

Identifiability from a Combination of Observations and Experiments

Sanghack Lee, Juan D. Correa, Elias Bareinboim

Causal Artificial Intelligence Laboratory
 Department of Computer Science
 Columbia University, USA

{sl4712, j.d.correa}@columbia.edu, eb@cs.columbia.edu

Abstract

We study the problem of causal identification from an arbitrary collection of observational and experimental distributions, and substantive knowledge about the phenomenon under investigation, which usually comes in the form of a causal graph. We call this problem *g-identifiability*, or gID for short. In this paper, we introduce a general strategy to prove non-gID based on *thickets* and *hedgelets*, which leads to a necessary and sufficient graphical condition for the corresponding decision problem. We further develop a procedure for systematically computing the target effect, and prove that it is sound and complete for gID instances. In other words, the failure of the algorithm in returning an expression implies that the target effect is not computable from the available distributions. Finally, as a corollary of these results, we show that do-calculus is complete for the task of g-identifiability.

1 Introduction

One of the main tasks in the empirical sciences and data-driven disciplines is to infer cause and effect relationships from a combination of observations, experiments, and substantive knowledge about the phenomenon under investigation. Causal relations are deemed desirable and valuable for constructing explanations and for contemplating novel interventions that were never experienced before (Pearl 2000; Spirtes, Glymour, and Scheines 2001; Bareinboim and Pearl 2016; Pearl and Mackenzie 2018).

In one line of investigation, this task is formalized through the question of whether the effect that an intervention on a set of variables \mathbf{X} will have on another set of outcome variables \mathbf{Y} , denoted by $P_{\mathbf{x}}(\mathbf{y})$, can be uniquely computed from the probability distribution P over the observed variables \mathbf{V} and a causal diagram \mathcal{G} . This is known as the problem of *identification* (Pearl 1995; 2000; Bareinboim and Pearl 2016), and has received great attention in the literature, starting with a number of sufficient conditions (Spirtes, Glymour, and Scheines 2001; Galles and Pearl 1995; Pearl and Robins 1995), and culminating in a complete graphical and algorithmic characterization (Tian and Pearl 2002;




Figure 1: $P_{x_1, x_2}(y)$ can be identified from $P_{x_1}(\mathbf{V})$ and $P_{x_2}(\mathbf{V})$ in (a) and (b), but not in (c) and (d). Differences among the causal diagrams are highlighted in red.

Shpitser and Pearl 2006; Huang and Valtorta 2006). Despite the generality of such results, it’s the case that in some real-world applications the quantity $P_{\mathbf{x}}(\mathbf{y})$ is not identifiable from the observational data and the causal diagram. On an alternative thread in the literature, causal effects ($P_{\mathbf{x}}(\mathbf{y})$) are obtained directly through controlled experimentation (Fisher 1951). In the biomedical sciences, for instance, considerable resources are spent every year by the FDA, the NIH, and others, in supporting large-scale, systematic, and controlled experimentation, which comes under the rubric of Randomized Controlled Trials. Despite all the inferential power entailed by this approach, there are real-world settings where controlling the variables in \mathbf{X} is not feasible.

In this paper, we note that these two approaches can be seen as extremes in a spectrum of possible research designs, which can be combined to solve natural, albeit non-trivial, causal inference problems. This generalized setting has been investigated in the literature under the rubric of *z-identifiability* (zID, for short) (Bareinboim and Pearl 2012). Unlike zID, gID does not assume the availability of an observational distribution. Furthermore, zID assumed that experiments on every possible subset of $\mathbf{Z} \subset \mathbf{V}$ are available, yet in many practical situations, however, it may be unfeasible to intervene simultaneously on some set of variables, while other set may not be intervened on separately. gID relaxes this assumption by considering an arbitrary combination of observational or experimental distributions over \mathbf{V} in identifying an unconditional causal quantities.

