Identifying causal effects in maximally oriented partially directed acyclic graphs

Emilija Perković
Department of Statistics
University of Washington
Seattle, WA 98195-4322

Abstract

We develop a necessary and sufficient causal identification criterion for maximally oriented partially directed acyclic graphs (MPDAGs). MPDAGs as a class of graphs include directed acyclic graphs (DAGs), completed partially directed acyclic graphs (CPDAGs), and CPDAGs with added background knowledge. As such, they represent the type of graph that can be learned from observational data and background knowledge under the assumption of no latent variables. Our identification criterion can be seen as a generalization of the g-formula of Robins (1986). We further obtain a generalization of the truncated factorization formula (Pearl, 2009) and compare our criterion to the generalized adjustment criterion of Perković et al. (2017) which is sufficient, but not necessary for causal identification.

1 INTRODUCTION

The gold standard method for answering causal questions are randomized controlled trials. In some cases, however, it may be impossible, unethical, or simply too expensive to perform a desired experiment. For this purpose, it is of interest to consider whether a causal effect can be identified from observational data.

We consider the problem of identifying causal effects from a causal graph that represents the observational data under the assumption of causal sufficiency. If the causal directed acyclic graph (DAGs, e.g. Pearl, 2009) is known, then all causal effects can be identified and estimated from observational data (see e.g. Robins, 1986; Pearl, 1995; Pearl and Robins, 1995; Galles and Pearl, 1995).

In general, however, it is not possible to learn the underlying causal DAG from observational data. When all variables in the causal system are observed, one can at most learn a completed partially directed acyclic graph (CPDAG, Meek, 1995; Andersson et al., 1997; Spirtes et al., 2000; Chickering, 2002). A CPDAG uniquely represents a Markov equivalence class of DAGs (see Section 2 for definitions).

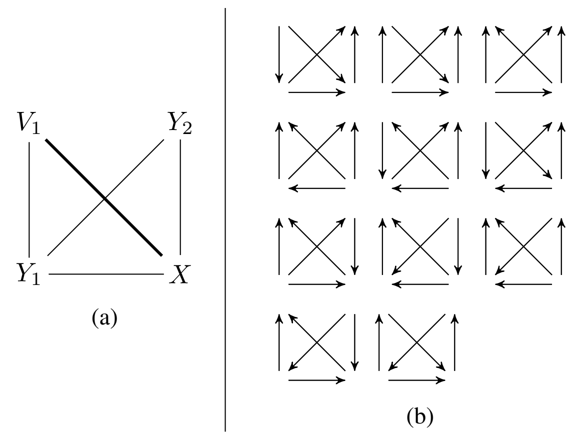

Figure 1: (a) CPDAG $\mathcal{C}$, (b) DAGs represented by $\mathcal{C}$.

If in addition to observational data one has background knowledge of some pairwise causal relationships, one can obtain a maximally oriented partially directed acyclic graph (MPDAG) which uniquely represents a refinement of the Markov equivalence class of DAGs (Meek, 1995). Other types of background knowledge, such as tiered orderings, data from previous experiments, or specific model restrictions also induce MPDAGs (Scheines et al., 1998; Hoyer et al., 2008; Hauser and Bühlmann, 2012; Eigenmann et al., 2017; Wang et al., 2017; Rothenhäusler et al., 2018).

To understand the difference and connections between DAGs, CPDAGs and MPDAGs, consider graphs in Figures 1 and 4. Graph $\mathcal{C}$ in Figure 1(a) is an example of a CPDAG that can be learned given enough observational

Proceedings of the 36th Conference on Uncertainty in Artificial Intelligence (UAI), PMLR volume 124, 2020.

data on variables X, V_1, Y_1, and Y_2. All DAGs in the Markov equivalence class represented by $\mathcal{C}$ are given in Figure 1(b). Graph $\mathcal{G}$ in Figure 4(a) is an MPDAG that can be obtained from CPDAG $\mathcal{C}$ in Figure 1(a) and background knowledge that Y_1 is a cause of X and that X is a cause of Y_2 (see Meek, 1995 for details on incorporating this type of background knowledge). All DAGs represented by $\mathcal{G}$ are given in Figure 4(b) and are a subset of DAGs in Figure 1(b).

One can consider MPDAGs as a graph class that is generally more causally informative than CPDAGs and less causally informative than DAGs. Conversely, a CPDAG can be seen as a special case of an MPDAG when the added background knowledge is not additionally informative compared to the observational data. Similarly, a DAG is a special case of an MPDAG when the additional background knowledge is fully causally informative. We will use MPDAGs to refer to all graphs in this paper.

The topic of identifying causal effects in MPDAGs has generated a wealth of research in recent years. The most relevant recent work on this topic is the generalized adjustment criterion of Perković et al. (2015, 2017, 2018) which is sufficient but not necessary for the identification of causal effects. Perković et al. (2015, 2018, 2017) build on prior work of Pearl (1993); Shpitser et al. (2010); van der Zander et al. (2014) and Maathuis and Colombo (2015).

One criterion that is necessary and sufficient for identifying causal effects in DAGs is the g-formula of Robins (1986). The g-formula is one of the causal identification methods that has seen considerable use in practice (see e.g. Taubman et al., 2009; Young et al., 2011; Westreich et al., 2012). However, the g-formula has not yet been generalized to other types of MPDAGs (including CPDAGs).

In this paper, we develop a necessary and sufficient graphical criterion for identifying causal effects in MPDAGs. We refer to our identification criterion (Theorem 3.6) as the causal identification formula. The causal identification formula is a generalization of the g-formula of Robins (1986) to MPDAGs. Consequently, we also obtain a generalization of the truncated factorization formula (Pearl, 2009), i.e. the manipulated density formula (Spirtes et al., 2000) in Corollary 3.7.

From a theoretical perspective, it is of interest to note that the proof of our causal identification formula does not consider intervening on additional variables in the graph (Section 3.5). This alleviates concerns of whether such additional interventions are reasonable to assume as possible (see e.g. VanderWeele and Robinson, 2014; Kohler-Hausmann, 2018).

We compare our result to the generalized adjustment criterion of Perković et al. (2017) in Section 4. Even though the generalized adjustment criterion is not complete for causal identification, we characterize a special case in which it is "almost" complete in Proposition 4.2.

Lastly, Jaber et al. (2019) recently constructed a graphical algorithm that is necessary and sufficient for identifying causal effects from observational data that allows for hidden confounders. The class of graphs that Jaber et al. (2019) consider is fully characterized by conditional independences in the observed probability distribution of the data. Their algorithm builds on the work of Tian and Pearl (2002); Shpitser and Pearl (2006); Huang and Valtorta (2006) and Richardson et al. (2017). To put our work into wider context, we compare our approach to the approach taken by Jaber et al. (2019) in the discussion. Omitted proofs can be found in the Supplement.

2 PRELIMINARIES

We use capital letters (e.g. X) to denote nodes in a graph as well as random variables that these nodes represent. Similarly, bold capital letters (e.g. $\mathbf{X}$) are used to denote both sets of nodes in a graph as well as the random vectors that these nodes represent.

Nodes, Edges And Subgraphs. A graph $\mathcal{G} = (\mathbf{V}, \mathbf{E})$ consists of a set of nodes (variables) $\mathbf{V} = \{X_1, \dots, X_p\}$ and a set of edges $\mathbf{E}$. The graphs we consider are allowed to contain directed ($\rightarrow$) and undirected ($-$) edges and at most one edge between any two nodes. An *induced subgraph* $\mathcal{G}_{\mathbf{V}'} = (\mathbf{V}', \mathbf{E}')$ of $\mathcal{G} = (\mathbf{V}, \mathbf{E})$ consists of $\mathbf{V}' \subseteq \mathbf{V}$ and $\mathbf{E}' \subseteq \mathbf{E}$ where $\mathbf{E}'$ are all edges in $\mathbf{E}$ between nodes in $\mathbf{V}'$. An *undirected subgraph* $\mathcal{G}_{undir} = (\mathbf{V}, \mathbf{E}')$ of $\mathcal{G} = (\mathbf{V}, \mathbf{E})$ consists of $\mathbf{V}$ and $\mathbf{E}' \subseteq \mathbf{E}$ where $\mathbf{E}'$ are all undirected edges in $\mathbf{E}$.

