On Causal Identification under Markov Equivalence*

Amin Jaber[1] , **Jiji Zhang**[2] and **Elias Bareinboim**[1]
[1]Department of Computer Science, Purdue University
[2]Department of Philosophy, Lingnan University
{jaber0, eb}@purdue.edu, jijizhang@ln.edu.hk

Abstract

In this work, we investigate the problem of computing an experimental distribution from a combination of the observational distribution and a partial qualitative description of the causal structure of the domain under investigation. This description is given by a partial ancestral graph (PAG) that represents a Markov equivalence class of causal diagrams, i.e., diagrams that entail the same conditional independence model over observed variables, and is learnable from the observational data. Accordingly, we develop a complete algorithm to compute the causal effect of an arbitrary set of intervention variables on an arbitrary outcome set.

1 Introduction

A prominent approach to infer causal effects leverages a combination of substantive knowledge about the domain under investigation, usually encoded in the form of a causal diagram, with observational data [Pearl, 2000; Bareinboim and Pearl, 2016]. A sample diagram is shown in Fig. 1a, where the nodes represent variables, directed edges represent direct causal relations from tails to heads, and bi-directed arcs represent the presence of unobserved (latent) variables that generate spurious association between the variables. The decision problem of whether an interventional distribution can be computed from a combination of observational and experimental data together with the causal diagram is known as the problem of identification of causal effects (identification, for short). For instance, a possible task is to identify the effect of $do(X{=}x)$ on $V_4{=}v_4$, i.e. $P_x(v_4)$, given the diagram in Fig. 1a and data from the distribution $P(X, V_1, ..., V_4)$.

This problem has been extensively studied in the literature with a number of criteria established [Pearl, 1993; Galles and Pearl, 1995; Kuroki and Miyakawa, 1999; Tian and Pearl, 2002; Huang and Valtorta, 2006; Shpitser and Pearl, 2006; Bareinboim and Pearl, 2012], which include the back-door criterion and the do-calculus [Pearl, 1995]. Despite their power, these techniques require a fully specified causal diagram, which is not always available in practical settings.

*This is an abridged version of [Jaber *et al.*, 2018a] which won the best student paper award at UAI-18 and [Jaber *et al.*, 2019].

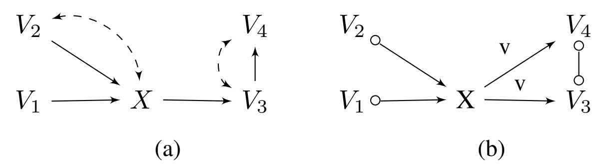

Figure 1: A causal diagram (left) and the inferred PAG (right).

One may try to learn the causal diagram from data. However, it is common that only an equivalence class of causal diagrams can be consistently inferred from observational data [Verma, 1993; Spirtes *et al.*, 2001]. A useful representation of such an equivalence class is a *partial ancestral graph (PAG)* [Zhang, 2008b]. Fig. 1b shows the PAG that can be inferred from observational data generated by Fig. 1a. In PAGs, directed edges signify (possibly indirect) causal relations and circle marks indicate structural uncertainty.

In this work, we analyze the marriage of these two lines of investigation. Identification from an equivalence class is considerably more challenging than from a single diagram due to the structural uncertainty regarding both the causal relations among the variables and the presence of hidden variables. Zhang [2007] extended the do-calculus to PAGs. In practice, however, it is computationally hard to decide whether there exists a sequence of applications of the rules of the generalized calculus to identify an effect. Perković *et al.* [2015] generalized the back-door criterion to PAGs, and provided a complete algorithm to find a back-door admissible set, should such a set exist. However, in practice, no adjustment set exists for many identifiable effects. In this extended abstract, we summarize the approach taken in [Jaber *et al.*, 2018b; Jaber *et al.*, 2018b; Jaber *et al.*, 2019] that culminated in a complete algorithm to identify causal effects given PAGs. Specifically, our contributions are as follows:

1. We revisit the original identification algorithm given causal diagrams and introduce a new formulation which is more amenable under structural uncertainties.
2. We derive crucial PAG properties including a novel graph-decomposition strategy that breaks a target causal distribution into an equal product of more tractable ones.
3. We develop a complete algorithm to compute the effect of an arbitrary set of intervention variables on an arbitrary outcome set from a PAG and observational data.

