); and (ii) there is no inducing path between any two non-adjacent vertices (i.e., maximal). In general, a MAG represents a set of causal diagrams with the same set of observed variables that entail the same conditional independence and ancestral relations among the observed variables. For each causal diagram, there exists a unique MAG over observed variables which represents its marginal independence relations, as well as its ancestral relations. However, a MAG is not fully testable with observational data since distinct MAGs can encode the same marginal independence relations. To illustrate, consider the causal diagrams $\mathcal{G}_1$ and $\mathcal{G}_2$ in Fig. 1. While they yield the same conditional independence relations, they correspond to distinct MAGs, $\mathcal{M}_1$ and $\mathcal{M}_2$, respectively.

A graph is a partial mixed graph (PMG) if it contains three types of marks: tails ($-$), arrowheads ($>$), and circles ($\circ$). A circle mark implies an uncertain mark that can be either an arrowhead or a tail. In addition, we use an asterisk ($*$) as a wildcard to denote any possible mark. In a PMG, if every edge mark on a path consists of circles, the path is

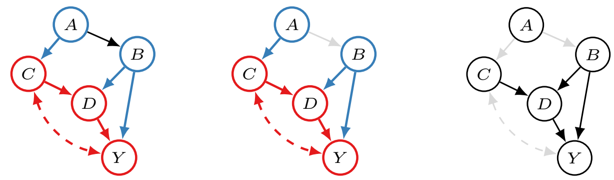

(a) $\mathcal{G} = \mathcal{G}_{\overline{\emptyset}}$ (b) $\mathcal{G}_{\overline{\{A,B\}}}$ (c) $\mathcal{G}_{\overline{\{A,B,C\}}}$

Figure 2: MUCT (red) and IB (blue) on subgraphs.

called a *circle path*, and each edge is called a *circle edge* ($\circ\!\!-\!\!\circ$). An edge is a *partially directed edge* ($\circ\!\!\rightarrow$) if it has both circle and arrowhead. A *circle component* is a subgraph of a PMG in which every edge is a circle edge. We use **?** mark to emphasize a wildcard that represents either a tail ($-$) or a circle ($\circ$), but not an arrowhead ($>$). Furthermore, $[\mathcal{Q}]$ denotes the set of MAGs represented by the PMG $\mathcal{Q}$, and similarly $[\mathcal{M}]$ denotes the set of causal diagrams conforming to the MAG $\mathcal{M}$.

A *partial ancestral graph* (PAG) denoted by $\mathcal{P}$, is a PMG such that it represents a Markov equivalence class of MAGs. Every MAG $\mathcal{M}$ represented by a PAG has the same skeleton as $\mathcal{P}$, and the non-circle marks in $\mathcal{P}$ are identical to those in $\mathcal{M}$. Every circle in $\mathcal{P}$ corresponds to a variant mark among the represented MAGs. The PAG $\mathcal{P}$ in Fig. 1e, for instance, is a PAG, as it encodes every MAG obtained by orienting circle marks incident to A and B as either $>$ or $-$, including both $\mathcal{M}_1$ and $\mathcal{M}_2$. In our work, we assume the absence of selection bias; therefore, there is no undirected edge in PAGs and MAGs we address.

Structural causal bandits. We follow the *structural causal bandit* (SCM-MAB) problem (Lee and Bareinboim, 2018), where an SCM models the target system with which an agent interacts, including a reward variable $Y \in \mathbf{V}$ and $\mathfrak{X}_Y \subseteq \mathbb{R}$. In the SCM-MAB setting, pulling each arm corresponds to intervening on a set of variables $\{\mathbf{x} \in \mathfrak{X}_{\mathbf{X}} \mid \mathbf{X} \subseteq \mathbf{V} \setminus \{Y\}\}$. The mean reward of an arm is denoted by $\mu_{\mathbf{x}} = \mathbb{E}[Y \mid do(\mathbf{x})]$ and the best expected reward by intervening on $\mathbf{X}$ is $\mu_{\mathbf{x}^*} = \max_{\mathbf{x} \in \mathfrak{X}_{\mathbf{X}}} \mu_{\mathbf{x}}$. We denote μ^* as the optimal expected reward. The goal of the agent is to minimize the cumulative regret after N rounds, which is given by $\mathrm{Reg}_N = \sum_{\mathbf{x} \in \mathfrak{X}_{\mathbf{X}}, \mathbf{X} \subseteq \mathbf{V} \setminus \{Y\}} \Delta_{\mathbf{x}} \mathbb{E}[T_{\mathbf{x}}(N)]$ where $T_{\mathbf{x}}(N)$ denotes the number of times the arm $\mathbf{x}$ was played after N rounds, $\Delta_{\mathbf{x}} = \mu^* - \mu_{\mathbf{x}}$ and $\mathfrak{X}_{\mathbf{X}} = \bigtimes_{X \in \mathbf{X}} \mathfrak{X}_X$.

MIS and POMIS. We review the notion of minimal intervention set (MIS) and possibly optimal minimal intervention set (POMIS) as well as their graphical characterizations for causal diagram by (Lee and Bareinboim, 2018; 2019). We denote by $\mathbf{x}[\mathbf{W}]$, values of $\mathbf{x}$ restricted to the subset of variables of $\mathbf{W} \cap \mathbf{X}$. We denote by $\mathbb{M}_{\mathcal{G},Y}$ and $\mathbb{P}_{\mathcal{G},Y}$, the sets of MISs and POMISs, respectively, given information $[\![\mathcal{G}, Y]\!]$.

Definition 1 (MIS (Lee and Bareinboim, 2018))**.** Given information $[\![\mathcal{G}, Y]\!]$, a set of variables $\mathbf{X} \subseteq \mathbf{V} \setminus \{Y\}$ is said to be a *minimal intervention set* (MIS) with respect to $[\![\mathcal{G}, Y]\!]$ if there is no $\mathbf{X}' \subsetneq \mathbf{X}$ such that $\mu_{\mathbf{x}[\mathbf{X}']} = \mu_{\mathbf{x}}$ for every SCM conforming to the causal diagram $\mathcal{G}$.

MIS leverages Rule 3 of do-calculus (Pearl, 1995) to eliminate variables that are irrelevant to the reward. Intuitively, a MIS can be understood as a set $\mathbf{X}$ in which there exists a directed path from any variable $X \in \mathbf{X}$ to Y, ensuring that each X can influence Y. In this context, the authors demonstrated that a set $\mathbf{X}$ is a MIS relative to $[\![\mathcal{G}, Y]\!]$ *if and only if* $\mathbf{X} \subseteq \mathtt{An}(Y)_{\mathcal{G}_{\overline{\mathbf{X}}}}$. For instance, consider $\mathcal{G}$ in Fig. 2. $\{A, B\}$ is a MIS relative to $[\![\mathcal{G}, Y]\!]$ since $\{A, B\} \subseteq \mathtt{An}(Y)_{\mathcal{G}_{\overline{\{A,B\}}}}$ holds. On the other hand, $\{A, B, C\}$ is a not a MIS since A is not an ancestor of Y in $\mathcal{G}_{\overline{\{A,B,C\}}}$.

Definition 2 (POMIS (Lee and Bareinboim, 2018))**.** Let $\mathbf{X} \subseteq \mathbf{V} \setminus \{Y\}$ be a MIS with respect to $[\![\mathcal{G}, Y]\!]$. If there exists an SCM conforming to $\mathcal{G}$ such that $\mu_{\mathbf{x}^*} > \forall_{\mathbf{W} \in \mathbb{M}_{\mathcal{G},Y} \setminus \{\mathbf{X}\}} \mu_{\mathbf{w}^*}$, then $\mathbf{X}$ is a *possibly-optimal minimal intervention set* (POMIS) with respect to $[\![\mathcal{G}, Y]\!]$.

If an agent is aware of POMIS, they should only explore and exploit actions consistent with those sets. When given a causal diagram $\mathcal{G}$, *minimal unobserved confounders' territory* (MUCT; Def. 12) and *interventional border* (IB; Def. 13) provide a graphical characterization of POMIS. In words, MUCT is the minimal set of variables that is closed under descendants and connected by a bidirected edge; and IB consists of the parents of MUCT, excluding MUCT itself. Intuitively, MUCT is the minimal closed mechanism that conveys all hidden information from unobserved confounders to the downstream reward, while IB consists of the nodes that directly affect this closed mechanism.

Theorem 1 (Theorem 6 in Lee and Bareinboim (2018))**.** *Let $\mathcal{G}$ be a causal diagram over the set of variables $\mathbf{V}$. Given information $[\![\mathcal{G}, Y]\!]$, a set $\mathbf{X} \subseteq \mathbf{V} \setminus \{Y\}$ is a POMIS if and only if IB for $\mathcal{G}_{\overline{\mathbf{X}}}$ with respect to Y equals $\mathbf{X}$.*

For instance, Fig. 2a shows MUCT (in red) and IB (in blue) for the subgraphs $\mathcal{G}_\emptyset$ and Fig. 2b shows $\mathcal{G}_{\overline{\{A,B\}}}$. The do-nothing action ($do(\emptyset)$) is a non-POMIS, as the IB of $\mathcal{G}_\emptyset$ is $\{A, B\}$ not $\emptyset$, while the set $\{A, B\}$ is identified as a POMIS, since the IB for $\mathcal{G}_{\overline{\{A,B\}}}$ is $\{A, B\}$.

In Appendix A, we provide detailed preliminaries for our work, along with brief descriptions of related works.

3. Generalizing Minimal Intervention Sets

Our goal is to find all sets that do not include variables irrelevant to the reward by ruling them out, referring to MAGs or PAGs. To achieve this, we first generalize minimal intervention set (MIS) to cover not only a causal diagram but also a MAG or a PAG over $\mathbf{V}$, denoted by $\mathcal{D}$.

Definition 3 (minimal intervention set)**.** Given information $[\![\mathcal{D}, Y]\!]$, a set of variables $\mathbf{X} \subseteq \mathbf{V} \setminus \{Y\}$ is called a *minimal intervention set* (MIS) relative to $[\![\mathcal{D}, Y]\!]$ if there is no $\mathbf{X}' \subsetneq \mathbf{X}$ such that $\mu_{\mathbf{x}[\mathbf{X}']} = \mu_{\mathbf{x}}$ for every SCM conforming to $\mathcal{D}$.

In the following parts, we provide complete graphical conditions for MIS in terms of MAGs and PAGs. Surprisingly, we then show in Sec. 3.2 that a MIS may include variables irrelevant to reward when dealing with PAGs. To address this issue, in Sec. 3.3, we propose the concept of *definitely* minimal intervention set (DMIS), which ensures that no further variables can be pruned from the set.

3.1. MIS for MAGs

We begin by examining whether MIS for MAGs is complete in encompassing all sets that have been fully trimmed.

Proposition 1. *Let $\mathcal{M}$ be a MAG over the set of variables $\mathbf{V}$. A set $\mathbf{X} \subseteq \mathbf{V} \setminus \{Y\}$ is a MIS relative to $[\![\mathcal{M}, Y]\!]$ if and only if there exists a causal diagram $\mathcal{G}$ conforming to $\mathcal{M}$ such that $\mathbf{X}$ is a MIS relative to $[\![\mathcal{G}, Y]\!]$.*

We provide all the proofs in Appendix E along with auxiliary results in Appendix D. The proposition guarantees the existence of a causal diagram $\mathcal{G}$ where $\mathbf{X}$ is a MIS relative to $[\![\mathcal{G}, Y]\!]$, provided that $\mathbf{X}$ is a MIS relative to $[\![\mathcal{M}, Y]\!]$ for the given MAG $\mathcal{M}$.

We now proceed to the graphical characterization of MIS for MAGs, in a manner similar to causal diagrams, utilizing the explicit ancestral relations among variables in MAGs and Rule 3 of do-calculus for MAGs (Zhang, 2008b).

Theorem 2. *Let $\mathcal{M}$ be a MAG over the set of variables $\mathbf{V}$. Given information $[\![\mathcal{M}, Y]\!]$, a set $\mathbf{X} \subseteq \mathbf{V} \setminus \{Y\}$ is a MIS relative to $[\![\mathcal{M}, Y]\!]$ if and only if $\mathbf{X} \subseteq \mathtt{An}(Y)_{\mathcal{M}_{\overline{\mathbf{X}}}}$ holds.*

For example, consider $\mathcal{G}'$ and $\mathcal{M}$ in Figs. 3a and 3b where $\mathcal{G}' \in [\mathcal{M}]$. A set $\{A, B, C\}$ is a MIS relative to $[\![\mathcal{M}, Y]\!]$ since $\{A, B, C\} \subseteq \mathtt{An}(Y)_{\mathcal{M}_{\overline{\{A,B,C\}}}}$ holds.

Remark 1. *Even though a set $\mathbf{X}$ is a MIS with respect to $[\![\mathcal{M}, Y]\!]$, there is no guarantee that $\mathbf{X}$ is a MIS with respect to $[\![\mathcal{G}, Y]\!]$ for every causal diagram $\mathcal{G}$ conforming to $\mathcal{M}$.*

The set $\{A, B, C\}$ is also a MIS relative to $[\![\mathcal{G}', Y]\!]$ in Fig. 3a since $\{A, B, C\} \subseteq \mathtt{An}(Y)_{\mathcal{G}'_{\overline{\{A,B,C\}}}}$ holds. However, while $\mathcal{G}$ in Fig. 2a is also represented by $\mathcal{M}$, it is not a MIS with respect to $[\![\mathcal{G}, Y]\!]$.[1]

3.2. MIS for PAGs and Its Possible Vacuousness

We proceed to the characterization of MIS for PAGs. Unfortunately, we cannot rely on Rule 3 for PAGs (Thm. 6 in Appendix) because the rule is applied when ancestral

[1]The inducing path $A \to C \leftrightarrow Y$ in $\mathcal{G}$ appears as $A \to Y$ in $\mathcal{M}$ since C is an ancestor of Y in $\mathcal{G}$.

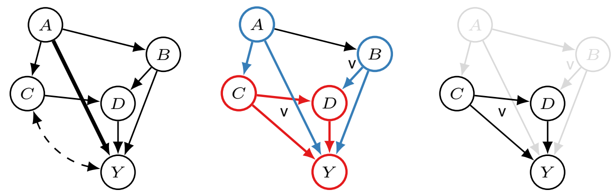

(a) $\mathcal{G}'$ (b) $\mathcal{M}$ (c) $\mathcal{M} \setminus \{A, B\}$

Figure 3: (a) The bold edge $A \rightarrow Y$ in $\mathcal{G}'$ represents an added edge from $\mathcal{G}$ in Fig. 2a. (b) MAG representing both $\mathcal{G}'$ and $\mathcal{G}$. (c) Induced graph of $\mathcal{M}$.

relations are apparent for *all* represented models, whereas a PAG might involve uncertainty reflected by circle marks. Hence, we define a specific type of path: A *proper* possibly-directed path from $X \in \mathbf{X}$ to Y with respect to $\mathbf{X}$, where only the first node X is in $\mathbf{X}$. This path is not disturbed by other intervening variables, thus aligning with the characterizations of MISs for causal diagrams and MAGs.

Proposition 2. *Let $\mathcal{P}$ be a PAG over the set of variables $\mathbf{V}$. A set $\mathbf{X} \subseteq \mathbf{V} \setminus \{Y\}$ is a MIS relative to $[\![\mathcal{P}, Y]\!]$ if and only if, for every variable $X \in \mathbf{X}$, there exists a proper possibly-directed path from X to Y with respect to $\mathbf{X}$ in $\mathcal{P}$.*

One might expect that if $\mathbf{X}$ is a MIS relative to $[\![\mathcal{P}, Y]\!]$, then it would also be a MIS relative to $[\![\mathcal{M}, Y]\!]$ for *some* MAG $\mathcal{M}$ conforming to the PAG $\mathcal{P}$. However, this is *not* always the case and no SCM associates $\mathbf{X}$ as MIS with respect to $[\![\mathcal{M}, Y]\!]$ as shown in the following example.

Consider a PAG $\mathcal{P}$ in Fig. 4. A set $\{A, B\}$ is a MIS with respect to $[\![\mathcal{P}, Y]\!]$ since each A and B has proper possibly-directed paths to Y (i.e., $A \circ\!\!-\!\!\circ Y$ and $B \circ\!\!-\!\!\circ Y$, respectively). However, we will demonstrate that at least one of A or B is irrelevant to reward in every conforming MAG. We consider two SCMs where the domains of variables are binary and $\forall_{U_V \in \{U_A, U_Y, U_B\}} P(U_V = 1) = \epsilon \approx 0$. For a proper subset $\mathbf{X}' = \{A\}$, we can construct an SCM $\mathcal{S}_1$ following that the mechanism for Y in $\mathcal{S}_1$ is $f_Y = b \oplus u_Y$, and the mechanism for B is $f_B = u_B$ where $\oplus$ denotes the exclusive-or function. Then, $\mu_a = \mu_\emptyset = 2\epsilon(1-\epsilon)$ while $\mu_{a,b^*} = \mu_{b^*} = 1 - \epsilon$ with $b^* = 1$. Thus, we find that $\mu_{a,b^*} > \mu_a$ holds in $\mathcal{S}_1$. This construction can be done for each proper subset of $\{A, B\}$, validating $\{A, B\}$ is a MIS relative to $[\![\mathcal{P}, Y]\!]$. However, the remarkable point here is that there is no representative SCM $\mathcal{S}^*$ that satisfies $\mu_{\mathbf{x}[\mathbf{X}']} \neq \mu_{\mathbf{x}}$ for *arbitrary* proper subset $\mathbf{X}' \subsetneq \mathbf{X}$, as doing so would require the mechanism f_y to depend on the values of both A and B. This setup would introduce an uncovered collider at Y in the underlying graph of $\mathcal{P}$, which is not consistent with the structure of $\mathcal{P}$. Therefore, we observed that $\{A, B\}$ is a MIS with respect to $[\![\mathcal{P}, Y]\!]$, but at least one of A or B is irrelevant to reward in all conforming MAGs.

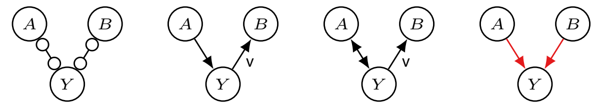

(a) $\mathcal{P}$ (b) $\mathcal{M}_1$ (c) $\mathcal{M}_2$ (d) $\mathcal{M}_3$

Figure 4: $\mathcal{M}_1$ and $\mathcal{M}_2$ are represented by $\mathcal{P}$. In contrast, $\mathcal{M}_3$ is not represented by $\mathcal{P}$.

3.3. Definitely MIS and Its Characterization

To address this vacuousness, we propose the concept of *definitely* MIS, which ensures the existence of a MAG where a set that is MIS in a PAG remains MIS. With the definition of MIS, we first choose $\mathbf{X}' \subsetneq \mathbf{X}$, and then check whether $\mu_{\mathbf{x}[\mathbf{X}']} \neq \mu_{\mathbf{x}}$ holds across all SCMs conforming to $\mathcal{D}$; here, we first choose an SCM $\mathcal{S}^*$ conforming to $\mathcal{D}$, and then check whether the inequality holds across all subsets $\mathbf{X}'$.

Definition 4 (definitely minimal intervention set)**.** Given information $[\![\mathcal{D}, Y]\!]$, a set $\mathbf{X} \subseteq \mathbf{V} \setminus \{Y\}$ is called a *definitely minimal intervention set* (DMIS) relative to $[\![\mathcal{D}, Y]\!]$, denoted by $\mathbb{D}_{\mathcal{D},Y}$ if there exists an SCM compatible with $\mathcal{D}$ such that, for every proper subset $\mathbf{X}' \subsetneq \mathbf{X}$, $\mu_{\mathbf{x}[\mathbf{X}']} \neq \mu_{\mathbf{x}}$ holds.

We now relate characterizations for MIS and DMIS.

Proposition 3. *If a set $\mathbf{X}$ is a DMIS with respect to $[\![\mathcal{D}, Y]\!]$, then $\mathbf{X}$ is a MIS with respect to $[\![\mathcal{D}, Y]\!]$.*

Proof. To see this, let $\mathcal{S}^*$ be an SCM associated with $\mathcal{D}$ such that $\mu_{\mathbf{x}[\mathbf{X}']} \neq \mu_{\mathbf{x}}$ holds for every $\mathbf{X}' \subsetneq \mathbf{X}$. Since such an $\mathcal{S}^*$ ensures that $\mu_{\mathbf{x}[\mathbf{X}']} \neq \mu_{\mathbf{x}}$ for all proper subsets, it guarantees that $\mathbf{X}$ satisfies the definition of a MIS relative to $[\![\mathcal{D}, Y]\!]$. □

Proposition 4. *Let $\mathcal{D}$ be either a causal diagram or a MAG. If a set $\mathbf{X}$ is a MIS with respect to $[\![\mathcal{D}, Y]\!]$, then $\mathbf{X}$ is a DMIS with respect to $[\![\mathcal{D}, Y]\!]$.*

proof sketch. To derive a contradiction, we can construct an SCM $\mathcal{S}^*$ where all mechanisms consist of the sum of the values of their parents, which ensures that $\mathbf{X}$ is a DMIS. □

This equivalence between MIS and DMIS for a causal diagram or a MAG (Props. 3 and 4) is derived from *deterministic* ancestral relations, $\mathbf{X} \subseteq \mathtt{An}(Y)_{\mathcal{D}_{\overline{\mathbf{X}}}}$. We now move on to discuss DMIS for PAGs, where ancestral relations are not deterministic. Recall that $\{A, B\}$ is a MIS but not a DMIS with respect to $[\![\mathcal{P}, Y]\!]$, as illustrated in Fig. 4a.

Proposition 5. *Let $\mathcal{P}$ be a PAG over the set of variables $\mathbf{V}$. A set $\mathbf{X} \subseteq \mathbf{V} \setminus \{Y\}$ is in $\mathbb{D}_{\mathcal{P},Y}$ if and only if there exists a MAG $\mathcal{M}$ conforming to $\mathcal{P}$ such that $\mathbf{X} \in \mathbb{M}_{\mathcal{M},Y}$.*

Hence, DMIS provides a *truly* feasible space for actions associated with intervention sets that no longer contain variables to rule out. According to Props. 3 and 4, we focus

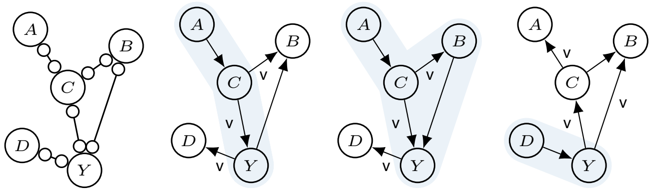

(a) $\mathcal{P}$ (b) $\mathcal{M}_1$ (c) $\mathcal{M}_2$ (d) $\mathcal{M}_3$

Figure 5: The nodes in the light blue region are the ancestors of Y. The three MAGs are represented by the PAG $\mathcal{P}$.

only on establishing the graphical criterion for DMIS only for PAGs. In Fig. 4a, we have observed that $A \circ\!\!-\!\!\circ Y$ and $B \circ\!\!-\!\!\circ Y$ cannot both be an ancestor of Y at the same time due to the uncovered path $A \circ\!\!-\!\!\circ Y \circ\!\!-\!\!\circ B$. To this end, we devise the notion of *relevance* among the edges in a PAG.

Definition 5 (relevant edges)**.** Let $\mathcal{P}$ be a PAG. For any edges $e_1(V_1 \ast\!\!-\!\!\ast V_2)$ and $e_2(V_{n-1} \ast\!\!-\!\!\ast V_n)$, we say that e_1 is *relevant* to e_2 in $\mathcal{P}$ if there exists an uncovered path $V_1 \ast\!\!-\!\!\circ V_2 \circ\!\!-\!\!\circ \cdots \circ\!\!-\!\!\circ V_{n-1} \circ\!\!-\!\!\ast V_n$ with $n \geq 3$ in $\mathcal{P}$.