Paths. A *path* p from X to Y in $\mathcal{G}$ is a sequence of distinct nodes $\langle X, \dots, Y \rangle$ in which every pair of successive nodes is adjacent. A path consisting of undirected edges in an *undirected* path. A *causal path* from X to Y is a path from X to Y in which all edges are directed towards Y, that is, $X \rightarrow \cdots \rightarrow Y$. Let $p = \langle X = V_0, \dots, V_k = Y \rangle$, $k \geq 1$ be a path in $\mathcal{G}$, p is a *possibly causal path* if no edge $V_i \leftarrow V_j, 0 \leq i < j \leq k$ is in $\mathcal{G}$. Otherwise, p is a *non-causal path* in $\mathcal{G}$ (see Definition 3.1 and Lemma 3.2 of Perković et al., 2017) (Lemma A.4 in the Supplement). For two disjoint subsets $\mathbf{X}$ and $\mathbf{Y}$ of $\mathbf{V}$, a path from $\mathbf{X}$ to $\mathbf{Y}$ is a path from some $X \in \mathbf{X}$ to some $Y \in \mathbf{Y}$. A path from $\mathbf{X}$ to $\mathbf{Y}$ is *proper* (w.r.t. $\mathbf{X}$) if only its first node is in $\mathbf{X}$.

Partially Directed And Directed Cycles. A causal path from X to Y and the edge $Y \to X$ form a *directed cycle*. A *partially directed cycle* is formed by a possibly causal path from X to Y, together with $Y \to X$.

Ancestral Relationships. If $X \to Y$, then X is a *parent* of Y. If there is a causal path from X to Y, then X is an *ancestor* of Y, and Y is a *descendant* of X. If there is a possibly causal path from X to Y, then X is a *possible ancestor* of Y. We use the convention that every node is a descendant, ancestor, and possible ancestor of itself. The sets of parents, ancestors, possible ancestors and descendants of X in $\mathcal{G}$ are denoted by $\mathrm{Pa}(X, \mathcal{G})$, $\mathrm{An}(X, \mathcal{G})$, $\mathrm{PossAn}(X, \mathcal{G})$ and $\mathrm{De}(X, \mathcal{G})$ respectively. For a set of nodes $\mathbf{X} \subseteq \mathbf{V}$, we let $\mathrm{Pa}(\mathbf{X}, \mathcal{G}) = (\cup_{X \in \mathbf{X}} \mathrm{Pa}(X, \mathcal{G})) \setminus \mathbf{X}$, $\mathrm{An}(\mathbf{X}, \mathcal{G}) = \cup_{X \in \mathbf{X}} \mathrm{An}(X, \mathcal{G})$, $\mathrm{PossAn}(\mathbf{X}, \mathcal{G}) = \cup_{X \in \mathbf{X}} \mathrm{PossAn}(X, \mathcal{G})$, and $\mathrm{De}(\mathbf{X}, \mathcal{G}) = \cup_{X \in \mathbf{X}} \mathrm{De}(X, \mathcal{G})$.

Undirected Connected Set. A node set $\mathbf{X}$ is an *undirected connected set* in graph $\mathcal{G}$ if for every two distinct nodes X_i and X_j in $\mathbf{X}$, there is an undirected path from X_i to X_j in $\mathcal{G}$.

Colliders, Shields, And Definite Status Paths. If a path p contains $X_i \to X_j \leftarrow X_k$ as a subpath, then X_j is a *collider* on p. A path $\langle X_i, X_j, X_k \rangle$ is an *unshielded triple* if X_i and X_k are not adjacent. A path is *unshielded* if all successive triples on the path are unshielded. A node X_j is a *definite non-collider* on a path p if the edge $X_i \leftarrow X_j$, or the edge $X_j \to X_k$ is on p, or if $X_i - X_j - X_k$ is a subpath of p and X_i is not adjacent to X_k. A node is of *definite status* on a path if it is a collider, a definite non-collider or an endpoint on the path. A path p is of definite status if every node on p is of definite status.

D-connection And Blocking. A definite status path p from X to Y is *d-connecting* given a node set $\mathbf{Z}$ ($X, Y \notin \mathbf{Z}$) if every definite non-collider on p is not in $\mathbf{Z}$, and every collider on p has a descendant in $\mathbf{Z}$. Otherwise, $\mathbf{Z}$ *blocks* p. If $\mathbf{Z}$ blocks all definite status paths between $\mathbf{X}$ and $\mathbf{Y}$ in MPDAG $\mathcal{G}$, then $\mathbf{X}$ is *d-separated* from $\mathbf{Y}$ given $\mathbf{Z}$ in $\mathcal{G}$ (Lemma C.1 of Henckel et al., 2019).

DAGs, PDAGs. A *directed graph* contains only directed edges. A *partially directed graph* may contain both directed and undirected edges. A directed graph without directed cycles is a *directed acyclic graph* (DAG). A *partially directed acyclic graph* (PDAG) is a partially directed graph without directed cycles.

Markov Equivalence And CPDAGs. (c.f. Meek, 1995; Andersson et al., 1997) All DAGs that encode the same d-separation relationships are *Markov equivalent* and form a *Markov equivalence class* of DAGs, which can be *represented* by a *completed partially directed acyclic graph* (CPDAG).

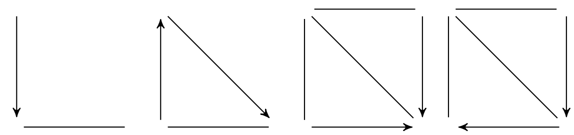

(a) (b) (c) (d)

Figure 2: Forbidden induced subgraphs of an MPDAG (see orientation rules in Meek, 1995).

MPDAGs. A PDAG $\mathcal{G}$ is a *maximally oriented* PDAG (MPDAG) if and only if the graphs in Figure 2 are **not** induced subgraphs of $\mathcal{G}$. Both a DAG and a CPDAG are types of MPDAG (Meek, 1995).

$\mathcal{G}$ And $[\mathcal{G}]$. A DAG $\mathcal{D}$ is *represented* by MPDAG $\mathcal{G}$ if $\mathcal{D}$ and $\mathcal{G}$ have the same adjacencies, same unshielded colliders and if for every directed edge $X \to Y$ in $\mathcal{G}$, $X \to Y$ is in $\mathcal{D}$ (Meek, 1995). If $\mathcal{G}$ is an MPDAG, then $[\mathcal{G}]$ denotes the set of all DAGs represented by $\mathcal{G}$.

Partial Causal Ordering. Let $\mathcal{D} = (\mathbf{V}, \mathbf{E})$ be a DAG. A total ordering, $<$, of nodes $\mathbf{V}' \subseteq \mathbf{V}$ is *consistent* with $\mathcal{D}$ and called a *causal ordering* of $\mathbf{V}'$ if for every $X_i, X_j \in \mathbf{V}'$, such that $X_i < X_j$ and such that X_i and X_j are adjacent in $\mathcal{D}$, $X_i \to X_j$ is in $\mathcal{D}$. There can be more than one causal ordering of $\mathbf{V}'$ in a DAG $\mathcal{D} = (\mathbf{V}, \mathbf{E})$. For example, in DAG $X_i \leftarrow X_j \to X_k$ both orderings $X_j < X_i < X_k$ and $X_j < X_k < X_i$ are consistent.

Let $\mathcal{G} = (\mathbf{V}, \mathbf{E})$ be an MPDAG. Since $\mathcal{G}$ may contain undirected edges, there is generally no causal ordering of $\mathbf{V}'$, for a node set $\mathbf{V}' \subseteq \mathbf{V}$ in $\mathcal{G} = (\mathbf{V}, \mathbf{E})$. Instead, we define a *partial causal ordering,* $<$, of $\mathbf{V}'$ in $\mathcal{G}$ as a total ordering of pairwise disjoint node sets $\mathbf{A_1}, \dots, \mathbf{A_k}$, $k \geq 1$, $\cup_{i=1}^{k} \mathbf{A_i} = \mathbf{V}'$, that satisfies the following: if $\mathbf{A_i} < \mathbf{A_j}$ and there is an edge between $A_i \in \mathbf{A_i}$ and $A_j \in \mathbf{A_j}$ in $\mathcal{G}$, then $A_i \to A_j$ is in $\mathcal{G}$.

Do-intervention. We consider interventions $do(\mathbf{X} = \mathbf{x})$ (for $\mathbf{X} \subseteq \mathbf{V}$) or $do(\mathbf{x})$ for shorthand, which represent outside interventions that set $\mathbf{X}$ to $\mathbf{x}$.

Observational And Interventional Densities. A density f of $\mathbf{V}$ is *consistent* with a DAG $\mathcal{D} = (\mathbf{V}, \mathbf{E})$ if it factorizes as $f(\mathbf{v}) = \prod_{V_i \in \mathbf{V}} f(v_i | \operatorname{pa}(v_i, \mathcal{D}))$ (Pearl, 2009). A density f that is consistent with $\mathcal{D} = (\mathbf{V}, \mathbf{E})$ is also called an *observational density*.