2 Preliminaries

In this section, we introduce the basic notation and machinery used throughout the paper. Bold capital letters denote sets of variables, while bold lowercase letters stand for particular value assignments to those variables.

2.1 Structural Causal Models

We use the language of Structural Causal Models (SCM) [Pearl, 2000, pp. 204-207] as our basic semantic framework. Formally, an SCM M is a 4-tuple $\langle \mathbf{U}, \mathbf{V}, \mathbf{F}, P(\mathbf{U}) \rangle$, where $\mathbf{U}$ is a set of exogenous (latent) variables and $\mathbf{V}$ is a set of endogenous (measured) variables. $\mathbf{F}$ represents a collection of functions $\{f_i\}$ such that each endogenous variable $V_i \in \mathbf{V}$ is determined by a function $f_i \in \mathbf{F}$, where f_i is a mapping from the respective domain of $U_i \cup Pa_i$ to V_i, $U_i \subseteq \mathbf{U}$, $Pa_i \subseteq \mathbf{V} \setminus \{V_i\}$. The uncertainty is encoded through a probability distribution over the exogenous variables, $P(\mathbf{U})$, and the marginal distribution induced over the endogenous variables $P(\mathbf{V})$ is called observational. Every SCM is associated with one causal diagram where every variable $V_i \in \mathbf{V}$ is a node, and there exists a directed edge from every node in Pa_i to V_i. Also, for every pair $V_i, V_j \in \mathbf{V}$ such that $U_i \cap U_j \neq \emptyset$, there exists a bi-directed edge between V_i and V_j. We restrict our study to recursive systems, which means that the corresponding diagram will be acyclic.

Within the structural semantics, performing an action $X = x$ is represented through the do-operator, $do(X = x)$, which encodes the operation of replacing the original equation for X by the constant x and induces a submodel M_x. The resulting distribution is denoted by P_x, which is the main target for identification in this paper. For further details on structural models, we refer readers to [Pearl, 2000].

2.2 Ancestral Graphs

We now introduce a graphical representation of equivalence classes of causal diagrams. A *mixed* graph can contain directed and bi-directed edges. A is an ancestor of B if they share a directed path out of A. A is a *spouse* of B if $A \leftrightarrow B$ is present. An *almost directed cycle* happens when A is both a spouse and an ancestor of B. An *inducing path* is a path on which every node (except for the endpoints) is a collider on the path (i.e., both edges incident to X are into X) and every collider is an ancestor of an endpoint of the path. A mixed graph is *ancestral* if it doesn't contain a directed or almost directed cycle. It is *maximal* if there is no inducing path between any two non-adjacent nodes. A *Maximal Ancestral Graph* (MAG) is a graph that is both ancestral and maximal [Richardson and Spirtes, 2002].

In short, a MAG represents a set of causal diagrams that entail the same independence and ancestral relations among the observed variables. Different MAGs may be Markov equivalent in that they entail the exact same independence model. A partial ancestral graph (PAG) represents an equivalence class of MAGs $[\mathcal{M}]$, which shares the same adjacencies as every MAG in $[\mathcal{M}]$ and displays all and only the invariant edge marks. A circle indicates an uncommon edge mark. A PAG is learnable from the conditional independence relations among the observed variables and the FCI algorithm is a standard method to learn such an object [Zhang, 2008b].

2.3 Graphical Notions

Given a causal diagram, a MAG, or a PAG, a path between X and Y is *potentially directed (causal)* from X to Y if there is no arrowhead on the path pointing towards X. X is a *possible ancestor* of Y, i.e., $X \in \texttt{An}(Y)$, if there is a potentially directed path from X to Y. By stipulation, $X \in \texttt{An}(X)$. Y is called a *possible child* of X, i.e. $Y \in \texttt{Ch}(X)$, if they are adjacent and the edge is not into X. For a set of nodes $\mathbf{X}$, we have $\texttt{Ch}(\mathbf{X}) = \cup_{X \in \mathbf{X}} \texttt{Ch}(X)$. If the edge marks on a path between X and Y are all circles, we call the path a *circle path*. We refer to the closure of nodes connected with circle paths as a *bucket*. Obviously, given a PAG, nodes are partitioned into a unique set of buckets.

A directed edge $X \rightarrow Y$ in a MAG or PAG 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. This implies that a visible edge is not confounded ($X \leftarrow - - \rightarrow Y$ doesn't exist). Which directed edges are visible is easily decidable by a graphical condition [Zhang, 2008a, Def. 8], so we simply mark visible edges by v. For brevity, we refer to any edge that is not a visible directed edge as *invisible*.