Consider the one shown in Fig. 5. Similarly, $A \circ\!\!-\!\!\circ C$ and $D \circ\!\!-\!\!\circ Y$ are relevant in $\mathcal{P}$ because of the path $A \circ\!\!-\!\!\circ C \circ\!\!-\!\!\circ Y \circ\!\!-\!\!\circ D$. The key point here is that all triplets along the path are definite non-colliders so that the end nodes cannot be simultaneously ancestors of non-end nodes.

Theorem 3. *Let $\mathcal{P}$ be a PAG over the set of variables $\mathbf{V}$. A set $\mathbf{X} \subseteq \mathbf{V} \setminus \{Y\}$ is a DMIS relative to $[\![\mathcal{P}, Y]\!]$ if and only if, for any pair of vertices $X, Z \in \mathbf{X}$, there exist uncovered proper possibly-directed paths from X and Z to Y with respect to $\mathbf{X}$ such that their starting edges are not relevant.*

We now revisit Fig. 5 for example. Consider any MAGs represented by $\mathcal{P}$ where $A \circ\!\!-\!\!\circ C$ appears as a directed edge out of A (e.g., $\mathcal{M}_1$ and $\mathcal{M}_2$). Clearly, this results in $C \to Y \to D$, as the path is of definite status. On the other hand, if any MAG contains $D \to Y$, this leads to $Y \to C \to A$, as in $\mathcal{M}_3$. The important observation is that $A \to C$ ensures $D \notin \texttt{An}(Y)_{\mathcal{M}}$, and $D \to Y$ ensures $A \notin \texttt{An}(Y)_{\mathcal{M}}$ for all MAGs $\mathcal{M}$ represented by $\mathcal{P}$. This indicates that A and D cannot simultaneously be ancestors of Y in $\mathcal{M}$, thus $\{A, D\}$ is not a DMIS relative to $[\![\mathcal{P}, Y]\!]$.

4. Possibly Optimal Minimal Intervention Sets

We investigate into the graphical and algorithmic characterization of POMISs for MAGs and PAGs. POMISs are sets that include possibly optimal actions corresponding to interventions on specific values within the domains of the POMIS, implying that intervening on any set that is not a POMIS cannot yield a better outcome.

The main challenge in characterizing POMIS for ancestral graphs lies in the fact that induced paths by latent variables (or UCs) do not explicitly appear, which makes it impossible to directly identify the unobserved confounders' territory (Def. 12) as for causal diagrams. Instead, we leverage edge's *visibility* which indicates that the edge is not confounded in any underlying causal diagram (see Lem. 2 in Appendix for details). To generalize the UC-territory, we introduce a *possible c-component* (Jaber et al., 2018), which provides a necessary condition for nodes to belong to the same c-component in an underlying causal diagram.

Definition 6 (pc-component)**.** Two nodes are in the same *possible c-component* (pc-component) if there is a path between them such that (i) all non-endpoint nodes along the path are colliders, and (ii) none of the edges are visible.

We denote the pc-component of a PMG $\mathcal{Q}$ containing X as $\mathsf{PC}(X)_{\mathcal{Q}}$ and $\mathsf{PC}(\mathbf{X})_{\mathcal{Q}} \triangleq \bigcup_{X \in \mathbf{X}} \mathsf{PC}(X)_{\mathcal{Q}}$. For example, A and B are in the same pc-component in $\mathcal{P}$ of Fig. 1e because they are connected through an invisible colliding path $A \circ\!\!\rightarrow C \leftarrow\!\!\circ B$, i.e. $\mathsf{PC}(A)_{\mathcal{P}} = \{A, B, C\}$. Furthermore, due to $A \notin \mathsf{PC}(Y)_{\mathcal{P}} = \{B, Y\}$, A and Y cannot belong to the same c-component in any causal diagrams conforming to $\mathcal{P}$. We now generalize MUCT and IB for PMGs.

Definition 7 (unobserved-confounders' territory for PMGs)**.** Given information $[\![\mathcal{Q}, Y]\!]$ and intervention set $\mathbf{X} \subseteq \mathbf{V} \setminus \{Y\}$, let $\mathcal{H} = \mathcal{Q}[\texttt{PossAn}(Y)_{\mathcal{Q}} \setminus \mathbf{X}]$. A set of variables $\mathbf{T} \subseteq \texttt{PossAn}(Y)_{\mathcal{Q}} \setminus \mathbf{X}$ containing Y is called an *UC-territory* on $\mathcal{Q}$ with respect to Y if $\texttt{PossDe}(\mathbf{T})_{\mathcal{H}} = \mathbf{T}$ and $\mathsf{PC}(\mathbf{T})_{\mathcal{H}} = \mathbf{T}$. A UC-territory $\mathbf{T}$ is called a *minimal* UC-territory (MUCT) if no $\mathbf{T}' \subsetneq \mathbf{T}$ is a UC-territory, denoted as $\mathsf{MUCT}(\mathcal{Q}, Y, \mathbf{X})$.

Definition 8 (interventional border for PMGs)**.** Let $\mathbf{T}$ be a minimal UC-territory with respect to $[\![\mathcal{Q}, Y, \mathbf{X}]\!]$. Then $\mathbf{W} = \texttt{Pa}(\mathbf{T})_{\mathcal{Q}} \setminus \mathbf{T}$ is called an *interventional border* (IB) with respect to $[\![\mathcal{Q}, Y, \mathbf{X}]\!]$, denoted as $\mathsf{IB}(\mathcal{Q}, Y, \mathbf{X})$.

For concreteness, consider $\mathcal{M}$ and $\mathbf{X} = \{A, B\}$ in Fig. 3. Here, we omit `Poss`, as we discuss in the context of a MAG. Let $\mathcal{H}$ be the induced graph $\mathcal{M}[\texttt{An}(Y)_{\mathcal{M}} \setminus \mathbf{X}]$. In $\mathcal{H}$, C and Y are in the same pc component and D is a descendant of C. This implies that $\mathbf{T} = \{C, D, Y\}$ is the minimal closed set for $\texttt{De}_{\mathcal{H}}$ and $\mathsf{PC}_{\mathcal{H}}$, leading to $\mathsf{IB}(\mathcal{M}, Y, \mathbf{X}) = \{A, B\}$, derived from $\texttt{Pa}(\mathbf{T})_{\mathcal{M}} \setminus \mathbf{T}$.

4.1. POMIS for MAGs

We first establish a connection between a causal diagram and a MAG with respect to POMIS.

Proposition 6. *Let $\mathcal{M}$ be a MAG over the set of variables $\mathbf{V}$. A set $\mathbf{X} \subseteq \mathbf{V} \setminus \{Y\}$ is a POMIS relative to $[\![\mathcal{M}, Y]\!]$ if and only if there exists a causal diagram $\mathcal{G}$ conforming to $\mathcal{M}$ such that $\mathbf{X}$ is a POMIS relative to $[\![\mathcal{G}, Y]\!]$.*

Given the existence of an environment represented by the MAG in which the optimal action is an instance of a POMIS,

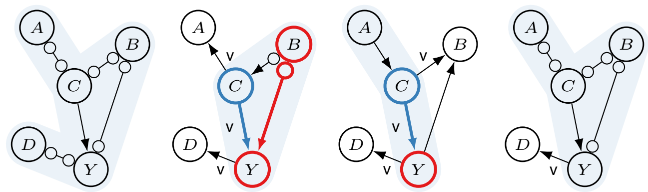

(a) $\mathcal{Q}_{\{C\}}$ (b) $\mathcal{Q}^1_{\{C\}}$ (c) $\mathcal{Q}^2_{\{C\}}$ (d) $\mathcal{Q}^*_{\{C\}}$

Figure 6: The light blue region indicates possible ancestors of Y. (a) PMG incorporating the information that $\{C\}$ is a DMIS. (b) $\mathbf{C}_Y^{\mathcal{Q}_{\{C\}}} = \{B\}$ and $\mathbf{C}_C^{\mathcal{Q}_{\{C\}}} = \{B\}$. (c) $\mathbf{C}_Y^{\mathcal{Q}_{\{C\}}} = \emptyset$ and $\mathbf{C}_C^{\mathcal{Q}_{\{C\}}} = \{A\}$. (d) PMG with orientation completeness.

POMISs for a MAG are indeed worth exploring. Equipped with the generalized MUCT and IB in the context of PMG, we proceed to characterize POMISs for MAGs.

Theorem 4. *Let $\mathcal{M}$ be a MAG over the set of variable $\mathbf{V}$. A set $\mathbf{X} \subseteq \mathbf{V} \setminus \{Y\}$ is a POMIS relative to $[\![\mathcal{M}, Y]\!]$ if and only if $\mathbf{X} = \mathsf{IB}(\mathcal{M}, Y, \mathbf{X})$.*

For example, see Fig. 3c where $\mathsf{IB}(\mathcal{M}, Y, \mathbf{X}) = \mathbf{X}$ holds for $\mathbf{X} = \{A, B\}$. Therefore, we get that $\{A, B\}$ is a POMIS with respect to $[\![\mathcal{M}, Y]\!]$, which is consistent with $\mathcal{G}$ in Fig. 2 where $\{A, B\}$ is a POMIS with respect to $[\![\mathcal{G}, Y]\!]$.

4.2. POMIS for PAGs

To begin with, we refine the *possibly-optimal minimal intervention set* over DMISs rather than MISs. This refinement ensures the existence of an underlying SCM represented by the PAG and implies the following proposition holds. Note that the refined POMIS aligns with established studies and serves as a natural extension, as discussed in Props. 3 and 4.

Definition 9 (possibly-optimal minimal intervention set)**.** Let $\mathbf{X} \subseteq \mathbf{V} \setminus \{Y\}$ be a DMIS relative to $[\![\mathcal{D}, Y]\!]$. If there exists an SCM conforming to $\mathcal{D}$ such that $\mu_{\mathbf{x}^*} > \forall_{\mathbf{W} \in \mathbb{D}_{\mathcal{D},Y} \setminus \{\mathbf{X}\}} \mu_{\mathbf{w}^*}$, then $\mathbf{X}$ is a *possibly-optimal minimal intervention set* (POMIS) relative to $[\![\mathcal{D}, Y]\!]$.

Proposition 7. *Let $\mathcal{P}$ be a PAG over the set of variables $\mathbf{V}$. A set $\mathbf{X} \subseteq \mathbf{V} \setminus \{Y\}$ is a POMIS relative to $[\![\mathcal{P}, Y]\!]$ if and only if there exists a MAG $\mathcal{M}$ conforming to $\mathcal{P}$ such that $\mathbf{X}$ is a POMIS relative to $[\![\mathcal{M}, Y]\!]$.*

The remainder of the main body focuses on characterizing POMIS for PAGs. A single PAG alone is insufficient to fully characterize POMIS due to significant structural uncertainty in circle marks. Since DMIS is already fully characterized, our strategy leverages this fact by incorporating the information that $\mathbf{X}$ is a DMIS into the PAG when $\mathbf{X}$ is a DMIS with respect to $[\![\mathcal{P}, Y]\!]$.

We first present necessary conditions for a PMG to represent MAGs $\mathcal{M}$ in which $\mathbf{X}$ is a MIS relative to $[\![\mathcal{M}, Y]\!]$.

Proposition 8. *Let $\mathcal{Q}_{\mathbf{X}}$ be a PMG representing MAGs where $\mathbf{X}$ is a MIS with respect to Y. Then, the following properties hold in $\mathcal{Q}_{\mathbf{X}}$, for $X \in \mathbf{X}$:*

1. *Every uncovered proper possibly-directed path from X to Y relative to $\mathbf{X}$ ends with an arrowhead (>).*
2. *If X is adjacent to Y, then the edge between X and Y is a directed edge.*

Recall the PAG $\mathcal{P}$ in Fig. 5a and $\mathbf{X} = \{C\}$, a DMIS relative to $[\![\mathcal{P}, Y]\!]$. Here, every MAG $\mathcal{M} \in [\mathcal{P}]$, satisfying that $\{C\}$ is a MIS relative to $[\![\mathcal{M}, Y]\!]$, conforms to $\mathcal{Q}_{\{C\}}$ in Fig. 6a.

Let $Q^*_{\mathbf{X}}$ be a PMG that satisfies the two conditions in Prop. 8, and the orientation completeness.[2] To characterize POMIS for PAGs, we partition $[Q^*_{\mathbf{X}}]$ based on the orientation of circle marks incident on $\mathbf{X} \cup \{Y\}$. We refer to a *local transformation* (Wang et al., 2023b) $\mathbf{C}_A^{\mathcal{Q}} \subseteq \{V \in \texttt{Adj}(A)_{\mathcal{Q}} \mid A \circ\!\!-\!\!* V\}$ as the vertices whose edges with a circle at A (i.e., $A \circ\!\!-\!\!* V$) will be oriented with arrowheads at A (i.e., $A \leftarrow\!\!* V$). The remaining vertices $V \in \{V \in \texttt{Adj}(A)_{\mathcal{Q}} \mid A \circ\!\!-\!\!* V\} \setminus \mathbf{C}_A^{\mathcal{Q}}$ will be oriented as $A \,-\!\!* V$. For clarity, consider Fig. 6 where all MAGs $\mathcal{M} \in [\mathcal{P}]$ that have $\{C\}$ as MIS relative to $[\![\mathcal{M}, Y]\!]$ are represented by $\mathcal{Q}^*_{\{C\}}$, which satisfies the orientation completeness. Each $\mathcal{Q}^1_{\{C\}}$ and $\mathcal{Q}^2_{\{C\}}$ illustrates a PMG where local transformations for C and Y are oriented, and both graphs are closed under the orientation rules.

Proposition 9. *Let $\mathcal{Q}^*_{\mathbf{X}}$ be a PMG which satisfies the conditions in Prop. 8 and the orientation completeness. For every MAG $\mathcal{M} \in [\mathcal{Q}^*_{\mathbf{X}}]$, if $\mathbf{X}$ is a MIS relative to $[\![\mathcal{M}, Y]\!]$, then there exists a PMG $\mathcal{Q}^i_{\mathbf{X}}$ representing $\mathcal{M}$ such that the following conditions are satisfied:*

1. *Every circle mark around nodes $\mathbf{X} \cup \{Y\}$ in $\mathcal{Q}_{\mathbf{X}}$ is oriented as either a tail (−) or an arrowhead (>) in $\mathcal{Q}^i_{\mathbf{X}}$ according to valid local transformations.* [3]
2. *Every $X \in \mathbf{X}$ is an ancestor of Y in $\mathcal{Q}^i_{\mathbf{X}}$.*
3. *$\mathcal{Q}^i_{\mathbf{X}}$ is closed under orientation rules.*[4]

In words, $\mathcal{Q}^i_{\mathbf{X}}$ is a *more* oriented PMG instance derived from $\mathcal{Q}_{\mathbf{X}}$ by applying the valid local transformations for circle marks around $\mathbf{X} \cup \{Y\}$, along with the orientation rules. Furthermore, each $\mathcal{Q}^i_{\mathbf{X}}$ confirms that $\mathbf{X}$ is a MIS relative to $[\![\mathcal{M}, Y]\!]$ for all MAGs $\mathcal{M} \in [\mathcal{Q}^i_{\mathbf{X}}]$.

[2]All circle marks are variant across the MAGs within $[\mathcal{Q}^*_{\{C\}}]$.

[3]For example, $\mathbf{C}_Y^{\mathcal{Q}_{\{A\}}} = \{B\}$ with $\mathbf{C}_B^{\mathcal{Q}_{\{A\}}} = \{C, Y\}$ is invalid, as it implies $B \leftrightarrow Y$ which introduces an almost directed cycle. In contrast, $\mathbf{C}_Y^{\mathcal{Q}_{\{A\}}} = \{B\}$ with $\mathbf{C}_B^{\mathcal{Q}_{\{A\}}} = \{C\}$ is valid, as depicted in Fig. 6b.

[4]The orientation rules refer to $\mathcal{R}_1$–$\mathcal{R}_3$, $\mathcal{R}'_4$, $\mathcal{R}_8$–$\mathcal{R}_{10}$ and $\mathcal{R}_{\text{SB}}$. $\mathcal{R}_5-\mathcal{R}_7$ are not considered since we assume no selection bias.

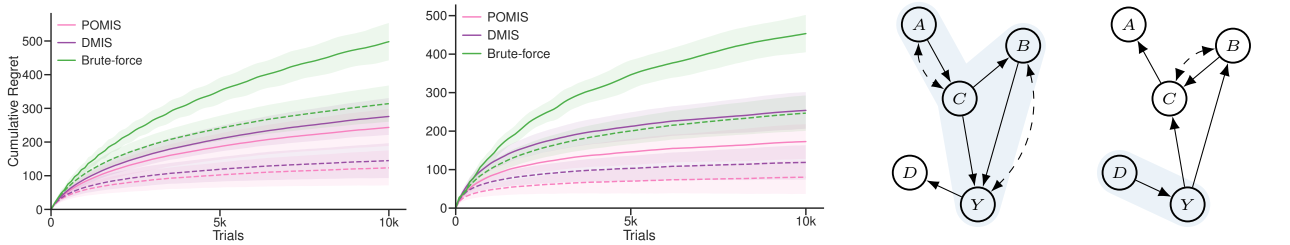

(a) (b) (c) $\mathcal{G}_1$ (d) $\mathcal{G}_2$

Figure 7: (a, b) Cumulative regret for the corresponding KL-UCB (solid) and TS (dashed) under distinct strategies. The number of trials is set to 10,000 and the shaded areas represent standard deviations based on 1,000 simulations. (c, d) The agent interacts with one of the models $\mathcal{S}_1$ or $\mathcal{S}_2$ which associate with $\mathcal{G}_1$ and $\mathcal{G}_2$, respectively. Both conform to $\mathcal{P}$ in Fig. 5a.

Algorithm 1 Identify whether a DMIS $\mathbf{X}$ is a POMIS for PAG.

function IsPOMIS$(\mathcal{P}, Y, \mathbf{X})$
Input: $\mathcal{P}$: PAG, Y: reward, $\mathbf{X}$: DMIS
1 Let $\mathcal{Q}_{\mathbf{X}}$ be a PMG oriented from $\mathcal{P}$ with $\mathbf{X}$ according to Prop. 8.
2 **return** subIsPOMIS$(\mathcal{Q}_{\mathbf{X}}, \mathbf{X} \cup \{Y\}, Y, \mathbf{X})$
3 **function** subIsPOMIS$(\mathcal{Q}, \mathbf{A}, Y, \mathbf{X})$
4   **if** $\mathbf{A}$ is empty **then**
5   **return** $\mathsf{IB}(\mathcal{Q}, Y, \mathbf{X}) = \mathbf{X}$
6     $A \leftarrow$ Pick a node from $\mathbf{A}$.
7   **for** each set $\mathbf{C}^{\mathcal{Q}}_A \subseteq \{V \in \mathtt{Adj}(A)_{\mathcal{Q}} \mid A \circ\!\!-\!\!* V\}$ **do**
8   **if** $\mathbf{C}^{\mathcal{Q}}_A$ satisfies Thm. 7 and $Y \in \mathtt{PossDe}(A)_{\mathcal{Q}\setminus\mathbf{C}^{\mathcal{Q}}_A}$ **then**
9   Let $\mathcal{Q}'$ be the PMG obtained by orienting the circle marks around A following $\mathbf{C}^{\mathcal{Q}}_A$ and completing the orientation rules from $\mathcal{Q}$.
10   **if** subIsPOMIS$(\mathcal{Q}', \mathbf{A} \setminus \{A\}, Y, \mathbf{X})$ **then**
11   **return** True
12   **return** False

We present an algorithm IsPOMIS, through which we can determine whether a given DMIS $\mathbf{X}$ is a POMIS relative to $[\![\mathcal{P}, Y]\!]$ based on our theoretical results Props. 8 and 9.

Theorem 5 (soundness and completeness)**.** *IsPOMIS returns* True *if and only if there exists a causal diagram $\mathcal{G}$ conforming to $\mathcal{P}$ such that $\mathbf{X}$ is a POMIS relative to $[\![\mathcal{G}, Y]\!]$.*

The algorithm begins by infusing the necessary condition for $\mathbf{X}$ to be a DMIS (Line 1). Then, the local transformations at $\mathbf{X} \cup \{Y\}$ are oriented recursively within subIsPOMIS. During each recursion, it evaluates the validity of a local transformation around a vertex and the ancestral relations between the vertex and the reward (Line 8), supported by Thm. 7 and lem. 22, respectively. The PMG updated according to the local transformation and the orientation rules proceeds to next recursive call (Lines 9–11). Finally, in the base case (Lines 4–5), we check whether the fully oriented PMG $\mathcal{Q}^i_{\mathbf{X}}$ satisfies $\mathsf{IB}(\mathcal{Q}^i_{\mathbf{X}}, Y, \mathbf{X}) = \mathbf{X}$. The key observation is that local transformations limited to $\mathbf{X} \cup \{Y\}$ are sufficient for this determination, thereby circumventing the need to enumerate all MAGs represented by the target PAG. To witness, consider $\mathcal{Q}^1_{\{C\}}$ with $\mathbf{X} = \{C\}$ in Fig. 6b where $\mathsf{IB}(\mathcal{Q}^1_{\{C\}}, Y, \mathbf{X}) = \mathbf{X}$ holds, and it follows that $\{C\}$ is a POMIS with respect to $[\![\mathcal{P}, Y]\!]$. Indeed, we can find a MAG $\mathcal{M}$ by orienting the circle marks around B in $\mathcal{Q}^1_{\{C\}}$ as tails, in which $\mathsf{IB}(\mathcal{M}, Y, \mathbf{X}) = \mathbf{X}$ also holds.

		$\mathcal{S}_1$	$\mathcal{S}_2$
TS	POMIS	$\mathbf{123.39} \pm 52.18$	$\mathbf{80.31} \pm 43.60$
	DMIS	144.84 ± 51.90	118.64 ± 44.73
	BF	313.97 ± 54.08	246.70 ± 46.09
KL-UCB	POMIS	$\mathbf{243.41} \pm 55.49$	$\mathbf{172.86} \pm 45.26$
	DMIS	275.87 ± 54.90	253.81 ± 47.66
	BF	497.90 ± 55.64	453.40 ± 48.69

Table 1: Mean and standard deviation of cumulative regret.

5. Experiments

We evaluate the cumulative regrets of SCM-MAB algorithm under different strategies to assess the effect of employing POMIS. In our setting, the deployed agent can only access the PAG $\mathcal{P}$ in Fig. 5a to obtain DMISs and POMISs. We consider three strategies for selecting arms: POMISs, DMISs, and Brute-force, combined with two prominent solvers: Thompson Sampling (TS) and Kullback-Leibler Upper Confidence Bound (KL-UCB). In the Brute-force strategy, all combinations of arms $\bigcup_{\mathbf{X}\subseteq\mathbf{V}\setminus\{Y\}} \mathfrak{X}_{\mathbf{X}}$ are evaluated. The environment in which an agent interacts is consistent with one of the causal diagrams $\mathcal{G}_1$ and $\mathcal{G}_2$.

Let us focus on an experiment comprising KL-UCB and $\mathcal{S}_2$ as a representative example. At the end of the trials, each of the three strategies–Brute-force (BF), DMIS and POMIS–yields cumulative regrets (mean $\pm$ standard deviation) of 453.40 ± 48.69, 253.81 ± 47.66, and $\mathbf{172.86} \pm 45.26$, respectively. The superiority of POMIS remains consistent across both $\mathcal{S}_1$ and $\mathcal{S}_2$, regardless of the solvers used. These results show that refining arms by taking the Markov equivalence class into account enhances the efficiency of agents. We provide detailed experimental settings along with additional experiments and discussions in Appendices B and C.