Let $\mathbf{X}$ be a subset of $\mathbf{V}$ and $\mathbf{V}' = \mathbf{V} \setminus \mathbf{X}$ in a DAG $\mathcal{D}$. A density over $\mathbf{V}'$ is denoted by $f(\mathbf{v}'|do(\mathbf{x}))$, or $f_{\mathbf{x}}(\mathbf{v}')$, and called an *interventional density consistent with* $\mathcal{D}$ if there is an observational density f consistent with $\mathcal{D}$ such that $f(\mathbf{v}'|do(\mathbf{x}))$ factorizes as

$$f(\mathbf{v}'|do(\mathbf{x})) = \prod_{V_i \in \mathbf{V}'} f(v_i | \operatorname{pa}(v_i, \mathcal{D})), \tag{1}$$

for values $\mathrm{pa}(v_i, \mathcal{D})$ of $\mathrm{Pa}(V_i, \mathcal{D})$ that are in agreement with $\mathbf{x}$. If $\mathbf{X} = \emptyset$, we define $f(\mathbf{v}|do(\emptyset)) = f(\mathbf{v})$. Equation (1) is known as the truncated factorization formula (Pearl, 2009), manipulated density formula (Spirtes et al., 2000) or the g-formula (Robins, 1986). A density f of $\mathbf{V}$ is consistent with an MPDAG $\mathcal{G} = (\mathbf{V}, \mathbf{E})$ if f is consistent with a DAG in $[\mathcal{G}]$.

A density $f(\mathbf{v}'|do(\mathbf{x}))$ of $\mathbf{V}' \subset \mathbf{V}$, $\mathbf{X} = \mathbf{V} \setminus \mathbf{V}'$ is an *interventional density* consistent with an MPDAG $\mathcal{G} = (\mathbf{V}, \mathbf{E})$ if it is an interventional density consistent with a DAG in $[\mathcal{G}]$. Let $\mathbf{Y} \subset \mathbf{V}'$, and let $f(\mathbf{v}'|do(\mathbf{x}))$ be an interventional density consistent with an MPDAG $\mathcal{G} = (\mathbf{V}, \mathbf{E})$ for some $\mathbf{X} \subset \mathbf{V}$, $\mathbf{V}' = \mathbf{V} \setminus \mathbf{X}$, then $f(\mathbf{y}|do(\mathbf{x}))$ denotes the marginal density of $\mathbf{Y}$ calculated from $f(\mathbf{v}'|do(\mathbf{x}))$.

Probabilistic Implications Of D-separation. Let f be any density over $\mathbf{V}$ consistent with an MPDAG $\mathcal{G} = (\mathbf{V}, \mathbf{E})$ and let $\mathbf{X}, \mathbf{Y}$, and $\mathbf{Z}$ be pairwise disjoint node sets in $\mathbf{V}$. If $\mathbf{X}$ and $\mathbf{Y}$ are d-separated given $\mathbf{Z}$ in $\mathcal{G}$, then $\mathbf{X}$ and $\mathbf{Y}$ are conditionally independent given $\mathbf{Z}$ in the observational probability density f consistent with $\mathcal{D}$ (Lauritzen et al., 1990; Pearl, 2009).

3 RESULTS

The causal effect of a set of treatments $\mathbf{X}$ on a set of responses $\mathbf{Y}$ is a function of the interventional density $f(\mathbf{y}|do(\mathbf{x}))$. For example, under the assumption of a Bernoulli distributed treatment variable X, the causal effect of X on a singleton response Y may be defined as the difference in expectation of Y under $do(X = 1)$ and $do(X = 0)$, that is, $E[Y|do(X = 1)] - E[Y|do(X = 0)]$ (Chapter 1 in Hernán and Robins, 2020).

We consider a causal effect to be identifiable in an MPDAG $\mathcal{G}$ if the interventional density of the response can be uniquely computed from $\mathcal{G}$. A precise definition is given in Definition 3.1. Definition 3.1 is analogous to the Definition 3 of Galles and Pearl (1995) and Definition 1 of Jaber et al. (2019).

Definition 3.1 (Identifiability of Causal Effects). *Let $\mathbf{X}$ and $\mathbf{Y}$ be disjoint node sets in an* MPDAG $\mathcal{G} = (\mathbf{V}, \mathbf{E})$. *The causal effect of $\mathbf{X}$ on $\mathbf{Y}$ is identifiable in $\mathcal{G}$ if $f(\mathbf{y}|do(\mathbf{x}))$ is uniquely computable from any observational density consistent with $\mathcal{G}$.*
Hence, there are no two DAGs $\mathcal{D}^1$, $\mathcal{D}^2$ in $[\mathcal{G}]$ such that

1. *$f_1(\mathbf{v}) = f_2(\mathbf{v}) = f(\mathbf{v})$, where f is an observational density consistent with $\mathcal{G}$, and*
2. *$f_1(\mathbf{y}|do(\mathbf{x})) \neq f_2(\mathbf{y}|do(\mathbf{x}))$, where $f_1(\cdot|do(\mathbf{x}))$ and $f_2(\cdot|do(\mathbf{x}))$ are interventional densities consistent with $\mathcal{D}^1$ and $\mathcal{D}^2$ respectively.*

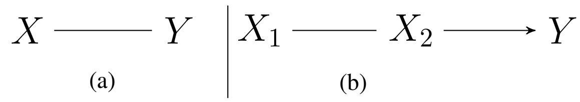

Figure 3: (a) MPDAG $\mathcal{C}$, (b) MPDAG $\mathcal{G}$.

3.1 A NECESSARY CONDITION FOR IDENTIFICATION

Proposition 3.2 presents a necessary condition for the identifiability of causal effects in MPDAGs. This necessary condition is referred to as amenability by Perković et al. (2015, 2017).

Proposition 3.2. *Let $\mathbf{X}$ and $\mathbf{Y}$ be disjoint node sets in an* MPDAG $\mathcal{G} = (\mathbf{V}, \mathbf{E})$. *If there is a proper possibly causal path from $\mathbf{X}$ to $\mathbf{Y}$ that starts with an undirected edge in $\mathcal{G}$, then the causal effect of $\mathbf{X}$ on $\mathbf{Y}$ is not identifiable in $\mathcal{G}$.*

Consider MPDAG $\mathcal{C}$ in Figure 3a. Since $X - Y$ is in $\mathcal{C}$, by Proposition 3.2, the causal effect of X on Y is not identifiable in $\mathcal{C}$. This is intuitively clear since both $X \to Y$ and $X \leftarrow Y$ are DAGs represented by $\mathcal{C}$. The DAG $X \leftarrow Y$ implies that there is no causal effect of X on Y. Conversely, the DAG $X \to Y$ implies that there is a causal effect of X on Y.

The condition in Proposition 3.2 is somewhat less intuitive for non-singleton $\mathbf{X}$. Consider MPDAG $\mathcal{G}$ in Figure 3b and let $\mathbf{X} = \{X_1, X_2\}$ and $\mathbf{Y} = \{Y\}$. The path $X_1 - X_2 \to Y$ in $\mathcal{G}$ is a possibly causal path from X_1 to Y that starts with an undirected edge. However, $X_1 - X_2 \to Y$ is not a proper possibly causal path from $\mathbf{X}$ to Y, since it contains X_2 and X_1. Hence, the causal effect of $\mathbf{X}$ on Y may still be identifiable in $\mathcal{G}$.

3.2 PARTIAL CAUSAL ORDERING IN MPDAGS

For the proof of our main result, it is necessary to determine a partial causal ordering for a set of nodes in an MPDAG. In order to compute a partial causal ordering of nodes in an MPDAG, we first define a bucket.

Definition 3.3 (Bucket). *Let $\mathbf{D}$ be a node set in an* MPDAG $\mathcal{G} = (\mathbf{V}, \mathbf{E})$. *If $\mathbf{B}$ is a maximal undirected connected subset of $\mathbf{D}$ in $\mathcal{G}$, we call $\mathbf{B}$ a bucket in $\mathbf{D}$.*

Definition 3.3 is similar to the definition of a bucket of Jaber et al. (2018a). One difference is that Definition 3.3 allows for directed edges between the nodes within the same bucket, whereas the definition of Jaber et al. (2018a) does not. For instance, $\{X, V_1, Y_1\}$ is a bucket in $\mathbf{V}$ in MPDAG $\mathcal{G} = (\mathbf{V}, \mathbf{E})$ in Figure 4(a). Note that since we require a bucket to be a maximal undirected connected set, $\{X, V_1\}$ is not a bucket in $\mathbf{V}$.