3 Revisit Identification in Causal Diagrams

Tian and Pearl [2002] presented an identification algorithm based on a decomposition strategy of the causal diagram into a set of so-called *c-components* (confounded components).

Definition 1 (C-Component)**.** *In a causal diagram, two nodes are said to be in the same c-component if and only if they are connected by a bi-directed path, i.e., a path composed solely of bi-directed edges.*

For any set $\mathbf{C} \subseteq \mathbf{V}$, the quantity $Q[\mathbf{C}]$ is defined to denote the post-intervention distribution of $\mathbf{C}$ under an intervention on $\mathbf{V} \setminus \mathbf{C}$, i.e. $P_{\mathbf{v} \setminus \mathbf{c}}(\mathbf{c})$. Given a diagram $\mathcal{D}$, $Q[\mathbf{C}]$ decomposes into a product of sub-queries over the c-components in $\mathcal{D}_{\mathbf{C}}$, where $\mathcal{D}_{\mathbf{C}}$ denotes the (induced) subgraph of $\mathcal{D}$ over $\mathbf{C}$. Hence, we get the following decomposition, where $\mathbf{C}_i$ is a c-component in $\mathcal{D}_{\mathbf{C}}$:

$$Q[\mathbf{C}] = \prod_i Q[\mathbf{C}_i] \tag{1}$$

The significance of c-components and their decomposition is evident from [Tian, 2002, Lemmas 10, 11], which are the basis of Tian's identification algorithm. Our goal is to reformulate the procedure with a more local, atomic criterion shown in Lem. 1. The step-wise algorithm is shown in Alg. 1.

Lemma 1. *Given a causal diagram $\mathcal{D}$ over $\mathbf{V}$, $X \in \mathbf{T} \subseteq \mathbf{V}$, and $P_{\mathbf{v} \setminus \mathbf{t}}$, i.e., an expression for $Q[\mathbf{T}]$. If X is not in the same c-component with a child in $\mathcal{D}_{\mathbf{T}}$, then $Q[\mathbf{T} \setminus \{X\}]$ is identifiable and given by*

$$Q[\mathbf{T} \setminus \{X\}] = \frac{P_{\mathbf{v} \setminus \mathbf{t}}}{Q[S^X]} \times \sum_x Q[S^X] \tag{2}$$

where S^X is the c-component of X in $\mathcal{D}_{\mathbf{T}}$ and $Q[S^X]$ is computable from $P_{\mathbf{v} \setminus \mathbf{t}}$ by [Tian, 2002, Lemma 11].

Algorithm 1 $\mathbf{ID}(\mathbf{x}, \mathbf{y})$ given Causal Diagram $\mathcal{G}$

Input: two disjoint sets $\mathbf{X}, \mathbf{Y} \subset \mathbf{V}$
Output: Expression for $P_{\mathbf{x}}(\mathbf{y})$ or `FAIL`

```
Let D = An(Y)_{G_{V\X}}
Let the c-components of G_D be D_i, i = 1, ..., k
P_x(y) = Σ_{d\y} Π_i IDENTIFY(D_i, V, P)
function IDENTIFY(C, T, Q = Q[T])
    if C = T then return Q[T]
    /* Let S^B denote the c-component of {B} in G_T */
    if ∃B ∈ T \ C such that S^B ∩ Ch(B) = ∅ then
        Compute Q[T \ {B}] from Q;        ▷ Lemma 1
        return IDENTIFY(C, T \ {B}, Q[T \ {B}])
    else
        throw FAIL
```

The revised algorithm requires checking an atomic criterion at every instance of the recursive routine IDENTIFY. This might not be crucial when the precise causal diagram is known and the induced subgraphs preserve complete information about the c-components and the ancestral relations between the nodes. However, it becomes significant when the domain description is an equivalence class represented by a PAG, in which structural information is partial.