6. Conclusion

We proposed a novel structured causal bandit strategy in the context of ancestral graphs to prevent an agent from considering the infinite number of underlying causal diagrams. We believe these results have practical implications for designing intelligent agents, providing a foundation for optimizing the action space when the environment is abstracted as a Markov equivalence class.

References

Yasin Abbasi-Yadkori, Dávid Pál, and Csaba Szepesvári. Improved algorithms for linear stochastic bandits. *Advances in neural information processing systems*, 24, 2011.

Shipra Agrawal and Navin Goyal. Analysis of thompson sampling for the multi-armed bandit problem. In *Conference on learning theory*, pages 39–1. JMLR Workshop and Conference Proceedings, 2012.

R Ayesha Ali. Towards characterizing Markov equivalence classes for directed acyclic graph models with latent variables. In *Proc. Conf. on Uncertainty in Artificial Intelligence (UAI-05)*, pages 10–17, 2005.

Bryan Andrews, Peter Spirtes, and Gregory F Cooper. On the completeness of causal discovery in the presence of latent confounding with tiered background knowledge. In *International Conference on Artificial Intelligence and Statistics*, pages 4002–4011. PMLR, 2020.

Elias Bareinboim, Andrew Forney, and Judea Pearl. Bandits with unobserved confounders: A causal approach. In *Proceedings of the 28th Annual Conference on Neural Information Processing Systems*, pages 1342–1350, 2015.

Elias Bareinboim, Sanghack Lee, and Junzhe Zhang. An introduction to causal reinforcement learning, 2021.

Alexis Bellot. Towards bounding causal effects under markov equivalence. In *The 40th Conference on Uncertainty in Artificial Intelligence*, 2024.

Blair Bilodeau, Linbo Wang, and Dan Roy. Adaptively exploiting d-separators with causal bandits. *Advances in Neural Information Processing Systems*, 35:20381–20392, 2022.

Olivier Cappé, Aurélien Garivier, Odalric-Ambrym Maillard, Rémi Munos, and Gilles Stoltz. Kullback-leibler upper confidence bounds for optimal sequential allocation. *The Annals of Statistics*, pages 1516–1541, 2013.

Ryan Carey, Sanghack Lee, and Robin J Evans. Toward a complete criterion for value of information in insoluble decision problems. *arXiv preprint arXiv:2407.09883*, 2024.

Nicolo Cesa-Bianchi and Gábor Lugosi. Combinatorial bandits. *Journal of Computer and System Sciences*, 78 (5):1404–1422, 2012.

Olivier Chapelle and Lihong Li. An empirical evaluation of thompson sampling. *Advances in neural information processing systems*, 24, 2011.

David Maxwell Chickering. Optimal structure identification with greedy search. *Journal of machine learning research*, 3(Nov):507–554, 2002.

Gabriel Andrew Dirac. On rigid circuit graphs. In *Abhandlungen aus dem Mathematischen Seminar der Universität Hamburg*, volume 25, pages 71–76. Springer, 1961.

Tom Everitt, Ryan Carey, Eric D Langlois, Pedro A Ortega, and Shane Legg. Agent incentives: A causal perspective. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 11487–11495, 2021.

Shi Feng and Wei Chen. Combinatorial causal bandits. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pages 7550–7558, 2023.

Andrew Forney, Judea Pearl, and Elias Bareinboim. Counterfactual data-fusion for online reinforcement learners. In Doina Precup and Yee Whye Teh, editors, *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 1156–1164. PMLR, 06–11 Aug 2017. URL https://proceedings.mlr.press/v70/forney17a.html.

Aurélien Garivier and Olivier Cappé. The kl-ucb algorithm for bounded stochastic bandits and beyond. In *Proceedings of the 24th annual conference on learning theory*, pages 359–376. JMLR Workshop and Conference Proceedings, 2011.

AmirEmad Ghassami, Saber Salehkaleybar, Negar Kiyavash, and Kun Zhang. Learning causal structures using regression invariance. *Advances in Neural Information Processing Systems*, 30, 2017.

AmirEmad Ghassami, Negar Kiyavash, Biwei Huang, and Kun Zhang. Multi-domain causal structure learning in linear systems. *Advances in neural information processing systems*, 31, 2018.

Martin Charles Golumbic. *Algorithmic graph theory and perfect graphs*. Elsevier, 2004.

Michel Habib, Ross Mac Connell, Christophe Paul, and Laurent Viennot. Lex-bfs a partition refining technique, application to transitive orientation and consecutive 1's testing. *Theoretical Computer Science*, 234, 2000.

Christina Heinze-Deml, Jonas Peters, and Nicolai Meinshausen. Invariant causal prediction for nonlinear models. *Journal of Causal Inference*, 6(2), 2018.

Biwei Huang, Kun Zhang, Jiji Zhang, Joseph Ramsey, Ruben Sanchez-Romero, Clark Glymour, and Bernhard Schölkopf. Causal discovery from heterogeneous/nonstationary data. *Journal of Machine Learning Research*, 21(89):1–53, 2020.

Amin Jaber, Jiji Zhang, and Elias Bareinboim. Causal identification under Markov equivalence. In *Proceedings of the 34th Conference on Uncertainty in Artificial Intelligence*, 2018.

Amin Jaber, Adele Ribeiro, Jiji Zhang, and Elias Bareinboim. Causal identification under Markov equivalence: calculus, algorithm, and completeness. *Advances in Neural Information Processing Systems*, 35:3679–3690, 2022.

Emilie Kaufmann, Nathaniel Korda, and Rémi Munos. Thompson sampling: An asymptotically optimal finite-time analysis. In *International conference on algorithmic learning theory*, pages 199–213. Springer, 2012.

Tze Leung Lai and Herbert Robbins. Asymptotically efficient adaptive allocation rules. *Advances in applied mathematics*, 6(1):4–22, 1985.

Finnian Lattimore, Tor Lattimore, and Mark D Reid. Causal bandits: Learning good interventions via causal inference. *Advances in neural information processing systems*, 29, 2016.

Tor Lattimore and Csaba Szepesvári. *Bandit algorithms*. Cambridge University Press, 2020.

Sanghack Lee and Elias Bareinboim. Structural causal bandits: Where to intervene? *Advances in neural information processing systems*, 31, 2018.

Sanghack Lee and Elias Bareinboim. Structural causal bandits with non-manipulable variables. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 4164–4172, 2019.

Sanghack Lee and Elias Bareinboim. Characterizing optimal mixed policies: Where to intervene and what to observe. *Advances in neural information processing systems*, 33: 8565–8576, 2020.

Lihong Li, Wei Chu, John Langford, and Robert E Schapire. A contextual-bandit approach to personalized news article recommendation. In *Proceedings of the 19th international conference on World wide web*, pages 661–670, 2010.

Yangyi Lu, Amirhossein Meisami, Ambuj Tewari, and William Yan. Regret analysis of bandit problems with causal background knowledge. In *Conference on Uncertainty in Artificial Intelligence*, pages 141–150. PMLR, 2020.

Marloes H Maathuis and Diego Colombo. A generalized back-door criterion. *The Annals of Statistics*, 43(3):1060–1088, 2015.

Marloes H Maathuis, Markus Kalisch, and Peter Bühlmann. Estimating high-dimensional intervention effects from observational data. *The Annals of Statistics*, pages 3133–3164, 2009.

Judea Pearl. Causal diagrams for empirical research. *Biometrika*, 82(4):669–710, 1995.

Judea Pearl. *Causality: Models, Reasoning, and Inference*. Cambridge University Press, New York, 2000. 2nd edition, 2009.

Judea Pearl and James Robins. Probabilistic evaluation of sequential plans from causal models with hidden variables. In *Proceedings of the 11th Conference on Uncertainty in Artificial Intelligence*, pages 444–453. Morgan Kaufmann Publishers Inc., 1995.

Emilija Perkovic, Johannes Textor, Markus Kalisch, Marloes H Maathuis, et al. Complete graphical characterization and construction of adjustment sets in Markov equivalence classes of ancestral graphs. *Journal of Machine Learning Research*, 18(220):1–62, 2018.

Ronan Perry, Julius Von Kügelgen, and Bernhard Schölkopf. Causal discovery in heterogeneous environments under the sparse mechanism shift hypothesis. *Advances in Neural Information Processing Systems*, 35:10904–10917, 2022.

Jonas Peters, Peter Bühlmann, and Nicolai Meinshausen. Causal inference by using invariant prediction: identification and confidence intervals. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 78 (5):947–1012, 2016.

Thomas Richardson and Peter Spirtes. Ancestral graph Markov models. *The Annals of Statistics*, 30(4):962–1030, 2002.

Herbert Robbins. Some aspects of the sequential design of experiments. *Bull. Amer. Math. Soc.*, 58(5):527–535, 1952.

Donald J Rose, R Endre Tarjan, and George S Lueker. Algorithmic aspects of vertex elimination on graphs. *SIAM Journal on computing*, 5(2):266–283, 1976.

Peter Spirtes and Thomas S. Richardson. A polynomial time algorithm for determining dag equivalence in the presence of latent variables and selection bias. In David Madigan and Padhraic Smyth, editors, *Proceedings of the Sixth International Workshop on Artificial Intelligence and Statistics*, volume R1 of *Proceedings of Machine Learning Research*, pages 489–500. PMLR, 04–07 Jan 1997. URL https://proceedings.mlr.press/r1/spirtes97b.html. Reissued by PMLR on 30 March 2021.

Peter Spirtes, Clark N Glymour, and Richard Scheines. *Causation, Prediction, and Search*. MIT Press, 2nd edition, 2001. ISBN 9780262194402.

William R Thompson. On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. *Biometrika*, 25(3-4):285–294, 1933.

Jin Tian and Judea Pearl. A general identification condition for causal effects. In *Proceedings of the 18th National Conference on Artificial Intelligence*, pages 567–573, 2002.

Ioannis Tsamardinos, Laura E Brown, and Constantin F Aliferis. The max-min hill-climbing bayesian network structure learning algorithm. *Machine learning*, 65:31–78, 2006.

Burak Varici, Karthikeyan Shanmugam, Prasanna Sattigeri, and Ali Tajer. Causal bandits for linear structural equation models. *Journal of Machine Learning Research*, 24(297):1–59, 2023.

Aparajithan Venkateswaran and Emilija Perković. Towards complete causal explanation with expert knowledge. *arXiv preprint arXiv:2407.07338*, 2024.

Thomas Verma and Judea Pearl. Equivalence and synthesis of causal models. In *Proceedings of the Sixth Annual Conference on Uncertainty in Artificial Intelligence*, pages 255–270, 1990.

Tian-Zuo Wang, Tian Qin, and Zhi-Hua Zhou. Sound and complete causal identification with latent variables given local background knowledge. *Advances in Neural Information Processing Systems*, 35:10325–10338, 2022.

Tian-Zuo Wang, Tian Qin, and Zhi-Hua Zhou. Estimating possible causal effects with latent variables via adjustment. In Andreas Krause, Emma Brunskill, Kyunghyun Cho, Barbara Engelhardt, Sivan Sabato, and Jonathan Scarlett, editors, *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pages 36308–36335. PMLR, 23–29 Jul 2023a. URL https://proceedings.mlr.press/v202/wang23ag.html.

Tian-Zuo Wang, Tian Qin, and Zhi-Hua Zhou. Sound and complete causal identification with latent variables given local background knowledge. *Artificial Intelligence*, 322:103964, 2023b.

Tian-Zuo Wang, Wen-Bo Du, and Zhi-Hua Zhou. An efficient maximal ancestral graph listing algorithm. In *Forty-first International Conference on Machine Learning*, 2024a.

Tian-Zuo Wang, Lue Tao, and Zhi-Hua Zhou. New rules for causal identification with background knowledge, 2024b. URL https://arxiv.org/abs/2407.15259.

Marcel Wienöbst, Max Bannach, and Maciej Liśkiewicz. A new constructive criterion for Markov equivalence of MAGs. In *Proceedings of the Thirty-Eighth Conference on Uncertainty in Artificial Intelligence*, volume 180 of *Proceedings of Machine Learning Research*, pages 2107–2116. PMLR, 2022. URL https://proceedings.mlr.press/v180/wienobst22a.html.

Yan Zeng, Shohei Shimizu, Ruichu Cai, Feng Xie, Michio Yamamoto, and Zhifeng Hao. Causal discovery with multi-domain lingam for latent factors. In *Causal Analysis Workshop Series*, pages 1–4. PMLR, 2021.

Jiji Zhang. *Causal inference and reasoning in causally insufficient systems*. PhD thesis, Citeseer, 2006.

Jiji Zhang. On the completeness of orientation rules for causal discovery in the presence of latent confounders and selection bias. *Artificial Intelligence*, 172(16):1873–1896, 2008a. ISSN 0004-3702. doi: https://doi.org/10.1016/j.artint.2008.08.001. URL https://www.sciencedirect.com/science/article/pii/S0004370208001008.

Jiji Zhang. Causal reasoning with ancestral graphs. *Journal of Machine Learning Research*, 9:1437–1474, 2008b.

Jiji Zhang. A characterization of markov equivalence classes for directed acyclic graphs with latent variables. *arXiv preprint arXiv:1206.5282*, 2012.

Junzhe Zhang and Elias Bareinboim. Transfer learning in multi-armed bandit: a causal approach. In *Proceedings of the 16th Conference on Autonomous Agents and Multi-Agent Systems*, pages 1778–1780, 2017.

Hui Zhao, Zhongguo Zheng, and Baijun Liu. On the Markov equivalence of maximal ancestral graphs. *Science in China Series A: Mathematics*, 48:548–562, 2005.

A. Additional Preliminaries and Background Results

In this section, we provide additional preliminaries from previous works (A.1), assumptions of our work (A.2), and related works relevant to our study (A.3).

A.1. Additional Preliminaries

Definite status. Let p be any path in a PMG, and $\langle X, Z, Y\rangle$ be any consecutive triple along p. Z is a *definite collider* on p if both edges are directed into Z. If one of the edges is out of Z, or both edges have a circle mark at Z (i.e., $X \ast\!\!-\!\!\circ Z \circ\!\!-\!\!\ast Y$) and there is no edge between X and Z, then we say that the Z is a *definite non-collider* on p. A path is said to have a *definite status* if every non-endpoint node along it is either a definite collider or a definite non-collider.

Markov Equivalence Class. Multiple MAGs can entail the same m-separation[5] relationships. Such MAGs constitute a Markov equivalence class (MEC). The Markov equivalence class of MAGs can be uniquely represented by a PMG which we refer to as a PAG.

Definition 10 (Markov equivalence (Zhang, 2012))**.** Two MAGs $\mathcal{M}_1$, $\mathcal{M}_2$ with $\mathbf{V}(\mathcal{M}_1) = \mathbf{V}(\mathcal{M}_2)$ are Markov equivalent if for any three disjoint sets of vertices $\mathbf{X}, \mathbf{Y}, \mathbf{Z}$, $\mathbf{X}$ and $\mathbf{Y}$ are m-separated by $\mathbf{Z}$ in $\mathcal{M}_1$ *if and only if* $\mathbf{X}$ and $\mathbf{Y}$ are m-separated by $\mathbf{Z}$ in $\mathcal{M}_2$.

A path between X and Y, $p = \langle X, \cdots, W, Z, Y\rangle$, is a *discriminating path* for Z if (i) p includes at least three edges; (ii) Z is a non-endpoint vertex on p, and is adjacent to Y on p; and (iii) X is not adjacent to Y, and every vertex between X and Z is a collider on p and is a parent of Y.

For two MAGs to be in the same Markov equivalence class, discriminating paths must either be present in both graphs or none of the graphs, as well as the same skeleton and unobserved colliders.

Lemma 1 (graphical characterization of MEC (Spirtes and Richardson, 1997; Zhang, 2012))**.** *Two MAGs $\mathcal{M}_1$ and $\mathcal{M}_2$ with $\mathbf{V}(\mathcal{M}_1) = \mathbf{V}(\mathcal{M}_2)$ are Markov equivalent if and only if*

(i) they have the same adjacencies;

(ii) they have the same uncovered colliders; and

(iii) if some path is a discriminating path for a vertex V in both graphs $\mathcal{M}_1$ and $\mathcal{M}_2$, then V is a collider on the path in $\mathcal{M}_1$ if and only if it is a collider on the path in $\mathcal{M}_2$.

A collider path $\langle V_1, \cdots, V_k\rangle$ is called a *minimal collider path* if V_1 is not adjacent to V_k, and no subsequence of the path is also a collider path.

The two conditions (ii) and (iii) can be expressed as a condition for two MAGs to share the same minimal colliding paths (Zhao et al., 2005). Identifying Markov equivalence of a pair of MAGs is tractable with worst-case runtime $\mathcal{O}(|\mathbf{V}|^3)$ (Wienöbst et al., 2022).

Visible edges. A directed edge $X \to Y$ is *visible* if there exists no causal diagram in the corresponding equivalence class where there is an inducing path between X and Y that is into X. We refer to any edge that is not visible as *invisible*.

Lemma 2 (graphical characterization of visibility (Zhang, 2006; Maathuis and Colombo, 2015))**.** *A directed edge $X \to Y$ is visible if*

(i) there is a vertex Z not adjacent to Y, such that there is an edge between Z and X that is into X ($Z \ast\!\!\rightarrow X$); or

(ii) there is a collider path between Z and X that is into X ($Z \ast\!\!\rightarrow \cdot \leftrightarrow \cdots \leftrightarrow X$) and every vertex on the path except Z is a parent of Y.

It is important to note that (i) an invisible edge $X \to Y$ does not necessarily imply that X and Y are confounded in every underlying causal graph; and (ii) invisible edges should not be considered independently. To witness, consider a scenario

[5] M-separation (Richardson and Spirtes, 2002) refers to an extension of d-separation (Pearl and Robins, 1995) for ancestral graphs.

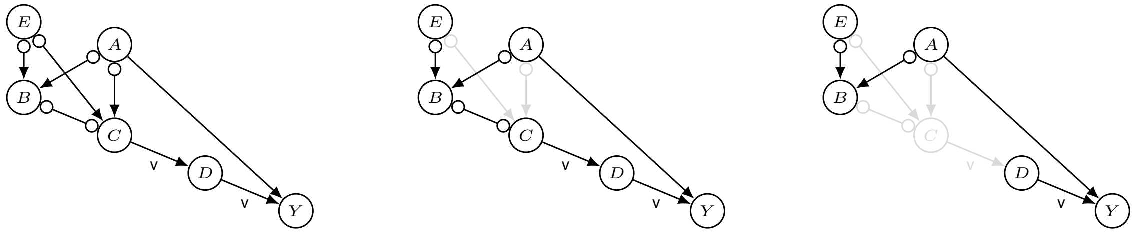

(a) $\mathcal{P}$ (b) $\mathcal{P}_{\overline{\{C\}}}$ (c) $\mathcal{P} \setminus \{C\}$

Figure 8: (a) PAG $\mathcal{P}$, (b) $\{C\}$-upper-manipulated graph, and (c) induced graph over $\mathbf{V}(\mathcal{P}) \setminus \{C\}$. In MAGs and PAGs, the visibility is preserved from $\mathcal{P}$ (see Lem. 15). For example, although there is no edge oriented into D in $\mathcal{P} \setminus \{C\}$, the directed edge $D \to Y$ remains visible.

where we have $X \leftarrow Y \to Z$ in a MAG $\mathcal{M}$, and X and Z are not adjacent. Since both edges, $X \leftarrow Y$ and $Y \to Z$, are invisible, causal diagrams can include at most one of the following structures added to $\mathcal{M}$: $X \leftarrow L_1 \leftarrow \cdots \to L_n \to Y$ or $Y \leftarrow L_1 \leftarrow \cdots \to L_n \to Z$ ($X \leftrightarrow Y$, or $Y \leftrightarrow Z$). Adding any one of these does not introduce a new collider between X and Z, thereby maintaining conformity with $\mathcal{M}$. However, if both are added simultaneously, a new collider is introduced at Y, resulting in a causal diagram that is not represented by $\mathcal{M}$.

Manipulations. Given a causal diagram $\mathcal{G}$ and a set of variables $\mathbf{X}$ therein, the $\mathbf{X}$-lower-manipulation of $\mathcal{G}$ deletes all edges in $\mathcal{G}$ that are out of the variables in $\mathbf{X}$. The resulting graph is denoted by $\mathcal{G}_{\underline{\mathbf{X}}}$. The $\mathbf{X}$-upper-manipulation of $\mathcal{G}$ deletes all edges in $\mathcal{G}$ that are into variables in $\mathbf{X}$. The resulting graph is denoted by $\mathcal{G}_{\overline{\mathbf{X}}}$.

Given a PMG $\mathcal{Q}$ and a set of variables $\mathbf{X}$ therein, the $\mathbf{X}$-lower-manipulation of $\mathcal{Q}$ deletes all those edges that are visible in $\mathcal{Q}$ and are out of variables in $\mathbf{X}$ and replaces all those edges that are out of variables in $\mathbf{X}$ but are invisible in $\mathcal{Q}$ with bidirected edges. The resulting graph is denoted as $\mathcal{Q}_{\overline{\mathbf{X}}}$. The $\mathbf{X}$-upper-manipulation of $\mathcal{Q}$ deletes all edges in $\mathcal{Q}$ that are into variables in $\mathbf{X}$, and otherwise keeps $\mathcal{Q}$ as it is.

The manipulated graphs play a crucial role in the derivation of do-calculus.

Do-calculus. Pearl (1995) devised *do-calculus* which acts as a bridge between observational and interventional distributions from a causal diagram without relying on any parametric assumptions. Zhang (2008b) proposed the do-calculus for MAGs and PAGs (also known as Zhang's calculus). Jaber et al. (2022) noted that there are cases where Pearl's do-calculus rules are applicable to every causal diagram within a given PAG, but Zhang's calculus cannot be applied to the same PAG. To address this, Jaber et al. (2022) proposed a refined version of do-calculus for PAGs and demonstrated that whenever the proposed rule is not applicable given a PAG, then the corresponding rule in Pearl's calculus is not applicable for some causal diagram in the Markov equivalence class represented by the PAG.

Here, we present do-calculus for PAGs, which encompasses that for MAGs.

Definition 11 (definite m-connecting path (Jaber et al., 2022))**.** In a PAG, a path p between X and Y is a *definite m-connecting path* relative to a set of nodes $\mathbf{Z}$ if p is definite status, every definite non-collider on p is not a member of $\mathbf{Z}$, and every collider on p is a ancestor of some member of $\mathbf{Z}$. X and Y are *m-separated* by $\mathbf{Z}$ if there is no definite m-connecting path between them relative to $\mathbf{Z}$.

Theorem 6 (do-calculus for PAGs (Jaber et al., 2022))**.** *Let $\mathcal{P}$ be the PAG over $\mathbf{V}$, and $\mathbf{X}$, $\mathbf{Y}$, $\mathbf{W}$, $\mathbf{Z}$ be disjoint subsets of $\mathbf{V}$. The following rules are valid, in the sense that if the antecedent of the rule holds, then the consequent holds in every MAG and consequently every causal diagrams represented by $\mathcal{P}$.*

Rule 1. $P(\mathbf{y} \mid do(\mathbf{w}), \mathbf{x}, \mathbf{z}) = P(\mathbf{y} \mid do(\mathbf{w}), \mathbf{z})$ *if $\mathbf{X}$ and $\mathbf{Y}$ are m-separated by $\mathbf{W} \cup \mathbf{Z}$ in $\mathcal{P}_{\overline{\mathbf{W}}}$*

Rule 2. $P(\mathbf{y} \mid do(\mathbf{w}), do(\mathbf{x}), \mathbf{z}) = P(\mathbf{y} \mid do(\mathbf{w}), \mathbf{x}, \mathbf{z})$ *if $\mathbf{X}$ and $\mathbf{Y}$ are m-separated by $\mathbf{W} \cup \mathbf{Z}$ in $\mathcal{P}_{\overline{\mathbf{W}},\underline{\mathbf{X}}}$*

Rule 3. $P(\mathbf{y} \mid do(\mathbf{w}), do(\mathbf{x}), \mathbf{z}) = P(\mathbf{y} \mid do(\mathbf{w}), \mathbf{z})$ *if $\mathbf{X}$ and $\mathbf{Y}$ are m-separated by $\mathbf{W} \cup \mathbf{Z}$ in $\mathcal{P}_{\overline{\mathbf{W}},\overline{\mathbf{X}(\mathbf{Z})}}$*

where $\mathbf{X}(\mathbf{Z}) \triangleq \mathbf{X} \setminus \texttt{PossAn}(\mathbf{Z})_{\mathcal{P}[\mathbf{V} \setminus \mathbf{W}]}$.