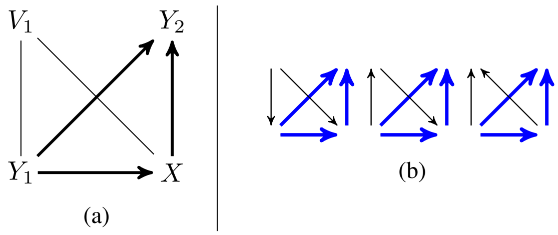

Figure 4: (a) MPDAG $\mathcal{G}$, (b) DAGs represented by $\mathcal{G}$.

Definition 3.3 can be used to induce a unique partition of any node set $\mathbf{D}$ in an MPDAG $\mathcal{G} = (\mathbf{V}, \mathbf{E})$, $\mathbf{D} \subseteq \mathbf{V}$. We refer to this partition as *the bucket decomposition* in the corollary of Definition 3.3 below.

Corollary 3.4 (**Bucket Decomposition**). *Let* $\mathbf{D}$ *be a node set in an* MPDAG $\mathcal{G} = (\mathbf{V}, \mathbf{E})$. *Then there is a unique partition of* $\mathbf{D}$ *into* $\mathbf{B_1}, \dots \mathbf{B_k}$, $k \geq 1$ *in* $\mathcal{G}$ *induced by Definition 3.3. That is*

- $\mathbf{D} = \cup_{i=1}^{k} \mathbf{B_i}$, *and*
- $\mathbf{B_i} \cap \mathbf{B_j} = \emptyset$, $i, j \in \{1, \dots, k\}$, $i \neq j$, *and*
- $\mathbf{B_i}$ *is a bucket in* $\mathbf{D}$ *for each* $i \in \{1, \dots, k\}$.

Algorithm 1: Partial causal ordering (PCO)

input : Node set $\mathbf{D}$ in MPDAG $\mathcal{G}$=$(\mathbf{V}, \mathbf{E})$.
output : An ordered list $\mathbf{B}$=$(\mathbf{B_1}, \dots, \mathbf{B_k})$, $k \geq 1$, of the bucket decomposition of $\mathbf{D}$ in $\mathcal{G}$.

1 Let $\mathbf{ConComp}$ be the bucket decomposition of $\mathbf{V}$ in $\mathcal{G}$;
2 Let $\mathbf{B}$ be an empty list;
3 **while** $\mathbf{ConComp} \neq \emptyset$ **do**
4   Let $\mathbf{C} \in \mathbf{ConComp}$;
5   Let $\overline{\mathbf{C}}$ be the set of nodes in $\mathbf{ConComp}$ that are not in $\mathbf{C}$;
6   **if** *all edges between* $\mathbf{C}$ *and* $\overline{\mathbf{C}}$ *are into* $\mathbf{C}$ *in* $\mathcal{G}$ **then**
7   Remove $\mathbf{C}$ from $\mathbf{ConComp}$;
8   Let $\mathbf{B}_* = \mathbf{C} \cap \mathbf{D}$;
9   **if** $\mathbf{B}_* \neq \emptyset$ **then**
10   Add $\mathbf{B}_*$ to the beginning of $\mathbf{B}$;
11   **end**
12   **end**
13 **end**
14 **return** $\mathbf{B}$;

Consider MPDAG $\mathcal{G} = (\mathbf{V}, \mathbf{E})$ in Figure 4a. In order to find the bucket decomposition of $\mathbf{V}$ in $\mathcal{G}$, let us consider the undirected subgraph $\mathcal{G}_{undir}$ of $\mathcal{G}$. The only path in $\mathcal{G}_{undir}$ is $Y_1 - V_1 - X$. Hence, the bucket decomposition of $\mathbf{V}$ is $\{\{X, V_1, Y_1\}, \{Y_2\}\}$.

Consider DAGs in Figure 4b, which are all DAGs represented by $\mathcal{G}$ in Figure 4a. Some total orderings of $\mathbf{V}$ that are consistent with DAGs in Figure 4b are: $V_1 < Y_1 < X < Y_2$, $Y_1 < V_1 < X < Y_2$, and $Y_1 < X < V_1 < Y_2$, from left to right respectively. These three orderings are consistent with the following partial causal ordering $\{X, V_1, Y_1\} < Y_2$, which is a total ordering of the buckets in the bucket decomposition of $\mathbf{V}$. This motivates Algorithm 1.

Algorithm 1 outputs an ordered bucket decomposition of a set of nodes $\mathbf{D}$ in an MPDAG $\mathcal{G}$. The proof that Algorithm 1 will always complete is given in Lemma C.1 in the Supplement. Next, we prove that the ordered list of buckets output by Algorithm 1 is a partial causal ordering of $\mathbf{D}$ in $\mathcal{G}$ (Lemma 3.5). Algorithm 1 and Lemma 3.5 are similar to the PTO algorithm and Lemma 1 of Jaber et al. (2018b).

Lemma 3.5. *Let* $\mathbf{D}$ *be a node set in an* MPDAG $\mathcal{G} = (\mathbf{V}, \mathbf{E})$ *and let* $(\mathbf{B_1}, \dots, \mathbf{B_k})$, $k \geq 1$, *be the output of* PCO$(\mathbf{D}, \mathcal{G})$. *Then for each* $i, j \in \{1, \dots k\}$, $\mathbf{B_i}$ *and* $\mathbf{B_j}$ *are buckets in* $\mathbf{D}$ *and if* $i < j$, *then* $\mathbf{B_i} < \mathbf{B_j}$ *in* $\mathcal{G}$.

Consider MPDAG $\mathcal{G} = (\mathbf{V}, \mathbf{E})$ in Figure 4a and let $\mathbf{D} = \{X, Y_1, Y_2\}$. We now explain how the output of PCO$(\mathbf{D}, \mathcal{G})$ is obtained.

In line 2, the bucket decomposition of $\mathbf{V}$ is obtained, $\mathbf{ConComp} = \{\{X, Y_1, V_1\}, \{Y_2\}\}$ (as noted above). In line 2, $\mathbf{B}$ is initialized as an empty list.

Let $\mathbf{C} = \{X, Y_1, V_1\}$ (line 4). Then $\overline{\mathbf{C}} = \{Y_2\}$ (line 5). Since $X \to Y_2$ and $Y_1 \to Y_2$ are in $\mathcal{G}$, $\mathbf{C}$ does not satisfy the condition in line 6 and hence, $\{X, Y_1, V_1\}$ cannot be removed from $\mathbf{ConComp}$ at this time.

Next, $\mathbf{C} = \{Y_2\}$ (line 4) and $\overline{\mathbf{C}} = \{X, Y_1, V_1\}$ (line 5). Since all edges between $\{Y_2\}$ and $\{X, Y_1, V_1\}$ in $\mathcal{G}$ are into $\{Y_2\}$, Algorithm 1 removes $\{Y_2\}$ from $\mathbf{ConComp}$ in line 7. Since $\mathbf{B}_* = \mathbf{C} \cap \mathbf{D} = \{Y_2\}$ (line 8), Algorithm 1 adds $\{Y_2\}$ to the beginning of list $\mathbf{B}$ (line 10).

Now, $\mathbf{C} = \{X, Y_1, V_1\}$ (line 4) and $\overline{\mathbf{C}} = \emptyset$ (line 5). Hence, $\mathbf{C}$ satisfies condition in line 6 and $\mathbf{C}$ is removed from $\mathbf{ConComp}$ (line 7). Then $\mathbf{B}_* = \mathbf{C} \cap \mathbf{D} = \{X, Y_1\}$ (line 8), and $\mathbf{B} = (\{X, Y_1\}, \{Y_2\})$ (line 10). Since $\mathbf{ConComp}$ is empty, Algorithm 1 outputs $\mathbf{B}$.

3.3 CAUSAL IDENTIFICATION FORMULA

We present our main result which we refer to as the causal identification formula in Theorem 3.6. Theorem 3.6 establishes that the condition from Proposition 3.2 is not only necessary, but also sufficient for the identification of causal effects in MPDAGs.