4 PAG Properties and Q-Decomposition

Evidently, induced subgraphs of the original causal diagram play a critical role in identification (cf Alg. 1). It is natural to expect that in the generalized setting we study here, induced subgraphs of the given PAG will also play an important role. An immediate challenge, however, is that a subgraph of a PAG $\mathcal{P}$ over $\mathbf{V}$ induced by $\mathbf{A} \subseteq \mathbf{V}$ is, in general, not a PAG that represents a full Markov equivalence class. In particular, if $\mathcal{D}$ is a diagram in the equivalence class represented by $\mathcal{P}$, $\mathcal{P}_{\mathbf{A}}$ is in general not the PAG that represents the equivalence class of $\mathcal{D}_{\mathbf{A}}$. For example, let $\mathcal{D}$ and $\mathcal{P}$ denote the diagram and the corresponding PAG in Figure 1, respectively, and let $\mathbf{A} = \{V_1, V_2, X, V_4\}$. Then, V_4 is disconnected in $\mathcal{D}_{\mathbf{A}}$ while it is a child of X in $\mathcal{P}_{\mathbf{A}}$. Despite this subtlety, we establish a few facts showing that for any $\mathbf{A} \subseteq \mathbf{V}$ and any diagram $\mathcal{D}$ in the equivalence class represented by $\mathcal{P}$, some information about $\mathcal{D}_{\mathbf{A}}$, which is particularly relevant to identification, can still be read off from $\mathcal{P}_{\mathbf{A}}$. In [Jaber *et al.*, 2018a], we discuss a key property related to c-components.

Definition 2 (PC-Component)**.** *In a MAG, a PAG, or any induced subgraph thereof, two nodes are in the same possible c-component (pc-component) if there is a path between them such that (1) all non-endpoint nodes along the path are colliders, and (2) none of the edges is visible.*

As mentioned earlier, a c-component in a causal diagram plays a central role in identification. The following proposition establishes a graphical condition in an induced subgraph $\mathcal{P}_{\mathbf{A}}$ that is necessary for two nodes being in the same c-component in $\mathcal{D}_{\mathbf{A}}$ for some diagram $\mathcal{D}$ represented by $\mathcal{P}$.

Proposition 1. *Let $\mathcal{P}$ be a PAG over $\mathbf{V}$, and $\mathcal{D}$ be any causal diagram in the equivalence class represented by $\mathcal{P}$. For any $X, Y \in \mathbf{A} \subseteq \mathbf{V}$, if X and Y are in the same c-component in $\mathcal{D}_{\mathbf{A}}$, then X and Y are in the same pc-component in $\mathcal{P}_{\mathbf{A}}$.*

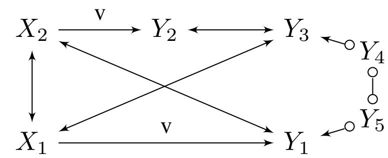

Figure 2: Sample PAG.

This result provides a sufficient condition for *not* belonging to the same c-component in any of the induced causal diagrams. In the aforementioned $\mathcal{P}_{\mathbf{A}}$ of Fig. 1b, for example, V_1 and V_4 or X and V_4 are not in the same pc-component, which implies by Prop. 1 that they are not in the same c-component in $\mathcal{D}_{\mathbf{A}}$ for any $\mathcal{D}$ in the equivalence class represented by the PAG in Fig. 1b. As a special case of Def. 2, we define the following notion, which will prove useful later on.

Definition 3 (DC-Component)**.** *In a MAG, a PAG, or any induced subgraph thereof, two nodes are in the same definite c-component (dc-component) if they are connected with a bi-directed path, i.e. a path composed of bi-directed edges.*

Next, we use the *pc-component* property to devise a decomposition for $Q[\mathbf{C}]$ in PAGs akin to that for causal diagrams presented in Eq. 1. However, the decomposition in PAGs is challenging due to the structural uncertainties; most relevant the presence of latent confounders. For instance, given the query $Q[\mathbf{Y}]$ over the PAG in Fig. 2, the sequence of nodes $\langle Y_2, Y_3, Y_4, Y_5, Y_1 \rangle$ is connected with invisible edges, which are possibly confounded. Hence, any naive decomposition of $Q[\mathbf{Y}]$ into a product of sub-queries over subsets of $\mathbf{Y}$ is invalid since we can construct a diagram in the equivalence class which violates this decomposition.