Induced graph. A subgraph $\mathcal{Q}[\mathbf{A}]$ is defined as a vertex-induced subgraph in which all the edges among the vertices in $\mathbf{A} \subseteq \mathbf{V}(\mathcal{Q})$ are preserved while maintaining the visibility from $\mathcal{Q}$ (see Fig. 8).

Chordal graph. We also introduce some useful graph theory and terminology, excerpted from Maathuis et al. (2009); Wang et al. (2023a). A graph is *chordal* if any cycle of length four or more has a chord, which refers to an edge joining two vertices that are not adjacent in the cycle. If a graph $\mathcal{G} = \langle \mathbf{V}, \mathbf{E} \rangle$ is chordal, then its subgraphs are also chordal. A vertex Z in $\mathbf{V}$ is called *simplicial* if $\mathcal{G}[\mathtt{Adj}(Z)_{\mathcal{G}}]$ induces a complete graph. As shown by Dirac (1961); Golumbic (2004), there are at least two non-adjacent simplicial vertices in any non-complete chordal graph with more than one vertex. A *perfect elimination order* of a graph $\mathcal{G}$ is an ordering $\sigma = (V_1, \cdots, V_{|\mathbf{V}|})$ of its vertices, so that each vertex V_i is a simplicial vertex in the subgraph $\mathcal{G} \setminus \{V_1, \cdots, V_{i-1}\}$. It is always possible to transform any circle component in a PAG into a *directed acyclic graph* (DAG) without introducing new unshielded colliders, as the circle component is chordal and every chordal graph has a perfect elimination order (Rose et al., 1976; Habib et al., 2000).

Orientation rules. *Fast Causal Inference* (FCI) is a causal discovery algorithm for identifying PAGs from conditional independence relationships derived from an observable distribution that follows underlying model. We present the complete orientation rules proposed by Zhang (2008a), omitting rules $\mathcal{R}_5$–$\mathcal{R}_7$ due to the absence of selection bias.

- $\mathcal{R}_0$ For each uncovered triple $\langle X, Z, Y \rangle$ in $\mathcal{P}$, orient it as a collider $X \ast\!\!\rightarrow Z \leftarrow\!\!\ast Y$ *if and only if* Z is not in $\mathsf{Sepset}(X, Y)$[6].
- $\mathcal{R}_1$ If $X \ast\!\!\rightarrow Z \circ\!\!-\!\!\ast Y$, and X and Y are not adjacent, then orient $Z \circ\!\!-\!\!\ast Y$ as $Z \rightarrow Y$.
- $\mathcal{R}_2$ If $X \rightarrow Z \ast\!\!\rightarrow Y$ or $X \ast\!\!\rightarrow Z \rightarrow Y$, and $X \ast\!\!-\!\!\circ Y$, then orient $X \ast\!\!-\!\!\circ Y$ as $X \ast\!\!\rightarrow Y$.
- $\mathcal{R}_3$ If $X \ast\!\!\rightarrow Z \leftarrow\!\!\ast Y$, $X \ast\!\!-\!\!\circ W \circ\!\!-\!\!\ast Y$, X and Y are not adjacent, and $W \ast\!\!-\!\!\circ Z$, then orient $W \ast\!\!-\!\!\circ Z$ as $W \ast\!\!\rightarrow Z$.
- $\mathcal{R}_4$ If $\langle X, \cdots, W, Z, Y \rangle$ is a discriminating path between X and Y for Z, and $Z \circ\!\!-\!\!\ast Y$; then if $Z \in \mathsf{Sepset}(X, Y)$, orient $Z \circ\!\!-\!\!\ast Y$ as $Z \rightarrow Y$; Otherwise orient the triple $\langle W, Z, Y \rangle$ as $W \leftrightarrow Z \leftrightarrow Y$.
- $\mathcal{R}_8$ If $X \rightarrow Z \rightarrow Y$, and $X \circ\!\!\rightarrow Y$, orient $X \circ\!\!\rightarrow Y$ as $X \rightarrow Y$.
- $\mathcal{R}_9$ If $X \circ\!\!\rightarrow Y$, and $p = \langle X, Z, W, \cdots, Y \rangle$ is an uncovered possibly directed path from X to Y such that Z and Y are not adjacent, then orient $X \circ\!\!\rightarrow Y$ as $X \rightarrow Y$.
- $\mathcal{R}_{10}$ Suppose $X \circ\!\!\rightarrow Y$, $Z \rightarrow Y \leftarrow W$, p_1 is an uncovered possibly directed path from X to Z, and p_2 is an uncovered possibly directed path from X to W. Let U be the vertex adjacent to X on p_1 (U could be Z), and V be the vertex adjacent to X on p_2 (V could be W). If U and V are distinct, and not adjacent, then orient $X \circ\!\!\rightarrow Y$ as $X \rightarrow Y$.

Incorporating background knowledge. Andrews et al. (2020) demonstrated that the ten rules $\mathcal{R}_1 - \mathcal{R}_{10}$ are complete for incorporating *tiered background knowledge*, which refers to background knowledge where the variables in a PAG can be partitioned into distinct groups with an explicit causal order defined among them.

Wang et al. (2022; 2023b) proposed that the rules $\mathcal{R}_1 - \mathcal{R}_3, \mathcal{R}_4', \mathcal{R}_8 - \mathcal{R}_{10}$ and $\mathcal{R}_{\mathrm{SB}}$ are complete for orienting a PAG when *local background knowledge* (i.e., *all* marks around a vertex) is available. The second additional rule $\mathcal{R}_{\mathrm{SB}}$ naturally follows from the absence of selection bias.[7]

- $\mathcal{R}_4'$ If $\langle X, \cdots, W, Z, Y \rangle$ is a discriminating path between X and Y for Z, and $Z \circ\!\!-\!\!\ast Y$; then orient $Z \circ\!\!-\!\!\ast Y$ as $Z \rightarrow Y$.
- $\mathcal{R}_{\mathrm{SB}}$ If $X -\!\!\circ Y$, then orient $X -\!\!\circ Y$ as $X \rightarrow Y$.

Furthermore, they built the necessary and sufficient conditions for validating local background knowledge (referred to here as *local transformation* in the context of our paper), which can be determined in $\mathcal{O}(|\mathbf{V}|^3)$.

[6] A set $\mathbf{Z} \in \mathsf{Sepset}(X, Y)$ if X and Y are independent given $\mathbf{Z}$.
[7] Wang et al. (2024a) proved that rules $\mathcal{R}_1$-$\mathcal{R}_{10}$ with one additional rule are sound and complete to incorporate local background knowledge to scenarios where selection bias is present.

Theorem 7 (Theorem 3 in Wang et al. (2023b)). *Denote $\mathcal{Q}$ the obtained PMG after some valid local transformations from a PAG $\mathcal{P}$ with orientation rules $\mathcal{R}_1$–$\mathcal{R}_3, \mathcal{R}'_4, \mathcal{R}_8$–$\mathcal{R}_{10}$ and $\mathcal{R}_{SB}$. Given a set $\mathbf{C}^{Q}_{X} \subseteq \{V \in \texttt{Adj}(X)_{\mathcal{Q}} \mid X \circ\!\!-\!\!* V\}$, there exists a MAG $\mathcal{M}$ consistent to $\mathcal{Q}$ with $X \leftarrow\!\!* V$ for all $V \in \mathbf{C}^{Q}_{X}$, and $X \rightarrow V$ for all $V \in \{V \in \texttt{Adj}(X)_{\mathcal{Q}} \mid X \circ\!\!-\!\!* V\} \setminus \mathbf{C}^{Q}_{X}$ if and only if $\mathbf{C}^{Q}_{X}$ satisfies the following conditions:*

1. $\texttt{PossDe}(X)_{\mathcal{Q}\setminus\mathbf{C}^{Q}_{X}} \cap \texttt{Pa}(\mathbf{C}^{Q}_{X})_{\mathcal{Q}} = \emptyset$;
2. $\mathcal{Q}[\mathbf{C}^{Q}_{X}]$ *is a complete graph;*
3. *Orient the subgraph* $\mathcal{Q}[\texttt{PossDe}(X)_{\mathcal{Q}\setminus\mathbf{C}^{Q}_{X}}]$ *as follows until no feasible updates:*
   For any vertices V_l and V_j such that $V_l \circ\!\!-\!\!\circ V_j$, orient it as $V_l \circ\!\!\rightarrow V_j$ if

   (i) $\mathcal{F}_{V_l} \setminus \mathcal{F}_{V_j} \neq \emptyset$, *or;*

   (ii) $\mathcal{F}_{V_l} = \mathcal{F}_{V_j}$ *as well as there is a vertex $V_m \in \texttt{PossDe}(X)_{\mathcal{Q}\setminus\mathbf{C}^{Q}_{X}}$ not adjacent to V_j such that $V_m \rightarrow V_l \circ\!\!-\!\!\circ V_j$*

   where $\mathcal{F}_{V_l} = \{V \in \mathbf{C}^{Q}_{X} \cup \{X\} \mid V *\!\!-\!\!\circ V_l \in \mathcal{Q}\}$. Then, no new uncovered colliders are introduced.

The PMG incorporating local transformations satisfies desirable properties as follows.

Theorem 8 (Theorem 1 in Wang et al. (2023b)). *Let $\mathcal{Q}$ be a PMG obtained from some valid local transformations from a PAG $\mathcal{P}$ and orientation rules $\mathcal{R}_1$–$\mathcal{R}_3, \mathcal{R}'_4, \mathcal{R}_8$–$\mathcal{R}_{10}$ and $\mathcal{R}_{SB}$. Then $\mathcal{Q}$ satisfies the following properties.*

(Closed). *$\mathcal{Q}$ is closed under the orientation rules.*

(Invariant). *The arrowheads (>) and tails (-) in $\mathcal{Q}$ are invariant in all the MAGs consistent with $\mathcal{Q}$.*

(Chordal). *The circle component in $\mathcal{Q}$ is chordal.*

(Balanced). *For any three nodes A, B, C in $\mathcal{Q}$, if $A *\!\!\rightarrow B \circ\!\!-\!\!* C$, then there is an edge between A and C with an arrowhead at C, namely, $A *\!\!\rightarrow C$. Furthermore, if the edge between A and B is $A \rightarrow B$, then the edge between A and C is either $A \rightarrow C$ or $A \circ\!\!\rightarrow C$ (i.e., it is not $A \leftrightarrow C$).*

(Complete). *For each circle at vertex A on any edge $A \circ\!\!-\!\!* B$ in $\mathcal{Q}$, there exist MAGs $\mathcal{M}_1$ and $\mathcal{M}_2$ consistent with $\mathcal{Q}$ such that $A \leftarrow\!\!* B$ in $\mathcal{M}_1$ and $A \rightarrow B$ in $\mathcal{M}_2$.*

Recently, Venkateswaran and Perković (2024); Wang et al. (2024b) devised additional rules for more general type of background knowledge. However, the completeness of the orientations in the resulting PMG after applying these rules remains an open problem.[8]

- $\widetilde{\mathcal{R}_4}$ If $\langle X = V_0, V_1, \cdots, V_k, V_{k+1} = Y\rangle$ with $k \geq 2$ is an *almost* discriminating path[9] for V_k in the graph and if $V_k \circ\!\!-\!\!* Y$ is in the graph, then orient $V_k \circ\!\!-\!\!* Y$ as $V_k \rightarrow Y$.
- $\mathcal{R}_{11}$ If $W *\!\!-\!\!* X *\!\!\rightarrow Y \rightarrow Z *\!\!-\!\!\circ W$ and $W \circ\!\!-\!\!\circ Y$, then orient $W \circ\!\!-\!\!* Z$ as $W \rightarrow Z$.
- $\mathcal{R}_{12}$ If there is an unshielded path of the form $V_1 \circ\!\!-\!\!\circ V_2 \circ\!\!-\!\!\circ \cdots \circ\!\!-\!\!\circ V_{n-1} \circ\!\!-\!\!* V_n$ with $i > 2$, as well as a path $V_n \rightarrow V_{n+1} \leftrightarrow V_1$, then orient $V_1 \circ\!\!-\!\!\circ V_2$ as $V_1 \leftarrow\!\!\circ V_2$.
- $\mathcal{R}_{13}$ Let W, X, Y, Z and $V_1, \cdots, V_k$ with $k > 1$ be distinct nodes. If $W \circ\!\!-\!\!* X$, $Y \leftrightarrow W \leftrightarrow Z$, and uncovered path $Y \leftarrow\!\!\circ V_1 \circ\!\!-\!\!\circ \cdots \circ\!\!-\!\!\circ V_k \circ\!\!\rightarrow Z$ are in the graph, and if there are uncovered possibly directed paths $\langle W, X, \cdots, V_i\rangle$ in the graph, for all $i \in \{1, \cdots, k\}$, then orient $W \circ\!\!-\!\!* X$ as $W \leftarrow\!\!* X$.

Wang et al. (2023a) leveraged the PMG incorporating local background knowledge to determine whether a given set of variables can be an adjustment set in some MAG consistent with the PMG, and Wang et al. (2024b) demonstrated that the additional rules can improve this process.

[8] The rules follow the version in Venkateswaran and Perković (2024).

[9] For more details on the concept of an almost discriminating path, we refer readers to Venkateswaran and Perković (2024).

Soundness and completeness of orientations. To eliminate ambiguity, we provide a formal description of soundness and completeness in the context of orientation within a PMG. Let $\mathcal{Q}$ be a PMG. We say that orientations in $\mathcal{Q}$ are *sound* if there is at least one MAG $\mathcal{M}$ conforming to $\mathcal{Q}$ such that invariant edge marks in $\mathcal{Q}$ are a subset of edge marks in $\mathcal{M}$. We say that the orientations in $\mathcal{Q}$ are *complete* if for every $A \circ\!\!-\!\!* B$ edge in $\mathcal{H}$, there are two MAGs $\mathcal{M}_1$ and $\mathcal{M}_2$ represented by $\mathcal{Q}$ containing the edges $A \rightarrow B$ and $A \leftarrow\!\!* B$, respectively, such that $\mathcal{M}_1$ and $\mathcal{M}_2$ conforming to $\mathcal{Q}$.

Structural causal bandit. Let $\mathcal{G}$ be a causal diagram and $\mathsf{CC}(X)_{\mathcal{G}}$ be the *c-component* (Tian and Pearl, 2002) of $\mathcal{G}$ that contains X where a c-component is a maximal set of vertices connected with bidirected edges. We denote $\mathsf{CC}(\mathbf{X})_{\mathcal{G}} = \bigcup_{X \in \mathbf{X}} \mathsf{CC}(X)_{\mathcal{G}}$. Let $\mathsf{MUCT}(\mathcal{G}, Y)$ and $\mathsf{IB}(\mathcal{G}, Y)$ be the MUCT and IB given $[\![\mathcal{G}, Y]\!]$, respectively.

Proposition 10 (Proposition 1 in Lee and Bareinboim (2018))**.** *Let $\mathcal{G}$ be a causal diagram over the set of variables $\mathbf{V}$. A set $\mathbf{X} \subseteq \mathbf{V} \setminus \{Y\}$ is a MIS relative to $[\![\mathcal{G}, Y]\!]$ if and only if $\mathbf{X} \subseteq \mathtt{An}(Y)_{\mathcal{G}_{\overline{\mathbf{X}}}}$.*

Definition 12 (unobserved-confounders' territory)**.** Given information $[\![\mathcal{G}, Y]\!]$, let $\mathcal{H} = \mathcal{G}[\mathtt{An}(Y)_{\mathcal{G}}]$. A set of variables $\mathbf{T} \subseteq \mathbf{V}(\mathcal{H})$ containing Y is called a *UC-territory* on $\mathcal{G}$ with respect to Y if $\mathtt{De}(\mathbf{T})_{\mathcal{H}} = \mathbf{T}$ and $\mathsf{CC}(\mathbf{T})_{\mathcal{H}} = \mathbf{T}$. A UC-territory $\mathbf{T}$ is said to be *minimal* if no $\mathbf{T}' \subsetneq \mathbf{T}$ is a UC-territory (MUCT).

Definition 13 (interventional border)**.** Let $\mathbf{T}$ be a minimal UC-territory on causal diagram $\mathcal{G}$ with respect to Y. Then $\mathbf{W} = \mathtt{Pa}(\mathbf{T})_{\mathcal{G}} \setminus \mathbf{T}$ is called an *interventional border* (IB) for $\mathcal{G}$ with respect to Y.

A.2. Assumptions

In this paper, we assume that there is *no selection bias* in the SCM-MAB system.

A.3. Background Results

We present useful results established in existing work.

A.3.1. Background Results in Zhang (2006; 2008a)

Lemma 3 (Lemma 0, as used in the proof of Lemma 5.1.7 in Zhang (2006))**.** *Let X and Y be distinct nodes in a MAG $\mathcal{M}$. If $p = \langle X, \cdots, Z, V, Y\rangle$ is a discriminating path from X to Y for V in a MAG $\mathcal{M}$, and the corresponding subpath between X and V in $\mathcal{P}$ is (also) a collider path, then the path corresponding to p in $\mathcal{Q}$ is also a discriminating path for V.*

Lemma 4 (Lemma A.1 in Zhang (2008a) & Lemma 5 in Jaber et al. (2018))**.** *Let $\mathcal{P}$ be a PAG over $\mathbf{V}$, and let $\mathcal{P}[\mathbf{A}]$ be the subgraph of $\mathcal{P}$ induced by $\mathbf{A} \subseteq \mathbf{V}$. For any three nodes A, B, C, if $A *\!\!\rightarrow B \circ\!\!-\!\!* C$, then there is an edge between A and C with an arrowhead at C, namely, $A *\!\!\rightarrow C$. Furthermore, if the edge between A and B is $A \rightarrow B$, then the edge between A and C is either $A \rightarrow C$ or $A \circ\!\!\rightarrow C$ (i.e., it is not $A \leftrightarrow C$).*

Lemma 5 (Lemma 3.3.2 in Zhang (2006))**.** *In a PAG $\mathcal{P}$, for any two nodes A and B, if there is a circle path, then following holds:*

1. *If there is an edge between A and B, the edge is not into A or B;*
2. *For any other node C, $C *\!\!\rightarrow A$ if and only if $C *\!\!\rightarrow B$. Furthermore, $C \leftrightarrow A$ if and only if $C \leftrightarrow B$.*

Lemma 6 (Theorem 2 in Zhang (2008a))**.** *Let $\mathcal{P}$ be a PAG. Let $\mathcal{M}$ be the graph resulting from the following procedure applied to a $\mathcal{P}$.*

Step 1. Replace all partially directed edges ($\circ\!\!\rightarrow$) in $\mathcal{P}$ with directed edges ($\rightarrow$).

Step 2. Orient the circle component of $\mathcal{P}$ into a DAG with no unshielded colliders.

Then, the result graph $\mathcal{M}$ conforms to $\mathcal{P}$.

Lemma 7 (Lemma B.1 in Zhang (2008a)). *Let A and B be two distinct nodes in a PAG $\mathcal{P}$. If p is a possibly directed path from A to B in a PAG $\mathcal{P}$, then some subsequence of p forms an uncovered possibly directed path from A to B in $\mathcal{P}$.*

Lemma 8 (Lemma B.2 in Zhang (2008a)). *Let A and B be two distinct nodes in a PAG $\mathcal{P}$. If $p = \langle V_0(= A), \cdots V_n(= B)\rangle, n \geq 2$, is an uncovered possibly directed path from A to B in $\mathcal{P}$, and $V_{i-1} \ast\!\!\rightarrow V_i$ for some $i \in \{1, \cdots, n\}$, then $V_{j-1} \rightarrow V_j$ for all $j \in \{i+1, \cdots, n\}$.*

Lemma 9 (Lemma B.4 in Zhang (2008a)). *In a PAG $\mathcal{P}$, if there is a possibly directed path from A to B, then the edge between A and B, if any, is not into A.*

Lemma 10 (Lemma B.5 in Zhang (2008a)). *In a PAG $\mathcal{P}$, let A and B be two distinct nodes in a PAG $\mathcal{P}$. If there is a possibly directed path from A to B that is into B, then every uncovered possibly directed path from A to B is into B.*

Lemma 11 (Lemma B.7 in Zhang (2008a)). *In a PAG $\mathcal{P}$, if there is a circle path between two adjacent vertices in $\mathcal{P}$, then the edge between the two vertices is a circle edge ($\circ\!\!-\!\!\circ$).*

A.3.2. Background Results in Maathuis and Colombo (2015); Perkovic et al. (2018)

Lemma 12 (Lemma 7.6 in Maathuis and Colombo (2015)). *Let $\mathcal{P}$ be a PAG with k edges into X, $k \geq 0$. Then there exists at least one MAG $\mathcal{M}$ in the Markov equivalence class represented by $\mathcal{P}$ that has k edges into X.*

Lemma 13 (Lemma 48 in Perkovic et al. (2018)). *Let X be a node in a PAG $\mathcal{P}$. Let $\mathcal{M}$ be a MAG conforming $\mathcal{P}$ that satisfies Lem. 6. Then any edge that is either $X \circ\!\!-\!\!\circ Y$, $X \circ\!\!\rightarrow Y$, or invisible $X \rightarrow Y$ in $\mathcal{P}$ is invisible $X \rightarrow Y$ in $\mathcal{M}$.*

A.3.3. Background Results in Jaber et al. (2018; 2022)

Lemma 14 (Proposition 1 in Jaber et al. (2018)). *Let $\mathcal{P}$ be a PAG over $\mathbf{V}$, and $\mathcal{G}$ be any causal diagram in the equivalence class represented by $\mathcal{P}$. Let $X \neq Y$ be two nodes in $\mathbf{A} \subseteq \mathbf{V}$. If X is an ancestor of Y in $\mathcal{G}[\mathbf{A}]$, then X is a possible ancestor of Y in $\mathcal{P}[\mathbf{A}]$.*

Lemma 15 (Lemma 4 in Jaber et al. (2018)). *Let $\mathcal{P}$ be a PAG over $\mathbf{V}$. For every directed edge $X \rightarrow Y$ in induced subgraph $\mathcal{P}[\mathbf{A}]$ with $\mathbf{A} \subseteq \mathbf{V}$, if it is visible in $\mathcal{P}$, then it is also visible in $\mathcal{P}[\mathbf{A}]$.*

Lemma 16 (Proposition 2 in Jaber et al. (2018)). *Let $\mathcal{P}$ be a PAG over $\mathbf{V}$, and $\mathcal{G}$ be any causal diagram in the equivalence class represented by $\mathcal{P}$. Let $X \neq Y$ be two nodes in $\mathbf{A} \subseteq \mathbf{V}$. If X and Y are in the same c-component in $\mathcal{G}[\mathbf{A}]$, then X and Y are in the same pc-component in $\mathcal{P}[\mathbf{A}]$, i.e., $B \in \mathsf{CC}(A)_{\mathcal{G}[\mathbf{A}]} \Rightarrow B \in \mathsf{PC}(A)_{\mathcal{P}[\mathbf{A}]}$ where $A, B \in \mathbf{A}$.*

Algorithm 2 Partial Topological Order PTO (Jaber et al., 2018)

Input: $\mathcal{P}$, $\mathbf{A} \subseteq \mathbf{V}(\mathcal{P})$
Output: Partial Topological Order over $\mathcal{P}[\mathbf{A}]$

```
while there exists a bucket B in P[A] with only arrowheads incident on it do
    Extract B from P[A]
    A ← A \ B
end
The partial order is B^1 ≺ ··· ≺ B^m in reverse order of the bucket extraction, i.e., B^1 is the last bucket extracted and B^m is the first.
```

Lemma 17 (Proposition 4 in Jaber et al. (2018)). *Let $\mathcal{P}$ be a PAG over $\mathbf{V}$, and let $\mathcal{P}[\mathbf{A}]$ be the subgraph of $\mathcal{P}$ induced by $\mathbf{A} \subseteq \mathbf{V}$. Then, Alg. 2 is sound over $\mathcal{P}[\mathbf{A}]$, in the sense that the partial order is valid with respect to $\mathcal{G}[\mathbf{A}]$, for every causal diagram $\mathcal{G}$ in the equivalence class represented by $\mathcal{P}$.*[10]

Lemma 18 (Lemma 6 in Jaber et al. (2018)). *In $\mathcal{M}[\mathbf{A}]$, where $\mathcal{M}$ is a MAG over $\mathbf{V}$ and $\mathbf{A} \subseteq \mathbf{V}$, the following property holds:*

[10] A *bucket* refers to the closure of nodes connected with circle paths.