Theorem 3.6 (Causal identification formula). *Let* $\mathbf{X}$ *and* $\mathbf{Y}$ *be disjoint node sets in an* MPDAG $\mathcal{G} = (\mathbf{V}, \mathbf{E})$. *If there is no proper possibly causal path from* $\mathbf{X}$ *to* $\mathbf{Y}$ *in* $\mathcal{G}$ *that starts with an undirected edge, then for any observational density* f *consistent with* $\mathcal{G}$ *we have*

$$f(\mathbf{y}|do(\mathbf{x})) = \int \prod_{i=1}^{k} f(\mathbf{b_i}|\operatorname{pa}(\mathbf{b_i}, \mathcal{G}))d\mathbf{b}, \qquad (2)$$

for values $\operatorname{pa}(\mathbf{b_i}, \mathcal{G})$ *of* $\operatorname{Pa}(\mathbf{b_i}, \mathcal{G})$ *that are in agreement with* $\mathbf{x}$, *where* $(\mathbf{B_1}, \dots, \mathbf{B_k}) = \mathit{PCO}(\operatorname{An}(\mathbf{Y}, \mathcal{G}_{\mathbf{V}\setminus\mathbf{X}}), \mathcal{G})$ *and* $\mathbf{B} = \operatorname{An}(\mathbf{Y}, \mathcal{G}_{\mathbf{V}\setminus\mathbf{X}}) \setminus \mathbf{Y}$.

For a DAG $\mathcal{D} = (\mathbf{V}, \mathbf{E})$, it is well known that in order to identify a causal effect of $\mathbf{X}$ on $\mathbf{Y}$ in $\mathcal{D}$ it is enough to consider the set of ancestors of $\mathbf{Y}$, that is $\operatorname{An}(\mathbf{Y}, \mathcal{D})$ (see Theorem 4 of Tian and Pearl, 2002). The causal identification formula refines this notion by using a subset of ancestors of $\mathbf{Y}$ to identify the causal effect of $\mathbf{X}$ on $\mathbf{Y}$ in an MPDAG $\mathcal{G}$. The variables that appear on the right hand side of equation (2) are in $\operatorname{An}(\mathbf{Y}, \mathcal{G}_{\mathbf{V}\setminus\mathbf{X}})$, or in $\mathbf{X}$, for those $\mathbf{X}$ that have a proper causal path to $\mathbf{Y}$ in $\mathcal{G}$.

The causal identification formula is a generalization of the g-formula of Robins (1986), the truncated factorization formula of Pearl (2009), or the manipulated density formula of Spirtes et al. (2000) to the case of MPDAGs. To further exhibit this connection, we include the following corollary.

Corollary 3.7 (Factorization and truncated factorization formula in MPDAG**s).** *Let* $\mathbf{X}$ *be a node set in an* MPDAG $\mathcal{G} = (\mathbf{V}, \mathbf{E})$ *and let* $\mathbf{V}' = \mathbf{V} \setminus \mathbf{X}$. *Furthermore, let* $(\mathbf{V_1}, \dots, \mathbf{V_k})$ *be the output of* $\mathit{PCO}(\mathbf{V}, \mathcal{G})$. *Then for any observational density* f *consistent with* $\mathcal{G}$ *we have*

1. $f(\mathbf{v}) = \prod_{\mathbf{V_i}\subseteq\mathbf{V}} f(\mathbf{v_i}|\operatorname{pa}(\mathbf{v_i}, \mathcal{G}))$,
2. *If there is no pair of nodes* $V \in \mathbf{V}'$ *and* $X \in \mathbf{X}$ *such that* $X - V$ *is in* $\mathcal{G}$, *then*

$$f(\mathbf{v}'|do(\mathbf{x})) = \prod_{\mathbf{V_i}\subseteq\mathbf{V}'} f(\mathbf{v_i}|\operatorname{pa}(\mathbf{v_i}, \mathcal{G})),$$

for values $\operatorname{pa}(\mathbf{v_i}, \mathcal{G})$ *of* $\operatorname{Pa}(\mathbf{v_i}, \mathcal{G})$ *that are in agreement with* $\mathbf{x}$.

Whenever $f(\mathbf{v}'|do(\mathbf{x}))$ is identifiable in MPDAG $\mathcal{G} = (\mathbf{V}, \mathbf{E})$, we can also identify $f(\mathbf{y}|do(\mathbf{x}))$ as

$$f(\mathbf{y}|do(\mathbf{x})) = \int f(\mathbf{v}'|do(\mathbf{x}))d\overline{\mathbf{v}'},$$

where $\mathbf{X}$ and $\mathbf{Y}$ are disjoint subsets of $\mathbf{V}$, $\mathbf{V}' = \mathbf{V} \setminus \mathbf{X}$, and $\overline{\mathbf{V}'} = \mathbf{V} \setminus \{\mathbf{X} \cup \mathbf{Y}\}$. Since the necessary condition for identifying $f(\mathbf{v}'|do(\mathbf{x}))$ (Corollary 3.7) is generally stronger than the necessary condition for identifying

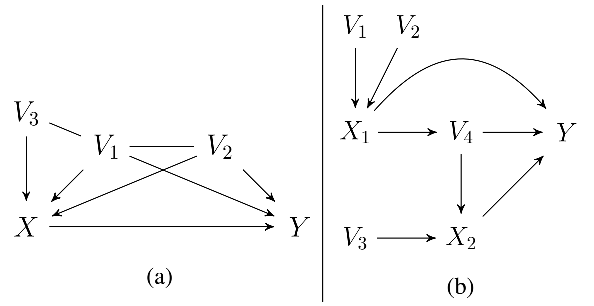

Figure 5: (a) MPDAG $\mathcal{G}$, (b) DAG $\mathcal{D}$.

$f(\mathbf{y}|do(\mathbf{x}))$ there are cases when $f(\mathbf{y}|do(\mathbf{x}))$ is identifiable and $f(\mathbf{v}'|do(\mathbf{x}))$ is not identifiable. One such case is explored in Example 3.8.

3.4 EXAMPLES

Example 3.8. *In this example, the causal effect of* $\mathbf{X}$ *on* $\mathbf{Y}$ *is identifiable in an* MPDAG $\mathcal{G} = (\mathbf{V}, \mathbf{E})$, *but the effect of* $\mathbf{X}$ *on* $\mathbf{V}' = \mathbf{V} \setminus \mathbf{X}$ *is not identifiable in* $\mathcal{G}$.

Consider MPDAG $\mathcal{G}$ *in Figure 4a and let* f *be an observational density consistent with* $\mathcal{G}$. *Let* $\mathbf{X} = \{X\}$ *and* $\mathbf{Y} = \{Y_1, Y_2\}$. *Note that path* $X - V_1 - Y_1$ *while proper is not possibly causal from* X *to* Y_1 *in* $\mathcal{G}$ *due to edge* $Y_1 \to X$. *The only possibly causal path from* X *to* $\mathbf{Y}$ *in* $\mathcal{G}$ *is* $X \to Y_2$. *Hence, by Theorem 3.6, the causal effect of* X *on* $\mathbf{Y}$ *is identifiable in* $\mathcal{G}$.

To use the causal identification formula we first determine that $\operatorname{An}(\{Y_1, Y_2\}, \mathcal{G}_{\mathbf{V}\setminus\{X\}}) = \{Y_1, Y_2\}$, *the bucket decomposition of* $\{Y_1, Y_2\}$ *is* $\{\{Y_1\}, \{Y_2\}\}$, *and* $\mathit{PCO}(\{Y_1, Y_2\}, \mathcal{G}) = (\{Y_1\}, \{Y_2\})$. *Next,* $\operatorname{Pa}(Y_1, \mathcal{G}) = \emptyset$, *and* $\operatorname{Pa}(Y_2, \mathcal{G}) = \{X, Y_1\}$. *Hence, by Theorem 3.6,* $f(y_1, y_2|do(x)) = f(y_2|x, y_1)f(y_1)$.

Now, let $\mathbf{V}' = \mathbf{V} \setminus \{X\}$. *Since* $X - V_1$ *is in* $\mathcal{G}$, *by Corollary 3.7,* $f(\mathbf{v}'|do(\mathbf{x}))$ *is not identifiable in* $\mathcal{G}$.

Example 3.9. *In this example, both the causal effect of* X *on* Y *and the causal effect of* X *on* $\mathbf{V}' = \mathbf{V} \setminus \{X\}$ *are identifiable in an* MPDAG $\mathcal{G} = (\mathbf{V}, \mathbf{E})$.

Consider MPDAG $\mathcal{G}$ *in Figure 5a and let* f *be an observational density consistent with* $\mathcal{G}$. *The only possibly causal path from* X *to* Y *in* $\mathcal{G}$ *is* $X \to Y$. *Hence, the causal effect of* X *on* Y *is identifiable in* $\mathcal{G}$.