To develop a valid decomposition, we start by introducing the notion of a *region*. In short, a *region* is the pc-component of a set $\mathbf{A}$ appended with the corresponding buckets of the nodes. We append the pc-component set with the corresponding buckets of the nodes to avoid non-identifiability of the sub-queries since no sufficient causal information is present within a bucket. Consider the PAG in Figure 2, and let $\mathbf{C} = \mathbf{Y}$ and $\mathbf{A} = \{Y_3\}$. Then, $\mathcal{R}^{\mathbf{C}}_{\mathbf{A}} = \{Y_3, Y_2, Y_4, Y_5\}$ since Y_2 and Y_4 are in the pc-component of Y_3 and Y_5 is in the same bucket as Y_4. For simplicity, we often drop $\mathbf{C}$, i.e. $\mathcal{R}_{\mathbf{A}}$, whenever it is clear from the context. Using this construction, we derive the decomposition in Theorem 1.

Definition 4 (Region $\mathcal{R}^{\mathbf{C}}_{\mathbf{A}}$)**.** *Given a PAG or a MAG $\mathcal{G}$ over $\mathbf{V}$, and $\mathbf{A} \subseteq \mathbf{C} \subseteq \mathbf{V}$. Let the region of $\mathbf{A}$ with respect to $\mathbf{C}$, denoted $\mathcal{R}^{\mathbf{C}}_{\mathbf{A}}$, be the union of the buckets that contain nodes in the pc-component of $\mathbf{A}$ in the induced subgraph $\mathcal{G}_{\mathbf{C}}$.*

Theorem 1. *Given a PAG $\mathcal{P}$ over $\mathbf{V}$ and set $\mathbf{C} \subseteq \mathbf{V}$, $Q[\mathbf{C}]$ decomposes as follows, where $\mathbf{A} \subset \mathbf{C}$ and $\mathcal{R}_{(.)} = \mathcal{R}^{\mathbf{C}}_{(.)}$.*

$$Q[\mathbf{C}] = \frac{Q[\mathcal{R}_{\mathbf{A}}].Q[\mathcal{R}_{\mathbf{C} \setminus \mathcal{R}_{\mathbf{A}}}]}{Q[\mathcal{R}_{\mathbf{A}} \cap \mathcal{R}_{\mathbf{C} \setminus \mathcal{R}_{\mathbf{A}}}]}$$

Back to the query $Q[\mathbf{Y}]$ over the PAG in Fig. 2, it can be decomposed as follows with $\mathbf{A} = \{Y_3\}$.

$$Q[\mathbf{Y}] = \frac{Q[\mathbf{Y} \setminus \{Y_1\}].Q[\{Y_1, Y_4, Y_5\}]}{Q[\{Y_4, Y_5\}]} \tag{3}$$

5 Identification in PAGs

We start by formally defining the notion of identification given a PAG, which generalizes the model-specific notion [Pearl, 2000, pp. 70].

Definition 5. *Given a PAG $\mathcal{P}$ over $\mathbf{V}$ and a query $P_{\mathbf{x}}(\mathbf{y})$ where $\mathbf{X}, \mathbf{Y} \subset \mathbf{V}$, $P_{\mathbf{x}}(\mathbf{y})$ is identifiable given $\mathcal{P}$ iff $P_{\mathbf{x}}(\mathbf{y})$ is identifiable given every causal diagram $\mathcal{D}$ (represented by a MAG) in the Markov equivalence class represented by $\mathcal{P}$, and with the same expression.*

We first derive an atomic identification criterion akin to Lemma 1. Due to the absence of any causal information within a bucket in a PAG, it is for our purpose analogous to a single node in a causal diagram. Therefore, the following criterion targets a bucket $\mathbf{X}$ rather than a single node. Note that a partial topological order is an order over the buckets. A detailed discussion can be found in [Jaber *et al.*, 2018b; Jaber *et al.*, 2018a].

Theorem 2. *Let $\mathcal{P}$ denote a PAG over $\mathbf{V}$, $\mathbf{T}$ be the union of a subset of the buckets in $\mathcal{P}$, and $\mathbf{X} \subset \mathbf{T}$ be a bucket. Given $P_{\mathbf{v} \setminus \mathbf{t}}$, and a partial topological order $\mathbf{B}_1 < \cdots < \mathbf{B}_m$ with respect to $\mathcal{P}_{\mathbf{T}}$, $Q[\mathbf{T} \setminus \mathbf{X}]$ is identifiable if and only if, in $\mathcal{P}_{\mathbf{T}}$, there does not exist $Z \in \mathbf{X}$ such that Z has a possible child $C \notin \mathbf{X}$ that is in the pc-component of Z. If identifiable, then the expression is given by*