For any three vertices A, B, C, if $A \ast\!\!\rightarrow B \rightarrow C$ and both edges are invisible, then we have $A \ast\!\!\rightarrow C$ and the edge is invisible.

Lemma 19 (Lemma 18 in Jaber et al. (2022))**.** *Let $\mathcal{P}$ be a PAG over $\mathbf{V}$, and let $\mathcal{P}[\mathbf{A}]$ be the subgraph of $\mathcal{P}$ induced by $\mathbf{A} \subseteq \mathbf{V}$. In $\mathcal{P}[\mathbf{A}]$, the following property holds:*

For any three vertices A, B, C, if $A \ast\!\!\rightarrow B \circ\!\!\rightarrow C$ and both edges are invisible, then we have $A \ast\!\!\rightarrow C$ and the edge is invisible.

A.3.4. Background Results in Wang et al. (2023b; 2024a)

Lemma 20 (Lemma 2 in Wang et al. (2023b))**.** *Let $\mathcal{Q}$ be a PMG obtained from some valid local transformations from a PAG $\mathcal{P}$ and the orientation rules. If p is a possibly directed path from A to B in $\mathcal{Q}$, then some subsequence of p is an uncovered possibly directed path from A to B in $\mathcal{Q}$.*

Lemma 21 (Lemma 3 in Wang et al. (2023b))**.** *Let $\mathcal{Q}$ be a PMG obtained from some valid local transformations from a PAG $\mathcal{P}$ and the orientation rules. In a PMG $\mathcal{Q}$, for any two nodes A and B, if there is a circle path, then following holds:*

1. *If there is an edge between A and B, the edge is not into A or B;*
2. *For any other node C, $C \ast\!\!\rightarrow A$ if and only if $C \ast\!\!\rightarrow B$. Furthermore, $C \leftrightarrow A$ if and only if $C \leftrightarrow B$.*

Lemma 22 (Lemma 4 in Wang et al. (2023b))**.** *Let $\mathcal{Q}$ be a PMG obtained from some valid local transformations from a PAG $\mathcal{P}$ and the orientation rules. Suppose a MAG $\mathcal{M}$ consistent to $\mathcal{Q}$ and the local transformation $\mathbf{C}_X^{\mathcal{Q}}$. Then $Y \in \texttt{PossDe}(X)_{\mathcal{Q}\setminus\mathbf{C}_X^{\mathcal{Q}}}$ if and only if $Y \in \texttt{De}(X)_{\mathcal{M}}$.*

Lemma 23 (Lemma 16.1 in Wang et al. (2023b))**.** *Let $\mathcal{Q}$ be a PMG obtained from some valid local transformations from a PAG $\mathcal{P}$ and the orientation rules. The MAG oriented according to Lem. 6 conforms to $\mathcal{Q}$.*

Lemma 24 (Lemma 2 in Wang et al. (2024a))**.** *Let $\mathcal{Q}$ be a PMG obtained from some valid local transformations from a PAG $\mathcal{P}$ and the orientation rules. If there is an uncovered circle path $p = \langle V_1, V_2, \cdots, V_n\rangle, n \geq 3$ in $\mathcal{Q}$, then any two non-consecutive vertices are not adjacent (minimal circle path).*

B. Experiments

We provide the detailed experimental settings in the main paper and present the results of additional experiments.

B.1. Experiment Details

We consider three strategies for selecting arms: POMISs, DMISs, and Brute-force, combined with two prominent MAB solvers: *Thompson Sampling* (TS) (Thompson, 1933; Chapelle and Li, 2011; Agrawal and Goyal, 2012; Kaufmann et al., 2012) and *Kullback-Leibler Upper Confidence Bound* (KL-UCB) (Garivier and Cappé, 2011; Cappé et al., 2013). In the Brute-force strategy, all combinations of arms $\bigcup_{\mathbf{X}\subseteq\mathbf{V}\setminus\{Y\}} \mathfrak{X}_{\mathbf{X}}$ are evaluated. The SCM representing the target system as follows:

$$\mathcal{S}_1 = \begin{cases} f_A &= U_A \wedge U_{AC} \\ f_B &= C \vee ((1-U_B) \wedge U_{BY}) \\ f_C &= A \vee ((1-U_C) \wedge U_{AC}) \\ f_D &= Y \wedge U_D \\ f_Y &= [(1-B) \vee \{(1-C) \wedge (1-U_{BY})\}] \wedge U_Y \end{cases} \qquad \mathcal{S}_2 = \begin{cases} f_A &= C \oplus U_A \\ f_B &= (C \vee Y) \wedge (U_B \oplus U_{BC}) \\ f_C &= \{(1-Y) \oplus (1-U_{BC})\} \wedge U_C \\ f_D &= U_D \\ f_Y &= D \oplus U_Y, \end{cases}$$

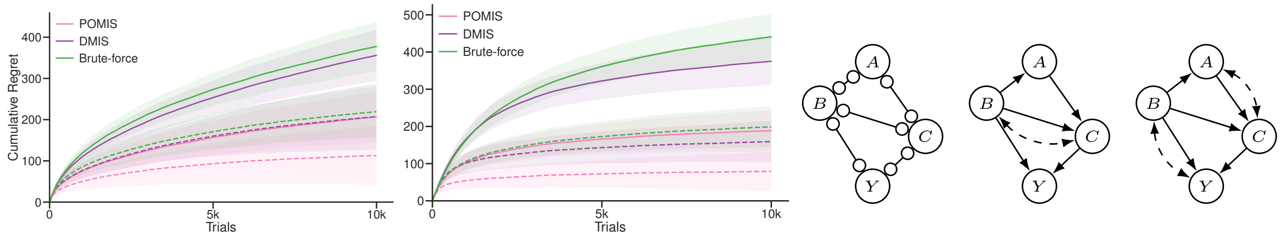

(a) (b) (c) $\mathcal{P}$ (d) $\mathcal{G}_1$ (e) $\mathcal{G}_2$

Figure 9: (a, b) Cumulative regret for the corresponding KL-UCB (solid) and TS (dashed) under distinct strategies: Brute-force, DMIS, and POMIS. (d, e) The agent interacts with one of the environments $\mathcal{S}_1$ and $\mathcal{S}_2$ which associate with $\mathcal{G}_1$ and $\mathcal{G}_2$, respectively. Both of causal diagrams conform to the PAG $\mathcal{P}$ in (c).

where the domains of all action variables A, B, C, D and the unobserved confounders are binary. For the first SCM $\mathcal{S}_1$, the exogenous variables are drawn from distributions $P(U_A = 1) = 0.44$, $P(U_B = 1) = 0.7$, $P(U_C = 1) = 0.4$, $P(U_D = 1) = 0.59$, $P(U_Y = 1) = 0.42$ and the unobserved confounder is drawn from $P(U_{BY} = 1) = 0.28$ and $P(U_{AC} = 1) = 0.77$. The second SCM $\mathcal{S}_2$ follows $P(U_A = 1) = 0.74$, $P(U_B = 1) = 0.74$, $P(U_C = 1) = 0.28$, $P(U_D = 1) = 0.32$, $P(U_Y = 1) = 0.23$ and the unobserved confounder is drawn from $P(U_{BC} = 1) = 0.46$. The DMISs and POMISs with respect to $[\![\mathcal{P}, Y]\!]$ are $\{\emptyset, \{A\}, \{B\}, \{C\}, \{D\}, \{A, B\}, \{B, C\}, \{B, D\}, \{C, D\}\}$ and $\{\emptyset, \{B\}, \{C\}, \{D\}, \{B, C\}, \{B, D\}, \{C, D\}\}$, respectively. The number of arms for each strategy—Brute-force, DMIS, and POMIS—is 81, 25, and 19, respectively.

B.2. Additional Experiments

We conducted additional experiments to confirm the effect of considering POMIS. The deployed agent does *not* have access to the underlying causal diagram and can only leverage the PAG $\mathcal{P}$ in Fig. 9, which abstracts the causal diagram of a true environment. The SCM representing the target system compatible with $\mathcal{P}$ as follows:

$$\mathcal{S}_1 = \begin{cases} f_A &= B \oplus U_A \\ f_B &= U_B \oplus U_{BC} \\ f_C &= (A \vee U_C) \wedge (B \oplus U_{BC}) \\ f_Y &= \{(1 - B) \oplus (1 - U_Y)\} \wedge C \end{cases} \qquad \mathcal{S}_2 = \begin{cases} f_A &= \{(1 - B) \oplus (1 - U_A)\} \wedge U_{AC} \\ f_B &= U_B \oplus U_{BY} \\ f_C &= (B \oplus U_{AC}) \wedge (A \vee U_C) \\ f_Y &= (B \oplus U_Y) \wedge (C \vee U_{BY}), \end{cases}$$

where the domains of all action variables A, B, C and the unobserved confounders are binary. For the first SCM $\mathcal{S}_1$, the exogenous variables are drawn from distributions $P(U_A = 1) = 0.23$, $P(U_B = 1) = 0.46$, $P(U_C = 1) = 0.22$, $P(U_Y = 1) = 0.47$, and the unobserved confounder is drawn from $P(U_{BC} = 1) = 0.59$. The second SCM $\mathcal{S}_2$ follows the same distribution for exogenous variables, while the unobserved confounders follow $P(U_{AC} = 1) = 0.59$ and $P(U_{BY} = 1) = 0.37$. The DMISs and POMISs with respect to $[\![\mathcal{P}, Y]\!]$ are $\{\emptyset, \{A\}, \{B\}, \{C\}, \{A, B\}, \{B, C\}, \{A, C\}\}$ and $\{\emptyset, \{A\}, \{B\}, \{C\}, \{B, C\}\}$, respectively.

		$\mathcal{S}_1$	$\mathcal{S}_2$
TS	POMIS	$\mathbf{112.99} \pm 70.90$	$\mathbf{79.17} \pm 51.61$
	DMIS	207.08 ± 79.75	159.13 ± 55.31
	BF	219.34 ± 61.98	198.96 ± 55.02
KL-UCB	POMIS	$\mathbf{206.96} \pm 61.89$	$\mathbf{188.60} \pm 55.10$
	DMIS	355.94 ± 62.07	375.13 ± 63.35
	BF	377.44 ± 59.35	440.86 ± 60.82

Table 2: Mean and standard deviation of cumulative regret.

The number of trials is set to 10,000, which is sufficient to observe the performance differences. Simulations are repeated 1,000 times to ensure consistent results.

The simulation results in Table 2 illustrate the clear performance gaps between the different strategies. We focus on a experiment comprising KL-UCB and $\mathcal{S}_2$ as a representative example. At the end of the trials, each of the three strategies–Brute-force (BF), DMIS, and POMIS–yields cumulative regrets (mean $\pm$ standard deviation) of 440.86 ± 60.82, 375.13 ± 63.35, and $\mathbf{188.60} \pm 55.10$, respectively. This clearly indicates the advantage of having a smaller number of arms (BF 27, DMIS 19, POMIS 11). Furthermore, the superiority of POMIS remains consistent across both $\mathcal{S}_1$ and $\mathcal{S}_2$, regardless of the solvers used. These results demonstrate that refining arms by considering the Markov equivalence class into account enhances the efficiency of agents when interacting with the underlying environment.

C. Discussions

We discuss circle mark transformations from the perspective of orientation completeness and topological order as well as the discussions on IsPOMIS. Finally, we discuss the expected future works of our study.

Local transformations. Let $\widetilde{\mathcal{Q}_{\mathbf{X}}}$ be a PMG that satisfies the two conditions in Prop. 8, and is closed under orientation rules $\mathcal{R}_1 - \mathcal{R}_3, \widetilde{\mathcal{R}_4}, \mathcal{R}_8 - \mathcal{R}_{13}$, and $\mathcal{R}_{\text{SB}}$. It is important to note that the completeness of $\widetilde{\mathcal{Q}_{\mathbf{X}}}$ remains an open problem. Therefore, $\widetilde{\mathcal{Q}_{\mathbf{X}}}$ is inadequate to completely characterize POMIS for PAGs.

Remark 2. *Every $\mathcal{Q}^i_{\mathbf{X}}$ is complete for orientations; for any $A \circ\!\!-\!\!* B$ in $\mathcal{Q}^i_{\mathbf{X}}$, there are two MAGs $\mathcal{M}_1$ and $\mathcal{M}_2$ represented by $\mathcal{Q}^i_{\mathbf{X}}$ containing $A \rightarrow B$ and $A \leftarrow\!\!* B$ respectively.*

Moreover, even though we have access to $\mathcal{Q}^*_{\mathbf{X}}$, a PMG that satisfies the two conditions in Prop. 8 and the orientation completeness, $\mathcal{Q}^*_{\mathbf{X}}$ is still insufficient to ensure $\mathbf{X} \subseteq \texttt{An}(Y)_{\mathcal{Q}^*_{\mathbf{X}}}$. To witness, consider a PAG $\mathcal{P}$ in Fig. 5 with $\mathbf{X} = \{A\}$.

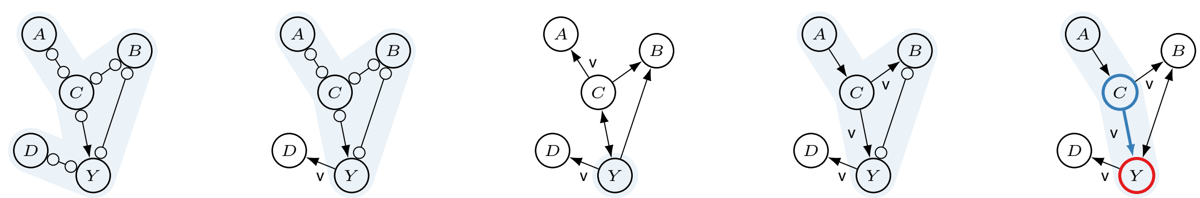

(a) $\mathcal{Q}_{\{A\}}$ (b) $\mathcal{Q}^*_{\{A\}}$ (c) $\mathcal{M}$ (d) $\bigwedge_{\forall i} \mathcal{Q}^i_{\{A\}}$ (e) $\mathcal{Q}^1_{\{A\}}$

Figure 10: The light blue region indicates possible ancestors of Y. (a) PMG incorporating the information that $\{A\}$ is a DMIS. (b) PMG with orientation completeness. (c) MAG represented by $\mathcal{Q}_{\{A\}}$ while $A \notin \texttt{An}(Y)_{\mathcal{M}}$. (d) PMG representing sound and complete orientations over MAGs satisfying that $\{A\}$ is a MIS. (e) PMG with $\mathbf{C}^{\mathcal{Q}_{\{A\}}}_{\{A\}} = \emptyset$ and $\mathbf{C}^{\mathcal{Q}_{\{A\}}}_{\{Y\}} = \{B\}$.

Then $C \circ\!\!-\!\!\circ Y$ in $\mathcal{P}$ corresponds to $C \circ\!\!\rightarrow Y$ in $\mathcal{Q}^*_{\mathbf{X}}$, according to the first condition in Prop. 8 supported by the uncovered proper possibly directed path $A \circ\!\!-\!\!\circ C \circ\!\!-\!\!\circ Y$. Moreover, $Y \rightarrow D$ is oriented by $\mathcal{R}_1$, and all remaining circle marks can vary across the underlying MAGs represented by $\mathcal{Q}^*_{\mathbf{X}}$. Here, we can find the MAG $\mathcal{M}$ where $\mathbf{X} \notin \texttt{An}(Y)_{\mathcal{M}}$ by orienting $C \circ\!\!\rightarrow Y$ as $C \leftrightarrow Y$, suggesting that additional information (orientation) is necessary.

Furthermore, neither $\mathcal{Q}^*_{\mathbf{X}}$ nor $\widetilde{\mathcal{Q}_{\mathbf{X}}}$ guarantees the *balanced property* (Lemmas 4 and 31). To witness, refer to $\mathcal{Q}^*_{\{C\}}$ in Fig. 6d, which is identical to $\widetilde{\mathcal{Q}_{\{C\}}}$. We can observe that there is $C \rightarrow Y \circ\!\!-\!\!\circ B$ while $C \circ\!\!-\!\!\circ B$, which violates the balanced property.

Complexity of IsPOMIS. A trivial approach is to enumerate all possible MAGs $\mathcal{M}$ conforming to a given PAG, and verify whether $\mathsf{IB}(\mathcal{M}, Y, \mathbf{X})$ holds for each $\mathcal{M}$. However, since the number of MAGs represented by the PAG is super-exponentially larger than the size of the PMG space (distinct PMGs $\mathcal{Q}^i_{\mathbf{X}}$ fully oriented by local transformations around $\mathbf{A} = \mathbf{X} \cup \{Y\}$), as discussed in Wang et al. (2022; 2023a;b; 2024a), such approach is prohibitive. Each local transformation for a vertex takes a $\mathcal{O}(2^p)$ complexity where p denotes the number of circle marks around the vertex, and each orientation rules take $\mathcal{O}(|\mathbf{V}|^3)$ complexity (Zhang, 2008a).

MUCT and IB in PAGs. One might surmise that $\mathbf{X} = \mathsf{IB}(\mathcal{P}, Y, \mathbf{X})$ is an appropriate characterization of POMIS for PAGs. However, this approach does *not* hold. For an illustration, consider a PAG $\mathcal{P}$ in Fig. 5a and a set $\mathbf{X} = \{A\}$, which is a DMIS with respect to $[\![\mathcal{P}, Y]\!]$. Moreover, we can simply derive $\mathsf{IB}(\mathcal{P}, Y, \mathbf{X}) = \{A\}$, and thus $\mathbf{X} = \mathsf{IB}(\mathcal{P}, Y, \mathbf{X})$ holds. For $\mathbf{X}$ to be a MIS for a MAG $\mathcal{M}$ represented by $\mathcal{P}$, the edge $A \circ\!\!-\!\!\circ C$ should correspond to $A \rightarrow C$ in $\mathcal{M}$, implying the visible edges $C \rightarrow B$ and $C \rightarrow Y$, as these are non-definite colliders (see $\mathcal{M}_1$ and $\mathcal{M}_2$ in Fig. 5, and Fig. 10d). Regardless of the edge orientation of $B \circ\!\!-\!\!\circ Y$, we find $\mathsf{IB}(\mathcal{M}, Y, \mathbf{X}) = \{C\}$, as in Fig. 10e. Thus, $\mathbf{X} = \{A\}$ is not a POMIS with respect to $[\![\mathcal{M}, Y]\!]$ for all $\mathcal{M} \in [\mathcal{P}]$. Therefore, the interventional border in PAGs $\mathsf{IB}(\mathcal{P}, Y, \mathbf{X})$ fails to characterize POMIS.

Characterizations in PMGs. MIS for PAGs can be generalized to PMGs obtained through local transformations, as they satisfy orientation completeness (Thm. 8). Furthermore, all characterizations for PAGs in the main paper can be applied to MAGs, since a MAG can be regarded as a PMG that is fully oriented through local transformations until no circle marks remain.

Future works. In future research, given the availability of an observational distribution, it becomes possible to identify specific causal effects and eliminate suboptimal arms (Jaber et al., 2022). Moreover, integrating this approach with *partial identification* (Bellot, 2024), enables the exclusion of arms where the upper bound is less than the lower bound of another arm, as proposed by Zhang and Bareinboim (2017). Finally, one can account for uncertainty in identification or bounds caused by a finite sample, which will lead to more robust analyses.

D. Auxiliary Results

In this section, we provide auxiliary results utilized throughout the paper.

Lemma 25. *Let $\mathcal{P}$ be a PAG over $\mathbf{V}$, and let $\mathcal{P}[\mathbf{A}]$ be the subgraph of $\mathcal{P}$ induced by $\mathbf{A} \subseteq \texttt{PossAn}(Y)_{\mathcal{P}} \subseteq \mathbf{V}$. If X and Z belong to different buckets over $\mathcal{P}[\mathbf{A}]$, then the starting edges of any uncovered proper possibly directed paths from X and Z to Y with respect to $\mathbf{X}$ are not relevant.*

Proof. Since X and Z are not in the same bucket, there is no circle path connecting the two nodes. Consequently, X and Z are not relevant. $\square$

Lemma 26. *Let $\mathcal{P}$ be a PAG over the set of variables $\mathbf{V}$. If a set $\mathbf{X} \subseteq \mathbf{V} \setminus \{Y\}$ is a DMIS relative to $[\![\mathcal{P}, Y]\!]$, then there exists a MAG $\mathcal{M}$ such that every $X \in \mathbf{X}$ has a proper directed path to Y with respect to $\mathbf{X}$ in $\mathcal{M}$.*

Proof. According to Prop. 5 and thm. 2, there exists a MAG $\mathcal{M}$ such that $\mathbf{X} \subseteq \texttt{An}(Y)_{\mathcal{M}_{\underline{\mathbf{X}}}}$. For the sake of contradiction, suppose that $\mathbf{X} \subseteq \texttt{An}(Y)_{\mathcal{M}_{\underline{\mathbf{X}}}}$ holds while there is no proper directed path from $X \in \mathbf{X}$ to Y with respect to $\mathbf{X}$ in $\mathcal{M}$. This implies that every directed path from X to Y must contain some node $Z \in \mathbf{X} \setminus \{X\}$. Consequently, such paths would be cut by the $\mathbf{X}$-lower manipulation, resulting in $X \notin \texttt{An}(Y)_{\mathcal{M}_{\underline{\mathbf{X}}}}$. This contradicts the assumption that $\mathbf{X} \subseteq \texttt{An}(Y)_{\mathcal{M}_{\underline{\mathbf{X}}}}$. $\square$

Lemma 27. *Let $\mathcal{Q}$ be a PMG obtained from some valid local transformations from a PAG $\mathcal{P}$ and the orientation rules. In $\mathcal{Q}$, the following property holds:*

If $A \to B$ is visible, then every $A \to C$ is also visible for every C connected as circle path with B.

Proof. For the sake of contradiction, assume that there exists a node C such that $A \to C$ is invisible while connected as circle path with B.

First, let $D \ast\!\!\to A$ be an arbitrary edge that makes $A \to B$ visible. Since $A \to C$ is invisible, D and C must be adjacent and the edge is into C by the orientation rule $\mathcal{R}_2$ (i.e., $D \ast\!\!\to C$). According to Lem. 31, this implies the existence of $D \ast\!\!\to B$, which contradicts the assumption that $A \to B$ is visible.