In fact, there are no undirected edges connected to X, *so the causal effect of* X *on* $\mathbf{V}'$, $\mathbf{V}' = \{V_1, V_2, V_3, Y\}$ *is also identifiable in* $\mathcal{G}$. *Thus, we can obtain the truncated factorization formula with respect to* X *in* $\mathcal{G}$.

We will first determine the causal identification formula for $f(y|do(x))$ *in* $\mathcal{G}$. *We first identify that* $\operatorname{An}(Y, \mathcal{G}_{\mathbf{V}\setminus\{X\}}) = \{V_1, V_2, Y\}$. *The bucket de-*

composition of $\{V_1, V_2, Y\}$ *is* $\{\{V_1, V_2\}, \{Y\}\}$ *and* $\mathcal{PCO}(\{V_1, V_2, Y\}, \mathcal{G})$ *is* $(\{V_1, V_2\}, \{Y\})$. *Furthermore,* $\mathrm{Pa}(\{V_1, V_2\}, \mathcal{G}) = \emptyset$, $\mathrm{Pa}(Y, \mathcal{G}) = \{X, V_1, V_2\}$. *Hence, by Theorem 3.6, the causal identification formula for* $f(y|do(x))$ *in* $\mathcal{G}$ *is* $f(y|do(x)) = \int f(y|x, v_1, v_2) f(v_1, v_2) dv_1 dv_2$.

To use Corollary 3.7, first note that the output of $\mathcal{PCO}(\mathbf{V}, \mathcal{G})$ *is* $(\{V_1, V_2, V_3\}, \{X\}, \{Y\})$ *and that the ordered bucket decoposition of* $\mathbf{V}'$ *is* $(\{V_1, V_2, V_3\}, \{Y\})$. *Further,* $\mathrm{Pa}(\{V_1, V_2, V_3\}, \mathcal{G}) = \emptyset$. *Then,* $f(\mathbf{v}'|do(x)) = f(y|x, v_1, v_2) f(v_1, v_2, v_3)$.

Example 3.10. *This example shows how the causal identification formula can be used to estimate the causal effect of* $\mathbf{X}$ *on* $\mathbf{Y}$ *in an* MPDAG $\mathcal{G}$ *under the assumption that the observational density* f *consistent with* $\mathcal{G}$ *is multivariate Gaussian.*

Consider DAG $\mathcal{D}$ *in Figure 5b and let* f *be an observational density consistent with* $\mathcal{D}$. *Further, let* $\mathbf{X} = \{X_1, X_2\}$ *and* $\mathbf{Y} = \{Y\}$. *Then* $\mathrm{An}(Y, \mathcal{D}_{\mathbf{V} \setminus \mathbf{X}}) = \{Y, V_4\}$, *the bucket decomposition of* $\{Y, V_4\}$ *is* $\{\{V_4\}, \{Y\}\}$, *and* $\mathcal{PCO}(\{Y, V_4\}, \mathcal{D}) = (\{V_4\}, \{Y\})$ *in* $\mathcal{D}$.

Since $\mathrm{Pa}(V_4, \mathcal{D}) = \{X_1\}$, *and* $\mathrm{Pa}(Y, \mathcal{D}) = \{X_1, X_2, V_4\}$, *by Theorem 3.6,*

$$f(y|do(x_1, x_2)) = \int f(y|x_1, x_2, v_4) f(v_4|x_1) dv_4.$$

Suppose that the density f *consistent with* $\mathcal{D}$ *is multivariate Gaussian. The causal effect of* $\mathbf{X}$ *on* Y *can then be defined as the vector*

$$\left(\frac{\partial E[Y|do(x_1, x_2)]}{\partial x_1}, \frac{\partial E[Y|do(x_1, x_2)]}{\partial x_2} \right)^T,$$

(Nandy et al., 2017). Hence, consider $E[Y|do(x_1, x_2)]$,

$$E[Y|do(x_1, x_2)] = \int y f(y|do(x_1, x_2)) dy$$
$$= \int \int y f(y|x_1, x_2, v_4) f(v_4|x_1) dv_4 dy$$
$$= \int E[Y|x_1, x_2, v_4] f(v_4|x_1) dv_4$$
$$= \alpha x_1 + \beta x_2 + \gamma \int v_4 f(v_4|x_1) dv_4$$
$$= \beta x_2 + x_1(\alpha + \gamma\delta),$$

where $E[Y|x_1, x_2, v_4] = \alpha x_1 + \beta x_2 + \gamma v_4$ *and* $E[V_4|x_1] = \delta x_1$ *(Theorem 3.2.4 of Mardia et al., 1980, see Theorem A.2 in the Supplement).*

The causal effect of $\mathbf{X}$ *on* $\mathbf{Y}$ *is equal to* $(\alpha + \gamma\delta, \beta)$. *Consistent estimators for* α, β, *and* γ *are the least-squares estimators of the respective coefficients of* X_1, X_2, *and* V_4 *in the regression of* Y *on* X_1, X_2, *and* V_4. *Analogously, the consistent estimator for* δ *is the least-squares coefficient of* X_1 *in the regression of* V_4 *on* X_1.

3.5 PROOF OF THEOREM 3.6

The proof of Theorem 3.6 relies on Lemma D.1 in the Supplement. Lemma D.1 is proven through use of do-calculus (Pearl, 2009) and basic probability calculus.

The proofs of Theorem 3.6 and Lemma D.1 do not require intervening on additional variables in $\mathcal{G}$. This fact alleviates any concerns of whether such additional interventions are reasonable to assume as possible (see e.g. VanderWeele and Robinson, 2014; Kohler-Hausmann, 2018).

Proof of Theorem 3.6. For $i \in \{2, \dots, k\}$, let $\mathbf{P_i} = (\cup_{j=1}^{i-1} \mathbf{B_i}) \cap \mathrm{Pa}(\mathbf{B_i}, \mathcal{G})$. For $i \in \{1, \dots, k\}$, let $\mathbf{X_{p_i}} = \mathbf{X} \cap \mathrm{Pa}(\mathbf{B_i}, \mathcal{G})$.

Then

$$f(\mathbf{y}|do(\mathbf{x})) = \int f(\mathbf{b}, \mathbf{y}|do(\mathbf{x})) d\mathbf{b}$$
$$= \int f(\mathbf{b_1}|do(\mathbf{x})) \prod_{i=2}^{k} f(\mathbf{b_i}|\mathbf{b_{i-1}}, \dots, \mathbf{b_1}, do(\mathbf{x})) d\mathbf{b}$$
$$= \int f(\mathbf{b_1}|do(\mathbf{x})) \prod_{i=2}^{k} f(\mathbf{b_i}|\mathbf{p_i}, do(\mathbf{x})) d\mathbf{b} \tag{3}$$
$$= \int f(\mathbf{b_1}|do(\mathbf{x_{p_1}})) \prod_{i=2}^{k} f(\mathbf{b_i}|\mathbf{p_i}, do(\mathbf{x_{p_i}})) d\mathbf{b} \tag{4}$$
$$= \int \prod_{i=1}^{k} f(\mathbf{b_i}|pa(\mathbf{b_i}, \mathcal{G})) d\mathbf{b}, \tag{5}$$

The first two equalities follow from the law of total probability and the chain rule. Equations (3), (4), and (5) follow by applying results (ii), (iii), and (iv) in Lemma D.1 in the Supplement. □

4 COMPARISON TO ADJUSTMENT

The current state-of-the-art method for identifying causal effects in MPDAGs is the generalized adjustment criterion of Perković et al. (2017) stated in Theorem 4.1.

Theorem 4.1 (Adjustment set, Generalized adjustment criterion; Perković et al., 2017**).** *Let* $\mathbf{X}, \mathbf{Y}$ *and* $\mathbf{Z}$ *be pairwise disjoint node sets in an* MPDAG $\mathcal{G} = (\mathbf{V}, \mathbf{E})$. *Let* f *be any observational density consistent with* $\mathcal{G}$.