$$Q[\mathbf{T} \setminus \mathbf{X}] = \frac{P_{\mathbf{v} \setminus \mathbf{t}}}{\prod_{\{i | \mathbf{B_i} \subseteq S^{\mathbf{X}}\}} P_{\mathbf{v} \setminus \mathbf{t}}(\mathbf{B_i} | \mathbf{B^{(i-1)}})} \times \sum_{\mathbf{x}} \prod_{\{i | \mathbf{B_i} \subseteq S^{\mathbf{X}}\}} P_{\mathbf{v} \setminus \mathbf{t}}(\mathbf{B_i} | \mathbf{B^{(i-1)}}), \tag{4}$$

where $S^{\mathbf{X}} = \bigcup_{Z \in \mathbf{X}} S^Z$, S^Z being the dc-component of Z in $\mathcal{P}_{\mathbf{T}}$, and $\mathbf{B}^{(i-1)}$ denoting the set of nodes preceding bucket $\mathbf{B}_i$ in the partial order.

For example, consider the query $P_x(\mathbf{v} \setminus \{x\})$ over the PAG in Fig. 1b. The intervention node X is not in the same pc-component with any of its possible children (V_3 and V_4), hence the effect is identifiable and given by

$$\begin{aligned} P_x(\mathbf{v} \setminus \{x\}) &= \frac{P(\mathbf{v})}{P(x | v_1, v_2)} \times \sum_{x'} P(x' | v_1, v_2) \\ &= P(v_1, v_2) P(v_3, v_4 | v_1, v_2, x) \end{aligned}$$

Using the identification criterion in Thm. 2 and the decomposition in Thm. 1, we formulate the procedure we call **IDP**, which is shown in Alg. 2. The main idea of **IDP** goes as follows. After receiving the sets $\mathbf{X}$, $\mathbf{Y}$, and a PAG $\mathcal{P}$, the algorithm pre-processes the query by computing $\mathbf{D}$, the set of possible ancestors of $\mathbf{Y}$ in $\mathcal{P}_{\mathbf{V} \setminus \mathbf{X}}$. Then, the procedure calls the subroutine IDENTIFY over $\mathbf{D}$ to compute $Q[\mathbf{D}]$ from the observational distribution $P(\mathbf{V})$. The recursive routine basically tests for one of two conditions. First, it checks for the presence of a bucket $\mathbf{B}$ in $\mathcal{P}_{\mathbf{T}}$ that is a subset of the intervention nodes, i.e. $\mathbf{B} \subseteq \mathbf{T} \setminus \mathbf{C}$, and satisfies the condition of Thm. 2. If found, it computes $Q[\mathbf{T} \setminus \mathbf{B}]$ using Eq. 4, and proceeds with a recursive call. Alternatively, if such a bucket does not exist in $\mathcal{P}_{\mathbf{T}}$, then **IDP** checks for a bucket $\mathbf{B}$ in $\mathcal{P}_{\mathbf{C}}$ such that the region of $\mathbf{B}$ with respect to $\mathbf{C}$, i.e. $\mathcal{R}^{\mathbf{C}}_{\mathbf{B}}$, does not span $\mathbf{C}$. If such a bucket exists, then **IDP** decomposes the query $Q[\mathbf{C}]$ according to Thm. 1. Finally, if both tests fail, then **IDP** throws a failure condition. In [Jaber *et al.*, 2019], we prove that **IDP** is complete and we devise a graphical characterization of non-identifiable causal effects.

Algorithm 2 $\mathbf{IDP}(\mathbf{x}, \mathbf{y})$ given PAG $\mathcal{P}$

Input: two disjoint sets $\mathbf{X}, \mathbf{Y} \subset \mathbf{V}$
Output: Expression for $P_{\mathbf{x}}(\mathbf{y})$ or `FAIL`

1: Let $\mathbf{D} = \mathtt{An}(\mathbf{Y})_{\mathcal{P}_{\mathbf{V} \setminus \mathbf{X}}}$
2: $P_{\mathbf{x}}(\mathbf{y}) = \sum_{\mathbf{d} \setminus \mathbf{y}} \textsc{Identify}(\mathbf{D}, \mathbf{V}, P)$