Next, consider the path $D \ast\!\!\to V_1 \leftrightarrow \cdots \leftrightarrow V_n \leftrightarrow A$ with $n \geq 1$ where V_i is a parent of B. By Lem. 31, we get that there exist edges $V_i \overset{?}{\to} C$ for all V_i. Furthermore, these edges must take the form $V_i \to C$, because if any edges $V_i \circ\!\!\to C$ existed, $\mathcal{R}'_4$ would be triggered, resulting in $V_i \to C$. Therefore, $A \to C$ is also visible, leading to a contradiction for the assumption that $A \to C$ is invisible. This concludes the proof. $\square$

Lemma 28. *Let $\mathcal{Q}$ be a PMG obtained from some valid local transformations from a PAG $\mathcal{P}$ and the orientation rules, and $\mathcal{G}$ be any causal diagram in the equivalence class represented by $\mathcal{Q}$. Let $X \neq Y$ be two nodes in $\mathbf{A} \subseteq \mathbf{V}(\mathcal{Q})$. If X is an ancestor of Y in $\mathcal{G}[\mathbf{A}]$, then X is a possible ancestor of Y in $\mathcal{Q}[\mathbf{A}]$.*

Proof. The lemma follows the proof of Lem. 14 (Prop. 1 in Jaber et al. (2018)). If X is an ancestor of Y in $\mathcal{G}[\mathbf{A}]$, then there exists a directed path $X \to \cdots \to Y$ in $\mathcal{G}[\mathbf{A}]$. This path is also present in $\mathcal{G}$, and consequently in the corresponding MAG $\mathcal{M}$. Hence, the path corresponds to a possibly directed path in $\mathcal{Q}$. Since all nodes along the path are in $\mathbf{A}$, they are also present in $\mathcal{Q}[\mathbf{A}]$, implying X is a possible ancestor of Y in $\mathcal{Q}[\mathbf{A}]$. □

Lemma 29. *Let $\mathcal{Q}$ be a PMG obtained from some valid local transformations from a PAG $\mathcal{P}$ and the orientation rules, and $\mathcal{G}$ be any causal diagram in the equivalence class represented by $\mathcal{Q}$. Let $X \neq Y$ be two nodes in $\mathbf{A} \subseteq \mathbf{V}(\mathcal{Q})$. For every $X \to Y$ in $\mathcal{Q}[\mathbf{A}]$, if it is visible in $\mathcal{Q}$, then it remains visible in $\mathcal{Q}[\mathbf{A}]$.*

Proof. The proof follows the argument of Lem. 15 (Lem 4. in Jaber et al. (2018)). Let $\mathcal{G}$ defined over $\mathbf{V}(\mathcal{Q}) \cup \mathbf{L}$. Let $X \to Y$ be a visible edge in $\mathcal{Q}$ where X and Y are in $\mathbf{A}$. Then, there are no inducing path between X and Y relative to $\mathbf{L}$ that is into X in $\mathcal{G}$. It follows that no such inducing path (relative to the latent nodes in $\mathcal{G}[\mathbf{A}]$) exists in the subgraph $\mathcal{G}[\mathbf{A}]$. □

Lemma 30. *Let $\mathcal{Q}$ be a PMG obtained from some valid local transformations from a PAG $\mathcal{P}$ and the orientation rules, and $\mathcal{G}$ be any causal diagram in the equivalence class represented by $\mathcal{Q}$. Let $X \neq Y$ be two nodes in $\mathbf{A} \subseteq \mathbf{V}(\mathcal{Q})$. If X and Y are in the same c-component in $\mathcal{G}[\mathbf{A}]$, then X and Y are in the same pc-component in $\mathcal{Q}[\mathbf{A}]$, i.e., $B \in \mathsf{CC}(A)_{\mathcal{G}[\mathbf{A}]} \Rightarrow B \in \mathsf{PC}(A)_{\mathcal{Q}[\mathbf{A}]}$ where $A, B \in \mathbf{A}$.*

Proof. The proof follows the argument of Lem. 16 (Prop. 2 in Jaber et al. (2018)). If X and Y are in the same c-component in $\mathcal{G}[\mathbf{A}]$, then there is a bidirected path p in $\mathcal{G}[\mathbf{A}]$.

> **Lemma I** (Lemma 6 in Jaber et al. (2018))**.** Let $\mathcal{M}$ be a MAG over $\mathbf{V}$ and $\mathcal{G}$ be a causal diagram represented by $\mathcal{M}$. For any X and Y in $\mathbf{V}$, if there is a bidirected path p between X and Y in $\mathcal{G}$, then there is a path p' between X and Y in $\mathcal{M}$ over a subsequence of p such that (1) all the non-endpoint nodes are colliders, and (2) all directed edges on p' are invisible.

> **Lemma II** (Lemma 7 in Jaber et al. (2018))**.** Let $\mathcal{M}$ be a MAG over $\mathbf{V}$ and $\mathcal{P}$ be a PAG representing $\mathcal{M}$. For any X and Y in $\mathbf{V}$, if there is a path p between X and Y in $\mathcal{M}$ such that (1) all non-endpoint nodes are colliders and (2) all directed edges, if any, are not visible, then there is a path p^* between X and Y in $\mathcal{P}$ over a subsequence of p such that (1) all non-endpoint nodes along the path are definite colliders, and (2) none of the edges are visible.

According to Lemma I, we choose a path p', which is the shortest subsequence of p between X and Y in $\mathcal{M}$, corresponding to p^* in $\mathcal{P}$, such that (1) all non-endpoint nodes along the path are colliders, and (2) none of the directed edges are visible. By Lemma II, the path p^* is a definite colliding path between X and Y, and none of the directed edges along the path are visible in $\mathcal{P}$. For contradiction, assume that $p^\dagger$ in $\mathcal{Q}$, which is corresponding to p^* in $\mathcal{P}$, includes a visible edge out of X. Then, the visible edge must appear in all MAGs represented by $\mathcal{Q}$. However, the edge along p' is invisible in $\mathcal{M}$, leading to a contradiction. Therefore, $p^\dagger$ is also of definite status, containing no visible edges, which implies that X and Y are in the same pc-component in $\mathcal{Q}$. Since all nodes along $p^\dagger$ are in $\mathbf{A}$, $p^\dagger$ is also present in $\mathcal{Q}[\mathbf{A}]$, ensuring that X and Y are in the same pc-component in $\mathcal{Q}[\mathbf{A}]$. □

Lemma 31. *Let $\mathcal{Q}$ be a PMG obtained from some valid local transformations from a PAG $\mathcal{P}$ and the orientation rules, and $\mathcal{Q}[\mathbf{A}]$ be the induced graph over $\mathbf{A} \subseteq \mathbf{V}(\mathcal{Q})$. For any three nodes A, B, C in $\mathcal{Q}$, if $A \ast\!\!\to B \circ\!\!-\!\!\ast C$, then there is an edge between A and C with an arrowhead at C, namely, $A \ast\!\!\to C$. Furthermore, if the edge between A and B is $A \to B$, then the edge between A and C is either $A \to C$ or $A \circ\!\!\to C$ (i.e., it is not $A \leftrightarrow C$).*

Proof. The balanced property holds in the PMG with local transformations as shown in Thm. 8 (Theorem 1 in Wang et al. (2023b)). By the definition of an induced graph, this property is preserved in $\mathcal{Q}[\mathbf{A}]$. □

Lemma 32. *Let $\mathcal{Q}$ be a PMG obtained from some valid local transformations from a PAG $\mathcal{P}$ and the orientation rules. In a PMG $\mathcal{Q}$, for any two nodes A and B, if there is a circle path, then following holds:*

1. *If there is an edge between A and B, the edge is not into A or B;*

2. *For any other node C, $C \ast\!\!\rightarrow A$ if and only if $C \ast\!\!\rightarrow B$. Furthermore, $C \leftrightarrow A$ if and only if $C \leftrightarrow B$.*

Proof. The proof follows the argument of Lem. 5 (Lem 3.3.2 in Zhang (2006)). The properties depend on the balanced property in Lem. 4, which holds in $\mathcal{Q}$ as demonstrated in Thm. 8 and lem. 31. □

Lemma 33. *Let $\mathcal{Q}$ be a PMG obtained from some valid local transformations from a PAG $\mathcal{P}$ and the orientation rules. PTO (Alg. 2) is also sound over $\mathcal{Q}[\mathbf{A}]$, in the sense that the partial order is valid with respect to $\mathcal{G}[\mathbf{A}]$, for every causal diagram $\mathcal{G}$ in the equivalence class represented by $\mathcal{Q}$.*

Proof. The proof follows the argument of Lem. 17 (Prop. 4 in Jaber et al. (2018)). By Lem. 28, the possible-ancestral relations in $\mathcal{Q}[\mathbf{A}]$ subsume those in $\mathcal{G}[\mathbf{A}]$. Hence, a partial topological order that is valid with respect to $\mathcal{Q}[\mathbf{A}]$ is also valid with respect to $\mathcal{G}[\mathbf{A}]$. The correctness of Alg. 2 relies solely on the balanced property, which is satisfied in the PMG with local transformations as per Thm. 8 and lem. 31. Thus, the algorithm is also sound with respect to $\mathcal{Q}[\mathbf{A}]$. □

E. Proofs

In this section, we provide detailed proofs of the propositions and theorems presented in the main body of the paper.

Theorem 2. *Let $\mathcal{M}$ be a MAG over the set of variables $\mathbf{V}$. Given information $[\![\mathcal{M}, Y]\!]$, a set $\mathbf{X} \subseteq \mathbf{V} \setminus \{Y\}$ is a MIS relative to $[\![\mathcal{M}, Y]\!]$ if and only if $\mathbf{X} \subseteq \texttt{An}(Y)_{\mathcal{M}_{\overline{\mathbf{X}}}}$ holds.*

Proof. (**If**) Suppose that $\mathbf{X}$ is *not* a MIS relative to $[\![\mathcal{M}, Y]\!]$. This implies that there exists some $\mathbf{X}' \subsetneq \mathbf{X}$ such that $\mu_{\mathbf{x}[\mathbf{X}']} = \mu_{\mathbf{x}}$ for every SCM conforming to the MAG $\mathcal{M}$. For the sake of contradiction, assume that $\mathbf{X} \subseteq \texttt{An}(Y)_{\mathcal{M}_{\overline{\mathbf{X}}}}$. To derive a contradiction, it suffices to construct a SCM such that $\mu_{\mathbf{x}[\mathbf{X}']} \neq \mu_{\mathbf{x}}$. Consider the causal diagram $\mathcal{G}$ generated by the following procedure:

Step 1. If $A \rightarrow B$ in $\mathcal{M}$, then add a directed edge $A \rightarrow B$ to $\mathcal{G}$.

Step 2. If $A \leftrightarrow B$ in $\mathcal{M}$, then add a bidirected edge $A \leftrightarrow B$ to $\mathcal{G}$.

From this construction, it is clear that the causal diagram $\mathcal{G}$ corresponds to $\mathcal{M}$. Furthermore, we have $\mathbf{X} \subseteq \texttt{An}(Y)_{\mathcal{G}_{\overline{\mathbf{X}}}}$ since $\mathcal{G}$ and $\mathcal{M}$ have the exact same edges.

Now consider the following SCM associated with $\mathcal{G}$: Each variable in $V_i \in \mathbf{V}(\mathcal{G})$ is associated with a unique latent variable U_i and the function of each endogenous variable in $\mathbf{V}(\mathcal{G})$ is the sum of the value of its parents. Since $\mathbf{X} \subseteq \texttt{An}(Y)_{\mathcal{G}_{\overline{\mathbf{X}}}}$ holds, there exist directed paths from $\mathbf{X} \setminus \mathbf{X}'$ to Y without passing through $\mathbf{X}'$. Let $\mathbf{W} = \mathbf{X} \setminus \mathbf{X}'$. Then, setting $\mathbf{W}$ to $\mathbb{E}[\mathbf{W} \mid do(\mathbf{x}')] + 1$ results in a larger outcome value for Y, i.e., $\mu_{\mathbf{x}} = \mu_{\mathbf{w},\mathbf{x}'} > \mu_{\mathbf{x}[\mathbf{X}']}$, which leads to a contradiction.

(**Only if**) Suppose that $\mathbf{X} \not\subseteq \texttt{An}(Y)_{\mathcal{M}_{\overline{\mathbf{X}}}}$ holds. This indicates that there exists a nonempty subset $\mathbf{Z} \triangleq \mathbf{X} \setminus \texttt{An}(Y)_{\mathcal{M}_{\overline{\mathbf{X}}}}$. Let $\mathbf{X}' = \mathbf{X} \setminus \mathbf{Z}$. Our goal is to show that Y and $\mathbf{Z}$ are m-separated by $\mathbf{X}'$ in $\mathcal{M}_{\overline{\mathbf{X}}}$. Once established, we can apply Rule 3 of do-calculus for MAGs (Zhang, 2008b) to derive $\mu_{\mathbf{x}'} = \mu_{\mathbf{x}',\mathbf{z}}$.

For contradiction, assume that there exists some variable $Z \in \mathbf{Z}$ such that Z and Y are m-connected conditioning on $\mathbf{X}'$ in $\mathcal{M}_{\overline{\mathbf{X}}}$. This means the existence of a m-connected path p between Z and Y. Since Z has its incoming edges removed, p must start with an edge outgoing from Z. If there were any collider along the path, it would be m-separated, as the collider cannot be an ancestor of a conditioned node $\mathbf{X}'$. However, if the path p begins with an outgoing edge from Z and has no colliders, then it must be a directed path from Z to Y. This implies that $Z \in \texttt{An}(Y)_{\mathcal{M}_{\overline{\mathbf{X}}}}$ holds, thus $\mathbf{Z}$ and Y are not m-separated by $\mathbf{X}'$ in $\mathcal{M}_{\overline{\mathbf{X}}}$, leading to a contradiction. Consequently, we have that $\mathbf{X}$ is not a MIS relative to $[\![\mathcal{M}, Y]\!]$. □

Proposition 1. *Let $\mathcal{M}$ be a MAG over the set of variables $\mathbf{V}$. A set $\mathbf{X} \subseteq \mathbf{V} \setminus \{Y\}$ is a MIS relative to $[\![\mathcal{M}, Y]\!]$ if and only if there exists a causal diagram $\mathcal{G}$ conforming to $\mathcal{M}$ such that $\mathbf{X}$ is a MIS relative to $[\![\mathcal{G}, Y]\!]$.*

Proof. (**If**) Let $\mathbf{X}$ be a MIS relative to $[\![\mathcal{G}, Y]\!]$ for some causal diagram $\mathcal{G}$ conforming to $\mathcal{M}$. By the definition of MIS for causal diagrams in Def. 1, there is no $\mathbf{X}' \subsetneq \mathbf{X}$ such that for all SCM conforming to $\mathcal{G}$, $\mu_{\mathbf{x}[\mathbf{X}']} = \mu_{\mathbf{x}}$. In other words, for every $\mathbf{X}' \subsetneq \mathbf{X}$, there exists an SCM $\mathcal{S}$ conforming to $\mathcal{G}$ such that $\mu_{\mathbf{x}[\mathbf{X}']} \neq \mu_{\mathbf{x}}$. Since any SCM conforming to $\mathcal{G}$ also conforms to $\mathcal{M}$, we know that $\mathcal{S}$ also conforms to $\mathcal{M}$. Thus, for any proper subset $\mathbf{X}' \subsetneq \mathbf{X}$, there exists an SCM associated with $\mathcal{M}$ in which $\mu_{\mathbf{x}[\mathbf{X}']} = \mu_{\mathbf{x}}$ holds.

(**Only if**) Let $\mathbf{X}$ be a MIS relative to $[\![\mathcal{M}, Y]\!]$. The causal diagram $\mathcal{G}$ constructed in the same manner as in the proof of thm. 2 conforms to $\mathcal{M}$ and satisfies $\mathbf{X} \subseteq \mathtt{An}(Y)_{\mathcal{G}_{\overline{\mathbf{X}}}}$. Therefore, we can conclude that $\mathbf{X}$ is a MIS relative to $[\![\mathcal{G}, Y]\!]$ supported by Prop. 10. $\square$

Proposition 2. *Let $\mathcal{P}$ be a PAG over the set of variables $\mathbf{V}$. A set $\mathbf{X} \subseteq \mathbf{V} \setminus \{Y\}$ is a MIS relative to $[\![\mathcal{P}, Y]\!]$ if and only if, for every variable $X \in \mathbf{X}$, there exists a proper possibly-directed path from X to Y with respect to $\mathbf{X}$ in $\mathcal{P}$.*

Proof. (**If**) Suppose that $\mathbf{X}$ is *not* a MIS relative to $[\![\mathcal{P}, Y]\!]$, which implies that there exists some proper subset $\mathbf{X}' \subsetneq \mathbf{X}$ such that $\mu_{\mathbf{x}[\mathbf{X}']} = \mu_{\mathbf{x}}$ for every SCM conforming to $\mathcal{P}$. For contradiction, suppose that for all $X \in \mathbf{X}$, there exist proper possibly-directed paths from X to Y with respect to $\mathbf{X}$ in $\mathcal{P}$. Let $\mathbf{W} = \mathbf{X} \setminus \mathbf{X}'$ and W be a vertex in $\mathbf{W}$. Suppose that p is an uncovered proper possibly-directed path from W to Y with respect to $\mathbf{X}$ in $\mathcal{P}$. Let $\mathcal{M} \in [\mathcal{P}]$ be a MAG constructed by the following procedure:

Step 1. Orient all edges along p as directed edges.

Step 2. Orient the remaining edges according to Lem. 6.

Then, p corresponds to a proper directed path from W to Y with respect to $\mathbf{X}$ in $\mathcal{M}$. Thus, $W \in \mathtt{An}(Y)_{\mathcal{M}_{\overline{\mathbf{X}}}}$ holds. We can then use the same construction in the proof of Thm. 2. In the constructed causal diagram $\mathcal{G}$, $W \in \mathtt{An}(Y)_{\mathcal{G}_{\overline{\mathbf{X}}}}$ holds. Furthermore, we know there exists an SCM $\mathcal{S}$ in which W has a positive causal effect on Y which is not mediated by any variable in $\mathbf{X}$. Thus, setting $\mathbf{W}$ to $\mathbb{E}[\mathbf{W} \mid do(\mathbf{x}')] + 1$ will result in a larger outcome for Y, i.e., $\mu_{\mathbf{x}} = \mu_{\mathbf{w},\mathbf{x}'} > \mu_{\mathbf{x}[\mathbf{X}']}$, meaning $\mu_{\mathbf{x}} \neq \mu_{\mathbf{x}[\mathbf{X}']}$, which contradicts the statement: $\mu_{\mathbf{x}[\mathbf{X}']} = \mu_{\mathbf{x}}$ for every SCM conforming to $\mathcal{P}$.

(**Only if**) Suppose that for some $Z \in \mathbf{X}$, there is no proper possibly directed path from Z to Y with respect to $\mathbf{X}$ in $\mathcal{P}$. Let $\mathbf{X}' = \mathbf{X} \setminus \{Z\}$. We aim to show that $P(y \mid do(\mathbf{x}')) = P(y \mid do(\mathbf{x}', z))$, which would imply $\mu_{\mathbf{x}'} = \mu_{\mathbf{x}',z}$. Unfortunately, we cannot apply Rule 3 of do-calculus for PAGs, since it is not guaranteed that X and Y are definitely m-separated by $\mathbf{X}'$ in $\mathcal{P}_{\overline{\mathbf{X}}}$. However, we can reason over the MAGs in the Markov equivalence class represented by $\mathcal{P}$.

All paths from Z to Y in $\mathcal{P}$ which do not pass through $\mathbf{X}$ must not be a directed path due to our assumption, i.e., they all contain an arrowhead pointing towards Z. Let $\mathcal{M}$ be a MAG conforming to $\mathcal{P}$. Then, all paths from Z to Y in $\mathcal{M}$ which do not pass through $\mathbf{X}$ must also be non-directed. Thus, using similar reasoning as in the proof of Thm. 2, Z and Y are m-separated by $\mathbf{X}'$ in $\mathcal{M}_{\overline{\mathbf{X}}}$. This is because any path out of Z to Y must contain a collider node, which must be blocked, since it cannot be an ancestor of any conditioned node. Therefore, we conclude that $P(y \mid do(\mathbf{x}')) = P(y \mid do(\mathbf{x}', z))$. Since this argument holds for every MAG conforming to $\mathcal{P}$, it holds for all SCMs conforming to $\mathcal{P}$. $\square$

Proposition 4. *Let $\mathcal{D}$ be either a causal diagram or a MAG. If a set $\mathbf{X}$ is a MIS with respect to $[\![\mathcal{D}, Y]\!]$, then $\mathbf{X}$ is a DMIS with respect to $[\![\mathcal{D}, Y]\!]$.*

Proof. Without loss of generality, assume that all nodes in $\mathcal{D}$ are ancestors of Y. For contradiction, assume that $\mathbf{X}$ is a MIS but *not* a DMIS relative to $[\![\mathcal{D}, Y]\!]$. By Thm. 2 and prop. 10, we have $\mathbf{X} \subseteq \mathtt{An}(Y)_{\mathcal{D}_{\overline{\mathbf{X}}}}$. Then, we can consider an SCM $\mathcal{S}^*$ compatible with $\mathcal{D}$, where all mechanisms consist of the sum of the values of their parents, i.e., $f_i = \sum_{|\mathbf{Pa}_i|} \mathbf{pa}_i + \mathbf{u}_i$. Let $\mathbf{X}'$ be an arbitrary proper subset of $\mathbf{X}$, and $\mathbf{W}$ denote $\mathbf{X} \setminus \mathbf{X}'$. Such a model $\mathcal{S}^*$ always ensures that setting $\mathbf{W}$ as $\mathbb{E}[\mathbf{W} \mid do(\mathbf{x}')] + 1$ results in $\mu_{\mathbf{x}} = \mu_{\mathbf{w},\mathbf{x}'} > \mu_{\mathbf{x}[\mathbf{X}']}$ for any proper subset $\mathbf{X}'$ since there exist directed paths from each $W \in \mathbf{W}$ to Y without passing through $\mathbf{X}'$. The existence of $\mathcal{S}^*$ leads to a contradiction. $\square$

Proposition 5. *Let $\mathcal{P}$ be a PAG over the set of variables $\mathbf{V}$. A set $\mathbf{X} \subseteq \mathbf{V} \setminus \{Y\}$ is in $\mathbb{D}_{\mathcal{P},Y}$ if and only if there exists a MAG $\mathcal{M}$ conforming to $\mathcal{P}$ such that $\mathbf{X} \in \mathbb{M}_{\mathcal{M},Y}$.*

Proof. (**If**) Suppose $\mathbf{X} \subseteq \mathbf{V} \setminus \{Y\}$ be a MIS relative to $[\![\mathcal{P}, Y]\!]$, and there exists a MAG $\mathcal{M}$ conforming to $\mathcal{P}$ where $\mathbf{X}$ is a MIS relative to $[\![\mathcal{M}, Y]\!]$. By Prop. 4, $\mathbf{X}$ is a DMIS relative to $[\![\mathcal{M}, Y]\!]$. Hence, there exists an SCM $\mathcal{S}$ such that for any proper subset $\mathbf{X}'$, $\mu_{\mathbf{x}[\mathbf{X}']} \neq \mu_{\mathbf{x}}$ holds. Since $\mathcal{S}$ conforms to $\mathcal{M}$, it also conforms to $\mathcal{P}$, thus concluding proof for this direction.