Then $\mathbf{Z}$ *is an adjustment set relative to* $(\mathbf{X}, \mathbf{Y})$ *in* $\mathcal{G}$ *and* *we have*

$$f(\mathbf{y}|do(\mathbf{x})) = \begin{cases} \int f(\mathbf{y}|\mathbf{x},\mathbf{z})f(\mathbf{z})d\mathbf{z} & \text{, if } \mathbf{Z} \neq \emptyset, \\ f(\mathbf{y}|\mathbf{x}) & \text{, otherwise.} \end{cases}$$

if and only if the following conditions are satisfied:

1. *There is no proper possibly causal path from* $\mathbf{X}$ *to* $\mathbf{Y}$ *that starts with an undirected edge in* $\mathcal{G}$.
2. $\mathbf{Z} \cap \mathrm{Forb}(\mathbf{X}, \mathbf{Y}, \mathcal{G}) = \emptyset$, *where*

   $\mathrm{Forb}(\mathbf{X}, \mathbf{Y}, \mathcal{G}) = \{W' \in \mathbf{V} : W' \in \mathrm{PossDe}(W, \mathcal{G}),$ *for some* $W \notin \mathbf{X}$ *which lies on a proper possibly causal path from* $\mathbf{X}$ *to* $\mathbf{Y}$ *in* $\mathcal{G}\}$.
3. *All proper non-causal definite status paths from* $\mathbf{X}$ *to* $\mathbf{Y}$ *are blocked by* $\mathbf{Z}$ *in* $\mathcal{G}$.

The generalized adjustment criterion is sufficient for identifying causal effects in an MPDAG, but it is not necessary. However, when $\mathbf{X}$ and $\mathbf{Y}$ are singleton sets, the generalized adjustment criterion identifies all non-zero causal effects of $\mathbf{X}$ on $\mathbf{Y}$ in an MPDAG $\mathcal{G}$. This is shown in the following proposition.

Proposition 4.2. *Let* X *and* Y *be distinct nodes in an* MPDAG $\mathcal{G} = (\mathbf{V}, \mathbf{E})$. *If* $Y \notin \mathrm{Pa}(X, \mathcal{G})$, *then the causal effect of* X *on* Y *is identifiable in* $\mathcal{G}$ *if and only if there is an adjustment set relative to* (X, Y) *in* $\mathcal{G}$.

Furthermore, if $Y \notin \mathrm{Pa}(X, \mathcal{G})$, *then* $\mathrm{Pa}(X, \mathcal{G})$ *is an adjustment set relative to* (X, Y) *in* $\mathcal{G}$ *whenever one such set exists.*

If $Y \in \mathrm{Pa}(X, \mathcal{G})$, then due to the acyclicity of $\mathcal{G}$, there is no causal path from X to Y in $\mathcal{G}$ and therefore no causal effect of X on Y (see Lemma E.1 in the Supplement). Hence, by Proposition 4.2, the generalized adjustment criterion is "almost" complete for the identification of causal effects of variable X on a response Y in MPDAGs.

If $\mathbf{X}$, or $\mathbf{Y}$ are non-singleton sets in $\mathcal{G}$, however, the generalized adjustment criterion will fail to identify some non-zero causal effects of $\mathbf{X}$ on $\mathbf{Y}$. We discuss this further in the two examples below.

Example 4.3. *Consider* MPDAG $\mathcal{G}$ *in Figure 4a and let* $\mathbf{X} = \{X\}$, *and* $\mathbf{Y} = \{Y_1, Y_2\}$ *as in Example 3.8.*

Path $X \leftarrow Y_1$ *is a non-causal path from* X *to* $\mathbf{Y}$ *that cannot be blocked by any set of nodes disjoint with* $\{X, Y_1\}$. *Hence, there is no adjustment set relative to* $(X, \mathbf{Y})$ *in* $\mathcal{G}$. *But there is a causal path from* X *to* $\mathbf{Y}$ *in* $\mathcal{G}$ *and as we have seen in Example 3.8, the causal effect of* X *on* $\mathbf{Y}$ *is identifiable in* $\mathcal{G}$.

Example 4.4. *Consider DAG* $\mathcal{D}$ *in Figure 5b and let* $\mathbf{X} = \{X_1, X_2\}$, *and* $\mathbf{Y} = \{Y\}$. *Then* $\mathrm{Forb}(\mathbf{X}, \mathbf{Y}, \mathcal{D}) = \{V_4, Y\}$. *For a set* $\mathbf{Z}$ *to satisfy the generalized adjustment criterion relative to* $(\mathbf{X}, Y)$ *in* $\mathcal{G}$, $\mathbf{Z}$ *cannot contain nodes in* $\{V_4, Y\}$, *or* $\{X_1, X_2\}$ *and* $\mathbf{Z}$ *must block all proper non-causal paths from* $\mathbf{X}$ *to* Y *in* $\mathcal{D}$.

However, $X_2 \leftarrow V_4 \rightarrow Y$ *is a proper non-causal path from* $\mathbf{X}$ *to* Y *in* $\mathcal{D}$ *that cannot be blocked by any set* $\mathbf{Z}$ *that satisfies* $\mathbf{Z} \cap \{X_1, X_2, V_4, Y\} = \emptyset$. *Hence, there is no adjustment set relative to* $(\mathbf{X}, Y)$ *in* $\mathcal{D}$. *But as we have seen in Example 3.10, the causal effect of* $\mathbf{X}$ *on* Y *is identifiable in* $\mathcal{D}$ *and furthermore, both* X_1 *and* X_2 *are causes of* Y *in* $\mathcal{D}$.

5 DISCUSSION

We introduced a causal identification formula that allows complete identification of causal effects in MPDAGs. Furthermore, we gave a comparison of our graphical criterion to the current state of the art method for causal identification in MPDAGs.

Since the causal identification formula comes in the familiar form of the g-formula of Robins (1986) for DAGs, our results can be used to generalize applications of the g-formula to MPDAGs. For example, Murphy (2003), Collins et al. (2004), and Collins et al. (2007) give criteria for estimating the optimal dynamic treatment regime from longitudinal data that are based on the g-formula. This idea can further be combined with recent work of Rahmadi et al. (2017) and Rahmadi et al. (2018) that establishes an approach for estimating the MPDAG using data from longitudinal studies.

Throughout the paper, we assume no latent variables. When latent variables are present, one can at most learn a partial ancestral graph (PAG) over the set of observed variables from the observed data (Richardson and Spirtes, 2002; Spirtes et al., 2000; Zhang, 2008a,b). PAGs represent an equivalence class of DAGs over the set of observed and unobserved variables.

Jaber et al. (2019) recently developed a recursive graphical algorithm that is both necessary and sufficient for identifying causal effects in PAGs. Our causal identification formula does not follow as a simplification of the result of Jaber et al. (2019). To see this, notice that the strategy of Jaber et al. (2019) for identifying causal effects in PAG $\mathcal{P}$ relies on the fact that the causal effect of $\mathbf{X}$ on $\mathbf{Y}$ is identifiable in $\mathcal{P}$ if and only if the causal effect of $\mathbf{V} \setminus \mathrm{PossAn}(\mathbf{Y}, \mathcal{P}_{\mathbf{V}\setminus\mathbf{X}})$ on $\mathrm{PossAn}(\mathbf{Y}, \mathcal{P}_{\mathbf{V}\setminus\mathbf{X}})$ is identifiable in $\mathcal{P}$ (see equation (8) of Jaber et al., 2019).

Consider applying this strategy to MPDAG $\mathcal{G}$ in Figure 4(a), with $\mathbf{X} = \{X\}$ and $\mathbf{Y} = \{Y_1, Y_2\}$.

Note that $\text{PossAn}(\mathbf{Y}, \mathcal{G}_{\mathbf{V}\setminus\mathbf{X}}) = \{V_1, Y_1, Y_2\}$, that is, $\text{PossAn}(\mathbf{Y}, \mathcal{G}_{\mathbf{V}\setminus\mathbf{X}}) = \mathbf{V} \setminus \mathbf{X}$. Then, $\mathbf{V} \setminus \text{PossAn}(\mathbf{Y}, \mathcal{G}_{\mathbf{V}\setminus\mathbf{X}}) = \mathbf{X}$. The strategy of Jaber et al. (2019) would dictate that we can identify the causal effect of $\mathbf{X}$ on $\mathbf{Y}$ by first identifying the causal effect of $\mathbf{X}$ on $\mathbf{V} \setminus \mathbf{X}$ in $\mathcal{G}$. As we have seen in Example 3.8, the causal effect of $\mathbf{X}$ on $\mathbf{V} \setminus \mathbf{X}$ in $\mathcal{G}$ is not identifiable, whereas the causal effect of $\mathbf{X}$ on $\mathbf{Y}$ is identifiable in $\mathcal{G}$. Therefore, the approach of Jaber et al. (2019) is not suitable for general MPDAGs. The above counter example arises as a consequence a partially directed cycle in the MPDAG. Hence, a modified approach of Jaber et al. (2019) may lead to a necessary causal identification algorithm in MPDAGs without partially directed cycles.