3: **function** $\textsc{Identify}(\mathbf{C}, \mathbf{T}, Q = Q[\mathbf{T}])$
4: **if** $\mathbf{C} = \emptyset$ **then return** 1
5: **if** $\mathbf{C} = \mathbf{T}$ **then return** Q
$\triangleright$ *In $\mathcal{P}_{\mathbf{T}}$, let $\mathbf{B}$ denote a bucket, and let $C^{\mathbf{B}}$ denote the pc-component of* $\mathbf{B}$
6: **if** $\exists \mathbf{B} \subset \mathbf{T} \setminus \mathbf{C}$ such that $C^{\mathbf{B}} \cap \mathtt{Ch}(\mathbf{B}) \subseteq \mathbf{B}$ **then**
7: Compute $Q[\mathbf{T} \setminus \mathbf{B}]$ from Q (via Thm. 2)
8: **return** $\textsc{Identify}(\mathbf{C}, \mathbf{T} \setminus \mathbf{B}, Q[\mathbf{T} \setminus \mathbf{B}])$
9: **else if** $\exists \mathbf{B} \subset \mathbf{C}$ such that $\mathcal{R}_{\mathbf{B}} \neq \mathbf{C}$ **then**
10: **return** $\frac{\textsc{Identify}(\mathcal{R}_{\mathbf{B}}, \mathbf{T}, Q) \cdot \textsc{Identify}(\mathcal{R}_{\mathbf{C} \setminus \mathcal{R}_{\mathbf{B}}}, \mathbf{T}, Q)}{\textsc{Identify}(\mathcal{R}_{\mathbf{B}} \cap \mathcal{R}_{\mathbf{C} \setminus \mathcal{R}_{\mathbf{B}}}, \mathbf{T}, Q)}$
11: **else**
12: **throw** `FAIL`

5.1 Illustrative Example

Consider the effect $P_{\mathbf{x}}(\mathbf{y})$ given PAG $\mathcal{P}$ in Fig. 2, where $\mathbf{X} = \{X_1, X_2\}$ and $\mathbf{Y} = \{Y_1, Y_2, Y_3, Y_4, Y_5\}$. We have $\mathbf{D} = \mathbf{Y}$, and the query reduces to computing $Q[\mathbf{D}]$ with $\textsc{Identify}(\mathbf{D}, \mathbf{V}, P)$. Neither X_1 nor X_2 satisfies the condition at line (6) of Alg. 2, hence **IDP** decomposes $Q[\mathbf{Y}]$, and we have the expression derived earlier in Eq. 3.

First, with the call $\textsc{Identify}(\{Y_2, Y_3, Y_4, Y_5\}, \mathbf{V}, P))$, node Y_1 satisfies the condition at line 6 as it has no children. Hence, we compute $Q[\mathbf{V} \setminus \{Y_1\}]$ from $P(\mathbf{V})$ using Thm. 2. Subsequent calls intervene on X_1 then X_2 and we obtain the simplified expression.

$$Q[\{Y_2, Y_3, Y_4, Y_5\}] = P(y_2, y_3, y_4, y_5 | x_2) \tag{5}$$

Similarly, we get the following expressions for $Q[\{Y_1, Y_4, Y_5\}]$ and $Q[\{Y_4, Y_5\}]$, respectively.

$$Q[\{Y_1, Y_4, Y_5\}] = P(y_1, y_4, y_5 | x_1) \tag{6}$$

and,

$$Q[\{Y_4, Y_5\}] = P(y_4, y_5) \tag{7}$$

Hence, the final expression for $P_{\mathbf{x}}(\mathbf{y})$ is a result of Eqs. 5, 6, and 7 as follows.

$$P_{\mathbf{x}}(\mathbf{y}) = \frac{P(y_2, y_3, y_4, y_5 | x_2) \;.\; P(y_1, y_4, y_5 | x_1)}{P(y_4, y_5)} \tag{8}$$

We can reformulate Eq. 8 using independence relations to,

$$P_{\mathbf{x}}(\mathbf{y}) = P(y_2, y_3|x_2, y_4) \cdot P(y_1|x_1, y_5) \cdot P(y_4, y_5)$$

6 Conclusion

In this work, we investigated the problem of identification of interventional distributions in Markov equivalence classes represented by PAGs. We first revisited the identification algorithm given a specific causal diagram and reformulated it using an atomic criterion. We then derived new graphical properties for induced subgraphs of PAGs over an arbitrary subset of nodes, including a novel decomposition strategy. Finally, building on these results, we proposed an identification criterion in PAGs and employed it in a complete identification algorithm. We believe the proposed approach enables empirical researchers to perform causal reasoning while lacking substantive background knowledge. This is considered more "data-driven", and more aligned with the zeitgeist in machine learning today.