(**Only if**) Suppose $\mathbf{X} \subseteq \mathbf{V} \setminus \{Y\}$ be a DMIS relative to $[\![\mathcal{P}, Y]\!]$. By the definition of DMIS (4), there exists an SCM $\mathcal{S}$ associated with $\mathcal{P}$ such that, for every $\mathbf{X}' \subsetneq \mathbf{X}$, $\mu_{\mathbf{x}[\mathbf{X}']} \neq \mu_{\mathbf{x}}$ holds. Therefore, $\mathbf{X}$ is a MIS, since for any proper subset $\mathbf{X}'$, $\mu_{\mathbf{x}[\mathbf{X}']} \neq \mu_{\mathbf{x}}$ holds under the SCM $\mathcal{S}$. $\square$

Theorem 3. *Let $\mathcal{P}$ be a PAG over the set of variables $\mathbf{V}$. A set $\mathbf{X} \subseteq \mathbf{V} \setminus \{Y\}$ is a DMIS relative to $[\![\mathcal{P}, Y]\!]$ if and only if, for any pair of vertices $X, Z \in \mathbf{X}$, there exist uncovered proper possibly-directed paths from X and Z to Y with respect to $\mathbf{X}$ such that their starting edges are not relevant.*

Proof. (**If**) Let p_X denote an uncovered proper possibly-directed path from X to Y with respect to $\mathbf{X}$ in $\mathcal{P}$. Suppose that $\mathbf{X}$ is *not* a DMIS, implying that, for all MAGs $\mathcal{M} \in [\mathcal{P}]$, it holds that $Z \notin \mathtt{An}(Y)_{\mathcal{M}_{\overline{\mathbf{X}}}}$ and $X \in \mathtt{An}(Y)_{\mathcal{M}_{\overline{\mathbf{X}}}}$ without loss of generality. In other words, if orienting p_X as $X \to \cdots \to Y$ is valid, it follows that orienting any possibly directed path from Z to Y as $Z \to \cdots \to Y$ is invalid in all MAGs conforming to $\mathcal{P}$. We will show that the starting edge of p_X is relevant to the starting edge of any uncovered possibly-directed path from Z to Y in $\mathcal{P}$.

Let p_Z be an arbitrary uncovered proper possibly-directed path from Z to Y with respect to $\mathbf{X}$ in $\mathcal{P}$. Note that such a path always exists, as established by Lem. 7. We know that the path p_Z must begin with one of the following edges: $\circ\!\!-\!\!\circ$, $\circ\!\!\rightarrow$, or $\rightarrow$. We will show that p_Z can *only* start with a circle edge ($\circ\!\!-\!\!\circ$).

(p_Z **only starts with a circle edge ($\circ\!\!-\!\!\circ$)).** Suppose p_Z starts with $\,^{?}\!\!\rightarrow$. Then, the path must take the form $Z \,^{?}\!\!\rightarrow \cdot \rightarrow \cdots \rightarrow Y$ in $\mathcal{P}$ by Lem. 8. In this case, we can construct a valid $\mathcal{M}$ by orienting any circle marks ($\circ$) along the path as tails ($-$) following Lem. 6. This contradicts the assumption that there is no MAG conforming $\mathcal{P}$ in which p_Z is a directed path from Z to Y. Therefore, we conclude that p_Z only can be $Z \circ\!\!-\!\!\circ \cdots *\!\!-\!\!* Y$.

For the sake of contradiction, assume $e_X(X *\!\!-\!\!* X')$ is not relevant to $e_Z(Z' \circ\!\!-\!\!\circ Z)$ where each denotes the starting edges of p_X and p_Z respectively; Then, we consider the following two cases separately: ① X and Z are *not* in the same bucket, or ② they are in the same bucket, and every circle path including e_X and e_Z is not uncovered, i.e., they are *not* relevant.

(**① X and Z do not belong to the same bucket).** Consider the orientation according to Lem. 6. In the second step of the construction, we always have a MAG $\mathcal{M}$ containing $Z \to Z'$ by the completeness of orientation in PAGs, which indicates p_Z corresponds to a directed path from Z to Y in $\mathcal{M}$, as it is uncovered. Therefore, we can construct a valid $\mathcal{M}$ according to Lem. 6, contradicting the assumption that $Z \notin \mathtt{An}(Y)_{\mathcal{M}_{\overline{\mathbf{X}}}}$ for all MAGs $\mathcal{M} \in [\mathcal{P}]$.

(**② X and Z are in the same bucket).** Suppose that X and Z are in the same bucket. Let $V_1(= X) \circ\!\!-\!\!\circ V_2(= X') \circ\!\!-\!\!\circ \cdots \circ\!\!-\!\!\circ V_{n-1}(= Z') \circ\!\!-\!\!\circ V_n(= Z)$ be an arbitrary non-uncovered circle path between X and Z in $\mathcal{P}$. By the definition of an uncovered circle path, such a path must include at least one non-uncovered triple $\langle V_i, V_{i+1}, V_{i+2}\rangle$ on the circle path. The existence of an edge between $V_i \circ\!\!-\!\!\circ V_{i+2}$ would induce an uncovered circle path $V_1 \circ\!\!-\!\!\circ \cdots \circ\!\!-\!\!\circ V_i \circ\!\!-\!\!\circ V_{i+2} \circ\!\!-\!\!\circ \cdots \circ\!\!-\!\!\circ V_n$. To avoid this, X and Z must be adjacent, and furthermore, the edge connecting X and Z must appear as a circle edge $X \circ\!\!-\!\!\circ Z$ by Lem. 11.

The existence of the edge $X \circ\!\!-\!\!\circ Z$ implies that there must be edges $X \circ\!\!-\!\!\circ V_i$ for all $3 \leq i \leq n-1$, or $Z \circ\!\!-\!\!\circ V_i$ for all $2 \leq i \leq n-2$ by chordality. In the former case, we orient the subgraph of $\mathcal{P}$ over $\{V_1, \cdots, V_n\}$ following a similar approach to the proof of Lemma 7.6 in Maathuis and Colombo (2015). We begin by selecting a vertex V_2 and orient all edges incident to V_2 as directed into V_2. Since the subgraph is chordal and V_2 is simplicial, this orientation does not create any uncovered colliders in the subgraph. We then remove V_2 and the oriented edges from the subgraph. The resulting graph remains chordal and therefore again choose a vertex V_3, and orient any edges incident to V_3 into V_3. We continue this procedure until all edges are oriented. The constructed subgraph does not create any directed cycle, almost directed cycle, or uncovered collider, thus it is valid orientations. Since $X \to X' \to \cdots \to Y$ is valid, we have a directed path $Z \to Z' \to \cdots \to X' \to \cdots \to Y$ which leads to a contradiction.

In the latter case, we can similarly orient the edges, starting from V_{n-1} and proceeding to V_2. Furthermore, this procedure can also be extended to cases where the graph takes on a superimposed form.

(**Only if**) Suppose that e_X is relevant to e_Z in $\mathcal{P}$. It follows that $V_1(= X) \circ\!\!-\!\!\circ V_2(= X') \circ\!\!-\!\!\circ \cdots \circ\!\!-\!\!\circ V_{n-1}(= Z') \circ\!\!-\!\!\circ V_n(=$

Z) is an uncovered circle path. For the sake of contradiction, assume that $\mathbf{X}$ is a DMIS relative to $[\![\mathcal{P}, Y,]\!]$. Then, there exists a MAG $\mathcal{M}$ conforming to $\mathcal{P}$ such that both p_X and p_Z are proper directed paths with respect to $\mathbf{X}$ in $\mathcal{M}$. Therefore, we can orient $V_1 \circ\!\!-\!\!\circ V_2$ as $V_1 \to V_2$, and $V_n \circ\!\!-\!\!\circ V_{n-1}$ as $V_n \to V_{n-1}$ to construct $\mathcal{M}$ from $\mathcal{P}$. Furthermore, since the circle path is uncovered, $V_i \circ\!\!-\!\!\circ V_{i+1}$ must be oriented $V_i \to V_{i+1}$ for $i = 2, \cdots, n-2$. However, this orientation introduces a new uncovered collider $V_{n-2} \to V_{n-1} \leftarrow V_n$, which leads to a contradiction. $\square$

Proposition 6. *Let $\mathcal{M}$ be a MAG over the set of variables $\mathbf{V}$. A set $\mathbf{X} \subseteq \mathbf{V} \setminus \{Y\}$ is a POMIS relative to $[\![\mathcal{M}, Y]\!]$ if and only if there exists a causal diagram $\mathcal{G}$ conforming to $\mathcal{M}$ such that $\mathbf{X}$ is a POMIS relative to $[\![\mathcal{G}, Y]\!]$.*

Proof. (**If**) Suppose that $\mathbf{X}$ is a POMIS relative to $[\![\mathcal{G}, Y]\!]$ for some $\mathcal{G}$ conforming to $\mathcal{M}$. Then there exists an SCM $\mathcal{S}$ conforming to $\mathcal{G}$ such that $\mu_{\mathbf{x}^*} > \forall_{\mathbf{W} \in \mathbb{M}_{\mathcal{G},Y} \setminus \{\mathbf{X}\}} \mu_{\mathbf{w}^*}$. Since any SCM conforming to $\mathcal{G}$ also conforms to $\mathcal{M}$, the SCM also conforms to $\mathcal{M}$, the SCM $\mathcal{S}$ also conforms to $\mathcal{M}$, and thus $\mathbf{X}$ is a POMIS relative to $[\![\mathcal{M}, Y]\!]$.

(**Only if**) Let $\mathbf{X}$ be a POMIS relative to $[\![\mathcal{M}, Y]\!]$. Then there exists an SCM $\mathcal{S}$ conforming to $\mathcal{M}$ such that $\mu_{\mathbf{x}^*} > \forall_{\mathbf{W} \in \mathbb{M}_{\mathcal{M},Y} \setminus \{\mathbf{X}\}} \mu_{\mathbf{w}^*}$. Let $\mathcal{G}$ be the causal diagram associated with $\mathcal{S}$. This $\mathcal{G}$ corresponds to $\mathcal{M}$, and $\mathbf{X}$ is a POMIS relative to $[\![\mathcal{G}, Y]\!]$, concluding the proof for this direction. $\square$

Theorem 4. *Let $\mathcal{M}$ be a MAG over the set of variable $\mathbf{V}$. A set $\mathbf{X} \subseteq \mathbf{V} \setminus \{Y\}$ is a POMIS relative to $[\![\mathcal{M}, Y]\!]$ if and only if $\mathbf{X} = \mathsf{IB}(\mathcal{M}, Y, \mathbf{X})$.*

Proof. **(Only if)** We will show contrapositive, i.e., if $\mathbf{X} = \mathsf{IB}(\mathcal{M}, Y, \mathbf{X})$ does not hold, then $\mathbf{X}$ is not a POMIS relative to $[\![\mathcal{M}, Y]\!]$. We denote $\mathbf{W} = \mathsf{IB}(\mathcal{M}, Y, \mathbf{X})$ and $\mathbf{T} = \mathsf{MUCT}(\mathcal{M}, Y, \mathbf{X})$, assuming $\mathbf{X} \neq \mathbf{W}$. Let $\mathbf{W}' \triangleq \mathbf{W} \setminus \mathbf{X}$. Before proceeding with the main proof, we first establish that the following conditional independence statement holds:

Claim 1. $(Y \perp\!\!\!\perp \mathbf{W}' \mid \mathbf{X})$ holds in $\mathcal{M}_{\overline{\mathbf{X}}\underline{\mathbf{W}'}}$.

Proof. Suppose that the negation of this statement holds: $(Y \not\perp\!\!\!\perp \mathbf{W}' \mid \mathbf{X})$ in $\mathcal{M}_{\overline{\mathbf{X}}\underline{\mathbf{W}'}}$. This would imply that there exists an m-connected path from some $W \in \mathbf{W}'$ to Y given $\mathbf{X}$ in $\mathcal{M}_{\overline{\mathbf{X}}\underline{\mathbf{W}'}}$. For the m-connected path to exist, there must be no colliders, as no node along the path can be an ancestor of $\mathbf{X}$ due to all incoming edges to $\mathbf{X}$ being cut in $\mathcal{M}_{\overline{\mathbf{X}}}$. Moreover, as all outgoing edges from $\mathbf{W}'$ are cut in $\mathcal{M}_{\underline{\mathbf{W}'}}$, the path cannot begin with an edge going out of W. Therefore, we get that the m-connected path must be of the following form: $W \leftarrow W_1 \leftarrow \cdots \leftarrow W_n \leftrightarrow R_1 \to \cdots \to R_m \to Y$ with $n, m \geq 0$ where no node along the path can be in $\mathbf{W}$; otherwise, it would either be part of $\mathbf{X}$, since we are conditioning on $\mathbf{X}$, or in $\mathbf{W}'$, in which case all of its outgoing arrows would have been removed. Since Y is contained in $\mathbf{T}$, the parent of Y, R_m, along the path must be either in $\mathbf{T}$ or $\mathbf{W}$. However, as previously argued, no node along the path can be in $\mathbf{W}$; therefore, it must be in $\mathbf{T}$. This reasoning can be applied iteratively up to R_1, implying that R_1 is also in $\mathbf{T}$. Since $\mathbf{T}$ is closed under PC, the inclusion of R_1 in $\mathbf{T}$ implies that W_n must also be in $\mathbf{T}$. Additionally, because $\mathbf{T}$ is closed under descendants, $W_{n-1}, \cdots, W_1$ must also be in $\mathbf{T}$. Consequently, W must be in $\mathbf{T}$ as well. However, this leads to a contradiction, since W is in $\mathbf{W}$, and $\mathbf{W}$ and $\mathbf{T}$ are disjoint by definition. Therefore, the conditional independence statement $(Y \perp\!\!\!\perp \mathbf{W}' \mid \mathbf{X})$ must hold in $\mathcal{M}_{\overline{\mathbf{X}}\underline{\mathbf{W}'}}$. $\square$

Claim 2. $(Y \perp\!\!\!\perp \mathbf{X}' \mid \mathbf{W})$ holds in $\mathcal{M}_{\overline{\mathbf{W}},\underline{\mathbf{X}'}}$ where $\mathbf{X}' \triangleq \mathbf{X} \setminus \mathbf{W}$.

Proof. Suppose this statement is *false*, i.e., $(Y \not\perp\!\!\!\perp \mathbf{X}' \mid \mathbf{W})$ holds in $\mathcal{M}_{\overline{\mathbf{W}},\underline{\mathbf{X}'}}$. Then, there exists an m-connected path from some $X \in \mathbf{X}'$ to Y given $\mathbf{W}$ in $\mathcal{M}_{\overline{\mathbf{W}},\underline{\mathbf{X}'}}$. Since all edges into $\mathbf{X}'$ are removed, the path must begin with an edge going out of X. The path cannot contain any colliders, as no node can be an ancestor of a node in the conditioned set $\mathbf{W}$, given that all incoming edges to $\mathbf{W}$ are cut. Thus, all edges along the path must be directed, pointing to Y: $X \to W_1 \to \cdots \to W_n \to Y (n \geq 0)$ where no node along the path can be in $\mathbf{W}$, since we are conditioning on $\mathbf{W}$. The parent of Y, W_n, along the path must be either in $\mathbf{T}$ or $\mathbf{W}$, as Y in T. However, as previously argued, no node along the path can be included in $\mathbf{W}$, which means it must be in $\mathbf{T}$. This reasoning can be applied iteratively up to W_1, implying that W_1 is also in $\mathbf{T}$. Therefore, X must be a parent of a node in $\mathbf{T}$, implying that X is in $\mathbf{W}$. This leads to a contradiction for $X \in \mathbf{X} \setminus \mathbf{W}$. $\square$

We are now ready to proceed to the main proof. We will show that $\mathbf{X}$ is *not* a POMIS by proving that $\mu_{\mathbf{x}^*} \leq \mu_{\mathbf{w}^*}$ in every SCM conforming to $\mathcal{M}$. We derive that the following holds:

$$
\begin{aligned}
\mu_{\mathbf{x}^*} &= \mathbb{E}[Y \mid do(\mathbf{x}^*)] \\
&= \sum_{\mathbf{w}'} \mathbb{E}[Y \mid do(\mathbf{x}^*), \mathbf{w}'] P(\mathbf{w}' \mid do(\mathbf{x}^*)) \\
&= \sum_{\mathbf{w}'} \mathbb{E}[Y \mid do(\mathbf{x}^*), do(\mathbf{w}')] P(\mathbf{w}' \mid do(\mathbf{x}^*)) && \because \textbf{Claim 1} \\
&= \sum_{\mathbf{w}'} \mathbb{E}[Y \mid do(\mathbf{x}^*[\mathbf{W}]), do(\mathbf{w}')] P(\mathbf{w}' \mid do(\mathbf{x}^*)) && \because \textbf{Claim 2} \\
&\leq \sum_{\mathbf{w}'} \mathbb{E}[Y \mid do(\mathbf{w}')] P(\mathbf{w}' \mid do(\mathbf{x}^*)) \\
&= \mathbb{E}[Y \mid do(\mathbf{w}^*)] \\
&= \mu_{\mathbf{w}^*}.
\end{aligned}
$$

Therefore, $\mathbf{X}$ is *not* a POMIS with respect to $[\![\mathcal{M}, Y]\!]$, which completes the proof.

(**If**) To prove this direction, we will show that if $\mathbf{X} = \mathsf{IB}(\mathcal{M}, Y, \mathbf{X})$, then $\mathbf{X}$ is a POMIS relative to $[\![\mathcal{M}, Y]\!]$. Suppose that $\mathbf{X} = \mathsf{IB}(\mathcal{M}, Y, \mathbf{X})$ holds. By Prop. 6, it suffices to show that there exists a causal diagram $\mathcal{G}$ such that $\mathbf{X}$ is a POMIS relative to $[\![\mathcal{G}, Y]\!]$. Consider the causal diagram $\mathcal{G}$ constructed by the following lemma:

Lemma 34. *Let $\mathcal{M}$ be a MAG. Let $\mathcal{G}$ be the graph resulting from the following procedure applied to $\mathcal{M}$.*

Step 1. For each visible edge $A \to B$ in $\mathcal{M}$, add $A \to B$ in $\mathcal{G}$.

Step 2. For each bidirected edge $A \leftrightarrow B$ in $\mathcal{M}$, add $A \leftrightarrow B$ in $\mathcal{G}$.

Step 3. For each invisible directed edge $A \to B$ in $\mathcal{M}$, if it is the unique invisible edge among directed edges outgoing from A in $\mathcal{M}$, then add both a directed edge $A \to B$ and bidirected edge $A \leftrightarrow B$ to $\mathcal{G}$.

Step 4. Let $\mathbf{T}_{\mathcal{G}} \triangleq \mathsf{MUCT}(\mathcal{G}, Y)$. Consider all nodes A for which there are invisible edges outgoing from A in $\mathcal{M}$.

1. *If there exists $B \in \mathtt{Ch}(A)_{\mathcal{M}}$ that is contained in $\mathbf{T}_{\mathcal{G}}$, add both a directed edge $A \to B$ and bidirected edge $A \leftrightarrow B$, and add directed edges $A \to C$ for all $C \in \mathtt{Ch}(A)_{\mathcal{M}} \setminus \{B\}$.*
2. *Otherwise, if there is no intersection with $\mathbf{T}_{\mathcal{G}}$, add directed edges $A \to C$ for all $C \in \mathtt{Ch}(A)_{\mathcal{M}}$.*

This step is repeated with the updated $\mathbf{T}_{\mathcal{G}} \leftarrow \mathsf{MUCT}(\mathcal{G}, Y)$ as long as $\mathcal{G}$ remains unchanged.

Then, the result graph $\mathcal{G}$ is a causal diagram conforming to $\mathcal{M}$.

Proof. We need to show that $\mathcal{G}$ and $\mathcal{M}$ have the same ancestral relations, and the same conditional independence relations.

(**① $\mathcal{G}$ and $\mathcal{M}$ have the same ancestral relations).** This is evident, as each directed edge is added to $\mathcal{G}$ *if and only if* it also exists in $\mathcal{M}$.

(**② $\mathcal{G}$ and $\mathcal{M}$ encode the same independence relations).** The graphs $\mathcal{G}$ and $\mathcal{M}$ differ only in the bidirected edges added to $\mathcal{G}$ corresponding to invisible edges in $\mathcal{M}$. Thus, it suffices to show that these additional bidirected edges added to $\mathcal{G}$ do not encode any additional independence between variables. Therefore, we need to show that these edges do not create any new uncovered colliders.

Consider a bidirected edge $A \leftrightarrow B$ added to $\mathcal{G}$ in **Step 3**. For this added edge to create a collider, there must be either a directed edge incoming to A (i.e., $C \to A \leftrightarrow B$), or bidirected edge incoming to A (i.e., $C \leftrightarrow A \leftrightarrow B$) in $\mathcal{G}$. In both cases, B and C are adjacent in $\mathcal{M}$, since $A \to B$ is invisible in $\mathcal{M}$ by Lem. 19. Therefore, this collider at A does not introduce any new independence.

Now consider a bidirected edge $A \leftrightarrow B$ added to $\mathcal{G}$ in **Step 4**. The previous argument can be reused here to argue that this edge does not encode any new independence, since we add only one bidirected among outgoing directed edges from A. For clarity, suppose that we have a MAG $\mathcal{M} = \langle A \to B, A \to B, A \to D\rangle$ where B,C, and D are mutually not adjacent in $\mathcal{M}$. Adding at most one of $A \leftrightarrow C$, $A \leftrightarrow B$, or $A \leftrightarrow D$ does not introduce a new collider at A, thereby preserving conditional independence. $\square$

Let $\mathcal{G}$ be the causal diagram constructed following Lem. 34. We will prove that $\mathbf{X}$ is a POMIS with respect to $[\![\mathcal{G}, Y]\!]$. Let X be any variable in $\mathbf{X}$. Then X is a parent of some $T \in \mathsf{MUCT}(\mathcal{M}, Y, \mathbf{X})$ in $\mathcal{M}$. It suffices to show that $T \in \mathsf{MUCT}(\mathcal{G}_{\overline{\mathbf{X}}}, Y)$ since this means that X is a parent of a member of $\mathsf{MUCT}(\mathcal{G}_{\overline{\mathbf{X}}}, Y)$, and is therefore in $\mathsf{IB}(\mathcal{G}_{\overline{\mathbf{X}}}, Y)$[11].

Let $\mathbf{T}_\mathcal{G} \triangleq \mathsf{MUCT}(\mathcal{G}, Y)$ and $\mathbf{T}_\mathcal{M} \triangleq \mathsf{MUCT}(\mathcal{M}, Y, \emptyset)$. We will show that $\mathbf{T}_\mathcal{M} \subseteq \mathbf{T}_\mathcal{G}$. Let T be a node in $\mathbf{T}_\mathcal{M}$. We know such a node always exists because Y is in both $\mathbf{T}_\mathcal{G}$ and $\mathbf{T}_\mathcal{M}$. Let $\mathcal{H} \triangleq \mathcal{G}[\mathtt{An}(Y)_\mathcal{G}]$ and $\mathcal{N} \triangleq \mathcal{M}[\mathtt{An}(Y)_\mathcal{M}]$. Since $\mathcal{M}$ and $\mathcal{G}$ share the same skeleton and the same ancestral relations among vertices, it follows that $\mathtt{An}(Y)_\mathcal{M} = \mathtt{An}(Y)_\mathcal{G}$, implying $\mathbf{V}(\mathcal{H}) = \mathbf{V}(\mathcal{N})$.

(① If $W \in \mathsf{PC}(T)_\mathcal{N}$, then $W \in \mathbf{T}_\mathcal{G}$). Suppose that another node W is in the same pc-component of T in $\mathcal{N}$, i.e., $W \in \mathsf{PC}(T)_\mathcal{N}$. This implies that there exists a path between T and W in $\mathcal{N}$ such that (i) all non-endpoint nodes along the path are colliders, and (ii) none of the edges are visible.

> For all directed edges $U \to V$ along this path, if there does not exist an edge $U \to Z(\neq V)$ in $\mathcal{N}$, a bidirected edge $U \leftrightarrow V$ is added to $\mathcal{G}$ in Step 3. Consequently, T and W are in the same c-component in $\mathcal{H}$.
>
> Otherwise, if there is some directed edge $U \to V$ along the path for which there exists $U \to Z(\neq V)$, then from Step 4, we know that one of these outgoing edges from U will have a corresponding bidirected edge in $\mathcal{H}$ which adds U to $\mathbf{T}_\mathcal{G}$. Since MUCT is closed under descendants, all descendants of U are also included in MUCT as well.