Consider the causal graphs in Fig. 1, where Y represents *cardiovascular disease*, W *blood pressure*, X_1 taking an *an-*

tihypertensive drug, and X_2 the use of an *anti-diabetic* drug. While it's currently understood that diabetes and hypertension do not affect each other (no direct link between them), it's common for patients with type 2 diabetes to be susceptible to hypertension, since both conditions share important confounding factors (graphically encoded through the bidirected arrows) (Ferrannini and Cushman 2012). The goal of the analysis is to assess the effect of prescribing a treatment for both conditions on the risk of developing cardiovascular diseases, $P_{x_1, x_2}(y)$. There are two RCTs that separately control for X_1 and X_2 , which means that $P_{x_1}(\mathbf{V})$ and $P_{x_2}(\mathbf{V})$ are available. It turns out that for the models in Figs. 1a, b, $P_{x_1, x_2}(y) = \sum_w P_{x_2}(y|w)P_{x_1}(w)$, which means that the experimental studies suffice to identify the joint effect. The same effect is not identifiable in Figs. 1c, d.

In this extended abstract, we summarize our approach in (Lee, Correa, and Bareinboim 2019) where the contributions in the paper are as follows: (i) We prove a necessary and sufficient graphical condition for gID, which follows from two new graphical constructs called *hedglets* and *thickets*. These structures constitute flexible and general building blocks that are helpful to understand and characterize general identification problems; (ii) Leveraging these results, we develop a sound and complete algorithm that returns any expression derivable from an arbitrary collection of observations and experiments. As a corollary, we prove that do-calculus is complete for g-identification.

2 Preliminaries

We denote variables by capital letters, X , and values by small letters, x . Bold letters, \mathbf{X} or \mathbf{x} , represent sets of variables or values. The domain of a variable X is denoted by \mathfrak{X}_X . Two values \mathbf{x} and \mathbf{z} are said to be consistent if they share the common values for $\mathbf{X} \cap \mathbf{Z}$. We also denote by $\mathbf{x} \setminus \mathbf{Z}$ the value of $\mathbf{X} \setminus \mathbf{Z}$ consistent with \mathbf{x} . We assume that domain of every variable is finite.

Our analysis heavily relies on causal graphs, which we often assign a calligraphic letter, e.g., \mathcal{G} , \mathcal{F} , or \mathcal{H} . We denote by $\mathbf{V}(\mathcal{H})$ the set of vertices (i.e., variables) in a graph \mathcal{H} . A vertex-induced subgraph is denoted by brackets, e.g., $\mathcal{G}[\mathbf{W}]$, which includes \mathbf{W} and the edges among its elements. We define $\mathcal{G} \setminus \mathbf{X}$ as $\mathcal{G}[\mathbf{V}(\mathcal{G}) \setminus \mathbf{X}]$. A root set of a graph is a set of variables that does not have outgoing edges. We use kinship notation for graphical relationships such as parents, descendants, and ancestors of a set of variables. For example, the set of parents of \mathbf{X} in \mathcal{G} is denoted by $pa(\mathbf{X})_{\mathcal{G}} := \bigcup_{X \in \mathbf{X}} pa(X)_{\mathcal{G}}$. Similarly, we define *de*, and *an*. Written as *Pa*, *De*, and *An* (i.e., capitalized), the argument is included as well, e.g., $De(\mathbf{X})_{\mathcal{G}} := de(\mathbf{X})_{\mathcal{G}} \cup \mathbf{X}$.

We use Structural Causal Models (SCMs) (Pearl 2000) as our basic semantical framework. A SCM \mathcal{M} is a 4-tuple $\langle \mathbf{U}, \mathbf{V}, \mathbf{F}, P(\mathbf{U}) \rangle$, where \mathbf{U} is a set of exogenous variables; \mathbf{V} is a set of endogenous variables; \mathbf{F} is a set of functions $\{f_V\}_{V \in \mathbf{V}}$, which determines the value of a variable, e.g., $v \leftarrow f_V(\mathbf{pa}_V, \mathbf{u}^V)$ is a function with $\mathbf{PA}^V \subseteq \mathbf{V} \setminus \{V\}$ and $\mathbf{U}^V \subseteq \mathbf{U}$; and $P(\mathbf{U})$ is a joint probability distribution over \mathbf{U} . A SCM \mathcal{M} induces a causal graph \mathcal{G} (also called a semi-Markovian graph) where \mathbf{V} is a set of vertices, directed