Acknowledgements

Bareinboim and Jaber are supported in parts by grants from NSF IIS-1704352, IIS1750807 (CAREER), IBM Research, and Adobe Research. Zhang's research was supported in part by the Research Grants Council of Hong Kong under the General Research Fund LU13602818.

References

[Bareinboim and Pearl, 2012] Elias Bareinboim and Judea Pearl. Causal inference by surrogate experiments: z-identifiability. In *Proceedings of the Twenty-Eighth Conference on Uncertainty in Artificial Intelligence*, UAI'12, pages 113–120, August 2012.

[Bareinboim and Pearl, 2016] Elias Bareinboim and Judea Pearl. Causal inference and the data-fusion problem. *Proceedings of the National Academy of Sciences*, 113:7345–7352, 2016.

[Galles and Pearl, 1995] David Galles and Judea Pearl. Testing identifiability of causal effects. In *Proceedings of the Eleventh Conference on Uncertainty in Artificial Intelligence*, UAI'95, pages 185–195, August 1995.

[Huang and Valtorta, 2006] Yimin Huang and Marco Valtorta. Pearl's calculus of intervention is complete. In *Proceedings of the Twenty-Second Conference on Uncertainty in Artificial Intelligence*, UAI'06, pages 217–224, July 2006.

[Jaber *et al.*, 2018a] Amin Jaber, Jiji Zhang, and Elias Bareinboim. Causal identification under Markov equivalence. In *Proceedings of the 34th Conference on Uncertainty in Artificial Intelligence*, UAI'18, pages 978–987, August 2018.

[Jaber *et al.*, 2018b] Amin Jaber, Jiji Zhang, and Elias Bareinboim. A graphical criterion for effect identification in equivalence classes of causal diagrams. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence*, IJCAI'18, pages 5024–5030, July 2018.

[Jaber *et al.*, 2019] Amin Jaber, Jiji Zhang, and Elias Bareinboim. Causal identification under markov equivalence: Completeness results. In *Proceedings of the 36th International Conference on Machine Learning*, ICML'19, pages 2981–2989, June 2019.

[Kuroki and Miyakawa, 1999] Manabu Kuroki and Masami Miyakawa. Identifiability criteria for causal effects of joint interventions. *Journal of the Japan Statistical Society*, 29(2):105–117, 1999.

[Pearl, 1993] Judea Pearl. Aspects of graphical models connected with causality. In *Proceedings of the 49th Session of the International Statistical Institute*, pages 391–401, Florence, Italy, 1993.

[Pearl, 1995] Judea Pearl. Causal diagrams for empirical research. *Biometrika*, 82(4):669–688, 1995.

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

[Perković *et al.*, 2015] Emilija Perković, Johannes Textor, Markus Kalisch, and Marloes H. Maathuis. A complete generalized adjustment criterion. In *Proceedings of the 31st Conference on Uncertainty in Artificial Intelligence*, UAI'15, pages 682–691, July 2015.

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

[Shpitser and Pearl, 2006] Ilya Shpitser and Judea Pearl. Identification of joint interventional distributions in recursive semi-Markovian causal models. In *Proceedings of the Twenty-First National Conference on Artificial Intelligence*, AAAI'06, pages 1219–1226, July 2006.

[Spirtes *et al.*, 2001] Peter Spirtes, Clark N Glymour, and Richard Scheines. *Causation, prediction, and search*, volume 81. MIT press, 2001.

[Tian and Pearl, 2002] Jin Tian and Judea Pearl. A general identification condition for causal effects. In *Proceedings of the Eighteenth National Conference on Artificial Intelligence*, AAAI'02, pages 567–573, 2002.

[Tian, 2002] Jin Tian. *Studies in causal reasoning and learning*. PhD thesis, University of California, Los Angeles, 2002.

[Verma, 1993] Thomas Verma. Graphical aspects of causal models. Technical report, R-191, UCLA, 1993.

[Zhang, 2007] Jiji Zhang. Generalized do-calculus with testable causal assumptions. In *Proceedings of the Eleventh International Conference on Artificial Intelligence and Statistics*, AISTATS'07, pages 667–674, 2007.

[Zhang, 2008a] Jiji Zhang. Causal reasoning with ancestral graphs. *Journal of Machine Learning Research*, 9(Jul):1437–1474, 2008.

[Zhang, 2008b] 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, 2008.