A natural question of interest is whether a similar approach to ours can be applied to PAGs. Another topic for future work is developing a complete identification formula for conditional causal effects in MPDAGs.

Acknowledgments

We thank Søren Wengel Mogensen, F. Richard Guo, Vincent Roulet, and the reviewers for their insightful comments and suggestions.

References

Andersson, S. A., Madigan, D., and Perlman, M. D. (1997). A characterization of Markov equivalence classes for acyclic digraphs. *Annals of Statistics*, 25:505–541.

Chickering, D. M. (2002). Learning equivalence classes of Bayesian-network structures. *Journal of Machine Learning Research*, 2:445–498.

Collins, L. M., Murphy, S. A., and Bierman, K. L. (2004). A conceptual framework for adaptive preventive interventions. *Prevention Science*, 5(3):185–196.

Collins, L. M., Murphy, S. A., and Strecher, V. (2007). The multiphase optimization strategy (MOST) and the sequential multiple assignment randomized trial (SMART): new methods for more potent ehealth interventions. *American Journal of Preventive Medicine*, 32(5):S112–S118.

Eigenmann, M., Nandy, P., and Maathuis, M. H. (2017). Structure learning of linear Gaussian structural equation models with weak edges. In *Proceedings of UAI 2017*.

Galles, D. and Pearl, J. (1995). Testing identifiability of causal effects. In *Proceedings of UAI 1995*, pages 185–195.

Hauser, A. and Bühlmann, P. (2012). Characterization and greedy learning of interventional Markov equivalence classes of directed acyclic graphs. *Journal of Maching Learning Research*, 13:2409–2464.

Henckel, L., Perković, E., and Maathuis, M. H. (2019). Graphical criteria for efficient total effect estimation via adjustment in causal linear models. arXiv:1907.02435.

Hernán, M. A. and Robins, J. M. (2020). *Causal Inference: What if.* Boca Raton: Chapman & Hall/CRC.

Hoyer, P. O., Hyvarinen, A., Scheines, R., Spirtes, P. L., Ramsey, J., Lacerda, G., and Shimizu, S. (2008). Causal discovery of linear acyclic models with arbitrary distributions. In *Proceedings of UAI 2008*, pages 282–289.

Huang, Y. and Valtorta, M. (2006). Pearl's calculus of intervention is complete. In *Proceedings of UAI 2006*.

Jaber, A., Zhang, J., and Bareinboim, E. (2018a). Causal identification under Markov equivalence. In *Proceedings of UAI 2018*.

Jaber, A., Zhang, J., and Bareinboim, E. (2018b). A graphical criterion for effect identification in equivalence classes of causal diagrams. In *Proceedings of IJCAI 2018*, pages 5024–5030.

Jaber, A., Zhang, J., and Bareinboim, E. (2019). Causal identification under Markov equivalence: Completeness results. In *Proceedings of ICML 2019*, volume 97, pages 2981–2989.

Kohler-Hausmann, I. (2018). Eddie Murphy and the dangers of counterfactual causal thinking about detecting racial discrimination. *Northwestern University Law Review*, 113.

Lauritzen, S. L., Dawid, A. P., Larsen, B. N., and Leimer, H.-G. (1990). Independence properties of directed Markov fields. *Networks*, 20(5):491–505.

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

Mardia, K. V., Kent, J. T., and Bibby, J. M. (1980). *Multivariate analysis (probability and mathematical statistics)*. Academic Press London.

Meek, C. (1995). Causal inference and causal explanation with background knowledge. In *Proceedings of UAI 1995*, pages 403–410.

Murphy, S. A. (2003). Optimal dynamic treatment regimes. *Journal of the Royal Statistical Society: Series B*, 65(2):331–355.

Nandy, P., Maathuis, M. H., and Richardson, T. S. (2017). Estimating the effect of joint interventions from observational data in sparse high-dimensional settings. *Annals of Statistics*, 45(2):647–674.

Pearl, J. (1993). Comment: Graphical models, causality and intervention. *Statistical Science*, 8:266–269.

Pearl, J. (1995). Causal diagrams for empirical research. *Biometrika*, 82:669–688.

Pearl, J. (2009). *Causality: Models, Reasoning, and Inference*. Cambridge University Press, New York, NY, second edition.

Pearl, J. and Robins, J. M. (1995). Probabilistic evaluation of sequential plans from causal models with hidden variables. In *Proceedings of UAI 1995*, pages 444–453.

Perković, E., Kalisch, M., and Maathuis, M. H. (2017). Interpreting and using CPDAGs with background knowledge. In *Proceedings of UAI 2017*.

Perković, E., Textor, J., Kalisch, M., and Maathuis, M. H. (2015). A complete generalized adjustment criterion. In *Proceedings of UAI 2015*, pages 682–691.

Perković, E., Textor, J., Kalisch, M., and Maathuis, M. H. (2018). Complete graphical characterization and construction of adjustment sets in Markov equivalence classes of ancestral graphs. *Journal of Machine Learning Research*, 18.

Rahmadi, R., Groot, P., Heins, M., Knoop, H., Heskes, T., et al. (2017). Causality on cross-sectional data: stable specification search in constrained structural equation modeling. *Applied Soft Computing*, 52:687–698.

Rahmadi, R., Groot, P., van Rijn, M. H., van den Brand, J. A., Heins, M., Knoop, H., Heskes, T., Initiative, A. D. N., Group, M. S., and consortium, O. (2018). Causality on longitudinal data: Stable specification search in constrained structural equation modeling. *Statistical Methods in Medical Research*, 27(12):3814–3834.

Richardson, T. S., Evans, R. J., Robins, J. M., and Shpitser, I. (2017). Nested Markov properties for acyclic directed mixed graphs. arXiv preprint arXiv:1701.06686.

Richardson, T. S. and Spirtes, P. (2002). Ancestral graph Markov models. *Annals of Statistics*, 30:962–1030.

Robins, J. M. (1986). A new approach to causal inference in mortality studies with a sustained exposure period-application to control of the healthy worker survivor effect. *Mathematical Modelling*, 7:1393–1512.

Rothenhäusler, D., Ernest, J., and Bühlmann, P. (2018). Causal inference in partially linear structural equation models: identifiability and estimation. *Annals of Statistics*, 46:2904–2938.

Scheines, R., Spirtes, P., Glymour, C., Meek, C., and Richardson, T. (1998). The TETRAD project: constraint based aids to causal model specification. *Multivariate Behavioral Research*, 33(1):65–117.

Shpitser, I. and Pearl, J. (2006). Identification of joint interventional distributions in recursive semi-Markovian causal models. In *Proceedings of AAAI 2006*, pages 1219–1226.

Shpitser, I., VanderWeele, T., and Robins, J. M. (2010). On the validity of covariate adjustment for estimating causal effects. In *Proceedings of UAI 2010*, pages 527–536.

Spirtes, P., Glymour, C., and Scheines, R. (2000). *Causation, Prediction, and Search*. MIT Press, Cambridge, MA, second edition.

Taubman, S. L., Robins, J. M., Mittleman, M. A., and Hernán, M. A. (2009). Intervening on risk factors for coronary heart disease: an application of the parametric g-formula. *International Journal of Epidemiology*, 38(6):1599–1611.

Tian, J. and Pearl, J. (2002). A general identification condition for causal effects. In *Proceedings of AAAI 2002*, pages 567–573.

van der Zander, B., Liśkiewicz, M., and Textor, J. (2014). Constructing separators and adjustment sets in ancestral graphs. In *Proceedings of UAI 2014*, pages 907–916.

VanderWeele, T. J. and Robinson, W. R. (2014). On causal interpretation of race in regressions adjusting for confounding and mediating variables. *Epidemiology*, 25(4).

Wang, Y., Solus, L., Yang, K. D., and Uhler, C. (2017). Permutation-based causal inference algorithms with interventions. In *Proceedings of NeurIPS 2017*.

Westreich, D., Cole, S. R., Young, J. G., Palella, F., Tien, P. C., Kingsley, L., Gange, S. J., and Hernán, M. A. (2012). The parametric g-formula to estimate the effect of highly active antiretroviral therapy on incident AIDS or death. *Statistics in Medicine*, 31(18):2000–2009.

Young, J. G., Cain, L. E., Robins, J. M., O'Reilly, E. J., and Hernán, M. A. (2011). Comparative effectiveness of dynamic treatment regimes: an application of the parametric g-formula. *Statistics in Biosciences*, 3(1):119.

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

Zhang, J. (2008b). On the completeness of orientation rules for causal discovery in the presence of latent confounders and selection bias. *Artificial Intelligence*, 172:1873–1896.