edges are formed satisfying $\mathbf{PA}^V = pa(V)_{\mathcal{G}}$, and each bidirected edge corresponds to an unobserved confounder between two variables, that is, $V_i \leftrightarrow V_j$ if $\mathbf{U}^i \cap \mathbf{U}^j \neq \emptyset$. Interventions are defined through an operator called $do(\mathbf{X} = \mathbf{x})$, which sets the intervened variables \mathbf{X} to specific values $\mathbf{x} \in \mathfrak{X}_{\mathbf{X}}$. Given a model \mathcal{M} , an intervention $do(\mathbf{X} = \mathbf{x})$ induces a submodel $\mathcal{M}_{\mathbf{x}}$, where f_X of \mathbf{F} is replaced by $f_X = x$ for every $X \in \mathbf{X}$ where x is consistent with \mathbf{x} . This submodel $\mathcal{M}_{\mathbf{x}}$ induces a causal graph $\mathcal{G}_{\overline{\mathbf{X}}}$, which reads as \mathcal{G} with edges onto any of \mathbf{X} removed.

We now revisit some key notions for deciding identifiability developed in the context of non-experimental settings. First, we define a special type of cluster of variables called *confounded components* (Tian and Pearl 2002).

Definition 1 (C-component). Let \mathcal{G} be a semi-Markovian graph such that a subset of its bidirected arcs forms a spanning tree over all vertices in \mathcal{G} . Then \mathcal{G} is a c-component.

Given a semi-Markovian graph \mathcal{G} over a set of variables \mathbf{V} , there exists a unique partition such that each subgraph is a maximal c-component. We denote by $\mathcal{C}(\mathcal{G})$ the set of c-components that partitions the vertices in \mathcal{G} . Armed with this definition, we build towards the hedge with the following notion adapted from (Shpitser and Pearl 2006).

Definition 2 (C-forest). A semi-Markovian graph \mathcal{G} with root set \mathbf{R} is said to be an \mathbf{R} -rooted c-forest if \mathcal{G} is a c-component with a minimal number of edges.

The minimality with respect to the number of edges guarantees that every vertex not in the root set of a c-forest has one child and its bidirected edges form exactly a spanning tree. We are now ready to define a hedge as follows.

Definition 3 (Hedge). A hedge is a pair of \mathbf{R} -rooted c-forests $\langle \mathcal{F}, \mathcal{F}' \rangle$ such that $\mathcal{F}' \subseteq \mathcal{F}$.

To realize the connection between definitions, note that given disjoint sets $\mathbf{X}, \mathbf{Y} \subset \mathbf{V}$, if $\mathbf{R} \subseteq An(\mathbf{Y})_{\mathcal{G}_{\mathbf{X}}}$, $\mathcal{F} \cap \mathbf{X} \neq \emptyset$, and $\mathcal{F}' \cap \mathbf{X} = \emptyset$, Def. 3 reduces to the original definition. The existence of such structure precludes the identifiability of $P_{\mathbf{x}}(\mathbf{y})$ from $P(\mathbf{V})$ (Shpitser and Pearl 2006). In the new theoretical treatment pursued in (Lee, Correa, and Bareinboim 2019), a hedge as a graphical structure itself is separated from its use as a witness of the non-identifiability of a specific causal distribution. We say that the a hedge structure $\langle \mathcal{F}, \mathcal{F}' \rangle$ is *formed for* $P_{\mathbf{x}}(\mathbf{y})$ in \mathcal{G} whenever referring to the original semantics, i.e., regarding the non-identifiability of $P_{\mathbf{x}}(\mathbf{y})$. Further, we'll distinguish two parts of a hedge $\langle \mathcal{F}, \mathcal{F}' \rangle$: the 'top' part, denoted by $\mathcal{F}'' = \mathcal{F} \setminus \mathbf{V}(\mathcal{F}')$, and the 'bottom' part, which is \mathcal{F}' . When the top is empty (i.e., $\mathcal{F} = \mathcal{F}'$), we will call this hedge *degenerate*.

3 G-Identifiability

We first introduce a new task that formalizes and generalizes the identifiability and z-identifiability settings by allowing a more flexible input consisting of any combination of observational and experimental distributions.