This logic applies along the entire path, ensuring that $T \in \mathbf{T}_\mathcal{G} \Rightarrow W \in \mathbf{T}_\mathcal{G}$.

(② If $W \in \mathtt{De}(T)_\mathcal{N}$, then $W \in \mathbf{T}_\mathcal{G}$). Now, suppose that W is a descendant of T in $\mathcal{N}$, i.e., $W \in \mathtt{De}(T)_\mathcal{N}$. Then W is a descendant of T in $\mathcal{H}$ as well, and so we have $T \in \mathbf{T}_\mathcal{G} \Rightarrow W \in \mathbf{T}_\mathcal{G}$.

(① + ② implies $\mathbf{T}_\mathcal{M} \subseteq \mathbf{T}_\mathcal{G}$). Thus, we have shown that any node which can be shown to be in $\mathbf{T}_\mathcal{M}$ can also be shown to be in $\mathbf{T}_\mathcal{G}$, and therefore $\mathbf{T}_\mathcal{M} \subseteq \mathbf{T}_\mathcal{G}$.

It can be applied to show that $\mathsf{MUCT}(\mathcal{M}, Y, \mathbf{X}) \subseteq \mathsf{MUCT}(\mathcal{G}_{\overline{\mathbf{X}}}, Y)$, as we can operate over $\mathcal{M} \setminus \mathbf{X}$ and $\mathcal{G} \setminus \mathbf{X}$ instead of $\mathcal{M}$ and $\mathcal{G}$, respectively. Thus, we have that $T \in \mathsf{MUCT}(\mathcal{M}, Y, \mathbf{X})$ implies $T \in \mathsf{MUCT}(\mathcal{G}_{\overline{\mathbf{X}}}, Y)$. Therefore, we can conclude that $\mathsf{IB}(\mathcal{G}_{\overline{\mathbf{X}}}, Y) = \mathbf{X}$ holds. $\square$

Proposition 7. *Let $\mathcal{P}$ be a PAG over the set of variables $\mathbf{V}$. A set $\mathbf{X} \subseteq \mathbf{V} \setminus \{Y\}$ is a POMIS relative to $[\![\mathcal{P}, Y]\!]$ if and only if there exists a MAG $\mathcal{M}$ conforming to $\mathcal{P}$ such that $\mathbf{X}$ is a POMIS relative to $[\![\mathcal{M}, Y]\!]$.*

Proof. (**If**) Suppose $\mathbf{X}$ is a POMIS relative to $[\![\mathcal{M}, Y]\!]$ for some $\mathcal{M}$ conforming to $\mathcal{P}$. Then there exists an SCM $\mathcal{S}$ conforming to $\mathcal{M}$ such that $\mu_{\mathbf{x}^*} > \forall_{\mathbf{W} \in \mathbb{D}_{\mathcal{M},Y} \setminus \{\mathbf{X}\}} \mu_{\mathbf{w}^*}$. Since any SCM conforming to $\mathcal{M}$ also conforms to $\mathcal{P}$, the SCM also conforms to $\mathcal{P}$, the SCM $\mathcal{S}$ also conforms to $\mathcal{P}$, and thus $\mathbf{X}$ is a POMIS relative to $[\![\mathcal{P}, Y]\!]$.

(**Only if**) Let $\mathbf{X}$ be a POMIS relative to $[\![\mathcal{P}, Y]\!]$. Then there exists an SCM $\mathcal{S}$ conforming to $\mathcal{P}$ such that $\mu_{\mathbf{x}^*} > \forall_{\mathbf{W} \in \mathbb{D}_{\mathcal{P},Y} \setminus \{\mathbf{X}\}} \mu_{\mathbf{w}^*}$. Let $\mathcal{G}$ be the causal diagram associated with the SCM $\mathcal{S}$. Then, there exists a MAG $\mathcal{M}$ representing $\mathcal{G}$ that corresponds to $\mathcal{P}$ with $\mathbf{X}$ as a POMIS relative to $[\![\mathcal{M}, Y]\!]$, since $\mathbb{P}_{\mathcal{P},Y} \subseteq \mathbb{D}_{\mathcal{P},Y}$. This concludes the proof for this direction. $\square$

Proposition 8. *Let $\mathcal{Q}_\mathbf{X}$ be a PMG representing MAGs where $\mathbf{X}$ is a MIS with respect to Y. Then, the following properties hold in $\mathcal{Q}_\mathbf{X}$, for $X \in \mathbf{X}$:*

[11] For convenience, we denote MUCT and IB relative to $[\![\mathcal{G}, Y]\!]$ as $\mathsf{MUCT}(\mathcal{G}, Y)$ and $\mathsf{IB}(\mathcal{G}, Y)$, respectively.

1. *Every uncovered proper possibly-directed path from X to Y relative to $\mathbf{X}$ ends with an arrowhead (>).*
2. *If X is adjacent to Y, then the edge between X and Y is a directed edge.*

Proof. **(First condition).** For the sake of contradiction, suppose that there exists an uncovered path ending with a tail mark at Y in a MAG $\mathcal{M} \in [\mathcal{Q}_{\mathbf{X}}]$. This implies the path must take the form $X \leftarrow \cdots \leftarrow Y$ in $\mathcal{M}$. Since $\mathbf{X}$ is a MIS relative to $[\![\mathcal{M}, Y]\!]$, there exists a directed path from X to Y in $\mathcal{M}$, which would introduce a directed cycle, leading to a contradiction.

(Second condition). We will first show $X \ast\!\!-\!\!\ast Y$ forms $X \ast\!\!\rightarrow Y$ in $\mathcal{Q}_{\mathbf{X}}$, and then demonstrate that it must be $X \rightarrow Y$ by proving that $X \leftrightarrow Y$ leads to a contradiction. For the sake of contradiction, assume that there exists $X \leftarrow Y$ in a MAG $\mathcal{M} \in [\mathcal{Q}_{\mathbf{X}}]$. In $\mathcal{M}$, any directed path from X to Y would violate the ancestral property, resulting in a contradiction. Similarly, assume that there exists $X \leftrightarrow Y$ in a MAG $\mathcal{M} \in [\mathcal{Q}_{\mathbf{X}}]$. This configuration would also violate the ancestral property by introducing an almost directed cycle, which leads to a contradiction. □

Proposition 9. *Let $\mathcal{Q}^*_{\mathbf{X}}$ be a PMG which satisfies the conditions in Prop. 8 and the orientation completeness. For every MAG $\mathcal{M} \in [\mathcal{Q}^*_{\mathbf{X}}]$, if $\mathbf{X}$ is a MIS relative to $[\![\mathcal{M}, Y]\!]$, then there exists a PMG $\mathcal{Q}^i_{\mathbf{X}}$ representing $\mathcal{M}$ such that the following conditions are satisfied:*

1. *Every circle mark around nodes $\mathbf{X} \cup \{Y\}$ in $\mathcal{Q}_{\mathbf{X}}$ is oriented as either a tail ($-$) or an arrowhead (>) in $\mathcal{Q}^i_{\mathbf{X}}$ according to valid local transformations.*
2. *Every $X \in \mathbf{X}$ is an ancestor of Y in $\mathcal{Q}^i_{\mathbf{X}}$.*
3. *$\mathcal{Q}^i_{\mathbf{X}}$ is closed under orientation rules.*

Proof. The first and third conditions are satisfied by the soundness and completeness of valid local transformations (Thm. 8). Furthermore, since $\mathbf{X}$ is a DMIS with respect to $[\![\mathcal{P}, Y]\!]$, the second condition is also satisfied. □

Theorem 5 (soundness and completeness). *$\mathsf{IsPOMIS}$ returns True if and only if there exists a causal diagram $\mathcal{G}$ conforming to $\mathcal{P}$ such that $\mathbf{X}$ is a POMIS relative to $[\![\mathcal{G}, Y]\!]$.*

Proof. **($\mathsf{IsPOMIS}$ returns True $\Rightarrow \exists\mathcal{G}$ such that $\mathsf{IB}(\mathcal{G}_{\overline{\mathbf{X}}}, Y) = \mathbf{X}$).** Suppose that $\mathsf{IsPOMIS}$ returns True. Then, there is a PMG $\mathcal{Q}^i_{\mathbf{X}}$ satisfying $\mathsf{IB}(\mathcal{Q}^i_{\mathbf{X}}, Y, \mathbf{X}) = \mathbf{X}$. We will demonstrate that there exists a MAG $\mathcal{M} \in [\mathcal{Q}^i_{\mathbf{X}}]$ such that $\mathsf{IB}(\mathcal{M}, Y, \mathbf{X}) = \mathbf{X}$ by constructing such a MAG. To do so, consider the following lemma:

Lemma 35. *Let $\mathcal{Q}^i_{\mathbf{X}}$ be a PMG in Alg. 1. Let $\mathcal{M}$ be the graph resulting from the following procedure applied to $\mathcal{Q}^i_{\mathbf{X}}$.*

Step 1. *Orient partial directed edges ($\circ\!\!\rightarrow$) as directed edges ($\rightarrow$).*

Step 2. *Consider $A \ast\!\!\rightarrow B$ in $\mathcal{Q}^i_{\mathbf{X}}$. Let $\mathbf{T}^{\mathbf{X}}_{\mathcal{M}} \triangleq \mathsf{MUCT}(\mathcal{M}, Y, \mathbf{X})$. If B is contained in $\mathbf{T}^{\mathbf{X}}_{\mathcal{M}}$, orient the circle component including A as a DAG where each circle edge involving A in $\mathcal{Q}^i_{\mathbf{X}}$ corresponds to a directed edge outgoing from A in $\mathcal{M}$ (i.e., $A \circ\!\!-\!\!\circ V$ corresponds to $A \rightarrow V$).*

This step is repeated with the updated $\mathbf{T}^{\mathbf{X}}_{\mathcal{M}} \leftarrow \mathsf{MUCT}(\mathcal{M}, Y, \mathbf{X})$ as long as $\mathcal{M}$ remains unchanged.

Step 3. *Orient remaining circle component into a DAG with no unshielded colliders.*

Then, the resulting graph $\mathcal{M}$ is a MAG conforming to $\mathcal{Q}^i_{\mathbf{X}}$.

Proof. The construction follows Lemmas 6 and 12, and the fact that every circle component can be oriented independently by Lem. 31. □

Now, we will show that the MAG $\mathcal{M}$ constructed according to Lem. 35 satisfies $\mathsf{IB}(\mathcal{M}, Y, \mathbf{X}) = \mathbf{X}$. Let X be any node in $\mathsf{IB}(\mathcal{Q}^i_{\mathbf{X}}, Y, \mathbf{X})$. Then, X is a parent of some $T_X \in \mathsf{MUCT}(\mathcal{Q}^i_{\mathbf{X}}, Y, \mathbf{X})$ in $\mathcal{Q}^i_{\mathbf{X}}$. By Lem. 24, there exists an uncovered possibly-directed path $T_X \circ\!\!-\!\!\circ \cdots \circ\!\!-\!\!\circ T^*_X \ ?\!\!\rightarrow \cdot \rightarrow \cdots \rightarrow Y$. Due to the balanced property in Lem. 31 a path $X \rightarrow T^*_X \ ?\!\!\rightarrow \cdot \rightarrow \cdots \rightarrow Y$ exists in $\mathcal{Q}^i_{\mathbf{X}}$, which corresponds to $X \rightarrow T^*_X \rightarrow \cdots \rightarrow Y$ in $\mathcal{M}$ by construction (see Step 1). Therefore, we have that for any nodes $X \in \mathbf{X}$, X and T^*_X are included in $\mathtt{An}(Y)_{\mathcal{M}}$. Our goal is to show that $T^*_X \in \mathsf{MUCT}(\mathcal{M}, Y, \mathbf{X})$ since this means $X \in \mathsf{IB}(\mathcal{M}, Y, \mathbf{X})$.

For convenience, we denote $\mathbf{T}_{\mathcal{M}} = \mathsf{MUCT}(\mathcal{M}, Y, \mathbf{X})$ and $\mathbf{T}_{\mathcal{Q}^i_{\mathbf{X}}} = \mathsf{MUCT}(\mathcal{Q}^i_{\mathbf{X}}, Y, \mathbf{X})$. Let $\mathcal{N} \triangleq \mathcal{M}[\mathtt{An}(Y)_{\mathcal{M}}]$ and $\mathcal{H} \triangleq \mathcal{Q}^i_{\mathbf{X}}[\mathtt{PossAn}(Y)_{\mathcal{Q}^i_{\mathbf{X}}}]$. Suppose that T is a node such that $T \in \mathbf{T}_{\mathcal{Q}^i_{\mathbf{X}}} \cap \mathtt{An}(Y)_{\mathcal{M}}$ and $T \in \mathbf{T}_{\mathcal{M}}$. We know such a node exists, as Y is in both $\mathbf{T}_{\mathcal{M}}$ and $\mathbf{T}_{\mathcal{Q}^i_{\mathbf{X}}} \cap \mathtt{An}(Y)_{\mathcal{M}}$.

(① If $W \in \mathsf{PC}(T)_{\mathcal{H}[\mathtt{An}(Y)_{\mathcal{M}}]}$, then $W \in \mathbf{T}_{\mathcal{M}}$). Suppose that another node W is in the same pc-component of T in $\mathcal{H}[\mathtt{An}(Y)_{\mathcal{M}}]$. This implies that there exists a path between T and W such that (i) all non-endpoint nodes along the path are colliders, and (ii) none of the edges are visible, i.e., $T \ast\!\!\rightarrow \cdot \leftrightarrow \cdots \leftrightarrow \cdot \leftarrow\!\!\ast W$ in $\mathcal{H}[\mathtt{An}(Y)_{\mathcal{M}}]$.

> For all edges $U \overset{?}{\rightarrow} V$ along this path, the edges correspond to directed edges $U \rightarrow V$ in $\mathcal{N}$. If there are no circle edges with U in $\mathcal{H}$, the edges remain invisible in $\mathcal{N}$ since orienting a tail mark alone does not introduce any visible edges.
>
> Otherwise, if there are any circle edges $U \circ\!\!-\!\!\circ Z$ in $\mathcal{H}$ that correspond to $U \rightarrow Z$ in $\mathcal{N}$, no additional visible edges are introduced. When the edges correspond to $U \leftarrow Z$ in $\mathcal{N}$, U would already have been included in $\mathbf{T}_{\mathcal{M}}$, which in turn ensures that V be included in $\mathbf{T}_{\mathcal{M}}$.

(② If $W \in \mathtt{PossDe}(T)_{\mathcal{H}[\mathtt{An}(Y)_{\mathcal{M}}]}$, then $W \in \mathbf{T}_{\mathcal{M}}$). This means that there exists an uncovered possibly-directed path from T to W in $\mathcal{H}[\mathtt{An}(Y)_{\mathcal{M}}]$ by Lem. 28. According to our construction, there is a node $S \in \mathbf{T}_{\mathcal{M}}$ (it could be T) in the same bucket as T and W such that all nodes in the bucket are descendants of S in $\mathcal{M}$. Since $W \in \mathtt{De}(S)_{\mathcal{M}}$ and $S \in \mathbf{T}_{\mathcal{M}}$, we have $W \in \mathbf{T}_{\mathcal{M}}$.

(① + ②). Thus, we have shown that any node in $\mathbf{T}_{\mathcal{Q}^i_{\mathbf{X}}} \cap \mathtt{An}(Y)_{\mathcal{M}}$ can also be shown to be in $\mathbf{T}_{\mathcal{M}}$, and therefore we can get $T^*_X \in \mathbf{T}_{\mathcal{M}}$.

The remaining task is to prove that $\mathbf{W} \triangleq \mathsf{IB}(\mathcal{M}, Y, \mathbf{X}) \setminus \mathsf{IB}(\mathcal{Q}^i_{\mathbf{X}}, Y, \mathbf{X})$ is empty. For the sake of contradiction, consider any vertex $W \in \mathbf{W}$. Then, there exists a node $T_W \in \mathbf{T}_{\mathcal{M}}$ where $W \in \mathtt{Pa}(T_W)_{\mathcal{M}}$. Note that $T_W \in \mathbf{T}_{\mathcal{Q}^i_{\mathbf{X}}} \cap \mathtt{An}(Y)_{\mathcal{M}}$ holds (see the proof of the reverse direction). If $W \rightarrow T_W$ is invisible, then W is included in $\mathbf{T}_{\mathcal{M}}$, leading to a contradiction for $W \in \mathsf{IB}(\mathcal{M}, Y, \mathbf{X})$. If $W \rightarrow T_W$ is visible in both $\mathcal{M}$ and $\mathcal{Q}^i_{\mathbf{X}}$, then we can find a visible edge $W \rightarrow T^*_W$ satisfying $W \rightarrow T^*_W \overset{?}{\rightarrow} \cdots \cdot \rightarrow Y$ in $\mathcal{Q}^i_{\mathbf{X}}$ corresponding to $W \rightarrow T^*_W \rightarrow \cdots \rightarrow Y$ in $\mathcal{M}$ by Lemmas 24 and 27. This implies $W \in \mathsf{IB}(\mathcal{Q}^i_{\mathbf{X}}, Y, \mathbf{X})$, resulting in a contradiction. If $W \rightarrow T_W$ appeared as an invisible edge, either $\circ\!\!-\!\!\circ$ or $\circ\!\!\rightarrow$, $W \rightarrow T^*_W$ should also appear as an invisible edge by our construction (see Step 2). Therefore, we conclude the proof of the soundness of IsPOMIS.

(IsPOMIS returns False $\Rightarrow \nexists\mathcal{G}$ **such that** $\mathsf{IB}(\mathcal{G}_{\overline{\mathbf{X}}}, Y) = \mathbf{X}$**).** Let $\mathcal{G}$ be a causal diagram consistent with $\mathcal{M}$. Suppose that $\mathbf{X}$ is a POMIS relative to $[\![\mathcal{G}, Y]\!]$. Then, we have $\mathbf{X} = \mathsf{IB}(\mathcal{M}, Y, \mathbf{X})$. Let $\mathcal{Q}^i_{\mathbf{X}}$ be a PMG representing $\mathcal{M}$. Moreover, we have that $\mathtt{An}(Y)_{\mathcal{M}} \subseteq \mathtt{PossAn}(Y)_{\mathcal{Q}^i_{\mathbf{X}}}$ holds by Lem. 28.

Let X be any variable in $\mathsf{IB}(\mathcal{M}, Y, \mathbf{X})$. Then, X is a parent of some $T_X \in \mathsf{MUCT}(\mathcal{M}, Y, \mathbf{X})$ in $\mathcal{M}$. Furthermore, this appears in $\mathcal{Q}^i_{\mathbf{X}}$ by the construction of IsPOMIS in Alg. 1 (outgoing edges from X are determined in $\mathcal{Q}^i_{\mathbf{X}}$). By Lem. 24, there exists an uncovered possibly-directed path $T_X \circ\!\!-\!\!\circ \cdots T'_X \overset{?}{\rightarrow} \cdot \rightarrow \cdots \rightarrow Y$ in $\mathcal{Q}^i_{\mathbf{X}}$. Due to Lemmas 20 and 31, the path $X \rightarrow T^*_X \rightarrow \cdot \rightarrow \cdots \rightarrow Y$ exists in $\mathcal{Q}^i_{\mathbf{X}}$. Now we will show that $T^*_X \in \mathsf{MUCT}(\mathcal{Q}^i_{\mathbf{X}}, Y, \mathbf{X})$ since this implies $X \in \mathsf{IB}(\mathcal{Q}^i_{\mathbf{X}}, Y, \mathbf{X})$.

Let $\mathbf{T}_{\mathcal{M}} \triangleq \mathsf{MUCT}(\mathcal{M}, Y, \mathbf{X})$ and $\mathbf{T}_{\mathcal{Q}^i_{\mathbf{X}}} \triangleq \mathsf{MUCT}(\mathcal{Q}^i_{\mathbf{X}}, Y, \mathbf{X})$. Let $\mathcal{N} \triangleq \mathcal{M}[\mathtt{An}(Y)_{\mathcal{M}}]$ and $\mathcal{H} \triangleq \mathcal{Q}^i_{\mathbf{X}}[\mathtt{PossAn}(Y)_{\mathcal{Q}^i_{\mathbf{X}}}]$. Suppose that T is a node satisfying $T \in \mathbf{T}_{\mathcal{Q}^i_{\mathbf{X}}} \cap \mathtt{An}(Y)_{\mathcal{M}}$ and $T \in \mathbf{T}_{\mathcal{M}}$. We know such a node exists since Y is in both $\mathbf{T}_{\mathcal{M}}$ and $\mathbf{T}_{\mathcal{Q}^i_{\mathbf{X}}} \cap \mathtt{An}(Y)_{\mathcal{M}}$.

(If $W \in \mathbf{T}_{\mathcal{M}}$, then $W \in \mathbf{T}_{\mathcal{Q}^i_{\mathbf{X}}} \cap \mathtt{An}(Y)_{\mathcal{M}}$). Since any invisible edges in $\mathcal{M}$ correspond to invisible ones in $\mathcal{Q}^i_{\mathbf{X}}$, we have $W \in \mathsf{PC}_{\mathcal{N}}(T)$ implies $W \in \mathbf{T}_{\mathcal{Q}^i_{\mathbf{X}}} \cap \mathtt{An}(Y)_{\mathcal{M}}$ according to Lem. 30. Furthermore, we know that $W \in \mathtt{De}(T)_{\mathcal{N}}$ implies $W \in \mathtt{PossDe}(T)_{\mathcal{H}[\mathtt{An}(Y)_{\mathcal{M}}]}$ by Lem. 28. Therefore, we get that $W \in \mathbf{T}_{\mathcal{Q}^i_{\mathbf{X}}} \cap \mathtt{An}(Y)_{\mathcal{M}}$. Thus, we have shown that any node in $\mathbf{T}_{\mathcal{M}}$ can also be shown to be in $\mathbf{T}_{\mathcal{Q}^i_{\mathbf{X}}} \cap \mathtt{An}(Y)_{\mathcal{M}}$, and therefore $T^*_X \in \mathbf{T}_{\mathcal{Q}^i_{\mathbf{X}}}$.

The remaining task is to prove that $\mathbf{W} \triangleq \mathsf{IB}(\mathcal{Q}^i_{\mathbf{X}}, Y, \mathbf{X}) \setminus \mathsf{IB}(\mathcal{M}, Y, \mathbf{X})$ is empty. For the sake of contradiction, consider any vertex $W \in \mathbf{W}$. Then, there exists a node $T_W \in \mathbf{T}_{\mathcal{Q}^i_{\mathbf{X}}}$ where $W \in \mathtt{Pa}(T_W)_{\mathcal{Q}^i_{\mathbf{X}}}$. If $W \rightarrow T_W$ is invisible in $\mathcal{Q}^i_{\mathbf{X}}$, then W is included in $\mathbf{T}_{\mathcal{Q}^i_{\mathbf{X}}}$, leading to a contradiction for $W \in \mathsf{IB}(\mathcal{Q}^i_{\mathbf{X}}, Y, \mathbf{X})$. If $W \rightarrow T_W$ is visible in $\mathcal{Q}^i_{\mathbf{X}}$, it is also visible in $\mathcal{M}$, and we can find a visible edge $W \rightarrow T^*_W$ satisfying $W \rightarrow T^*_W \rightarrow \cdots \cdot \rightarrow Y$ by Lemmas 24 and 27. This implies $W \in \mathsf{IB}(\mathcal{M}, Y, \mathbf{X})$, resulting in a contradiction. Therefore, we conclude the proof of the completeness of IsPOMIS. $\square$