Definition 4 (g-Identifiability). Let \mathbf{X}, \mathbf{Y} be disjoint sets of variables, $\mathbf{Z} = \{\mathbf{Z}_i\}_{i=1}^m$ be a collection of sets of variables, and let \mathcal{G} be a causal diagram. $P_{\mathbf{x}}(\mathbf{y})$ is said to be




Figure 2: Hedgelet decomposition of hedges and a thicket (color coded in blue and red with purple for shared elements). Each of (a) and (b) is a hedge formed for $P_{\mathbf{r}}(\mathbf{x})$ or a thicket with respect to $\mathbb{Z} = \{\{X_1\}, \{X_2\}\}$ while (c) is not a hedge but a thicket.

g-identifiable from \mathbb{Z} in \mathcal{G} , if $P_{\mathbf{x}}(\mathbf{y})$ is uniquely computable from distributions $\{P(\mathbf{V} \mid do(\mathbf{z}))\}_{\mathbf{z} \in \mathbb{Z}, \mathbf{z} \in \mathfrak{X}_{\mathbf{Z}}}$ in any causal model which induces \mathcal{G} .

A traditional and pervasive assumption made throughout the identification literature is that a probability distribution describing the natural state of the system is available, that is, $P(\mathbf{V})$. In the setting defined above, such distribution is not a priori required unless the empty set is explicitly included in \mathbb{Z} . The following statement can be shown based on the definition of g-identifiability:

Lemma 1. *Let \mathbf{X}, \mathbf{Y} be disjoint sets of variables, $\mathbb{Z} = \{\mathbf{Z}_i\}_{i=1}^m$ be a collection of sets of variables, and let \mathcal{G} be a causal diagram. $P_{\mathbf{x}}(\mathbf{y})$ is not g-identifiable from \mathbb{Z} in \mathcal{G} if there exist two causal models \mathcal{M}_1 and \mathcal{M}_2 compatible with \mathcal{G} such that $P_{\mathbf{z}}^1(\mathbf{v}) = P_{\mathbf{z}}^2(\mathbf{v})$ for all $\mathbf{Z} \in \mathbb{Z}, \mathbf{z} \in \mathfrak{X}_{\mathbf{Z}}$, but $P_{\mathbf{x}}^1(\mathbf{y}) \neq P_{\mathbf{x}}^2(\mathbf{y})$.*

Even though this statement formally characterizes non-g-identifiability of a certain data collection, it does not provide any insight on how to determine if such pair of models exists, or how to construct them when a given instance is not g-identifiable.

3.1 Hedgelets and Thickets

When considering multiple experimental distributions as inputs, a graphical structure that might be able to witness the non-g-identifiability has to account for all such experiments. To deal with the complexity added by a broader input, we introduce *hedgelets*, a unique decomposition of a hedge. We define how to obtain the set of hedgelets associated with any given hedge.

Definition 5 (hedgelet decomposition). The hedgelet decomposition of a hedge $\langle \mathcal{F}, \mathcal{F}' \rangle$ is the collection of hedgelets $\{\mathcal{F}(\mathbf{W})\}_{\mathbf{W} \in \mathcal{C}(\mathcal{F}'')}$ where each hedgelet $\mathcal{F}(\mathbf{W})$ is a subgraph of \mathcal{F} made of (i) $\mathcal{F}[\mathbf{V}(\mathcal{F}') \cup \mathbf{W}]$ and (ii) $\mathcal{F}[De(\mathbf{W})_{\mathcal{F}}]$ without bidirected edges.

Let $\mathbb{H}_{\mathcal{F}} = \{\mathcal{F}(\mathbf{W})\}_{\mathbf{W} \in \mathcal{C}(\mathcal{F}'')}$ be the set of *hedgelets* of $\langle \mathcal{F}, \mathcal{F}' \rangle$. For a degenerate hedge, $\mathbb{H}_{\mathcal{F}}$ contains a single hedgelet $\mathcal{F}(\emptyset) = \mathcal{F}$, which we call a *degenerate* hedgelet. Given a non-degenerate hedge, for every hedgelet \mathcal{H} in it, there exists at least one directed edge, and exactly one bidirected edge between \mathcal{H}' and \mathcal{H}'' by definition. For a simple example, see Fig. 2a, a hedge $\langle \mathcal{F}, \mathcal{F}' \rangle$ for $P_{\mathbf{x}}(\mathbf{r})$. This hedge can be decomposed into two hedgelets $\mathcal{F}(\{X_1\})$ in blue

(i.e., $\mathcal{G}[\{X_1, R\}]$ and $\mathcal{F}(\{X_2\})$ in red (i.e., $\mathcal{G}[\{X_2, R\}]$). Fig. 2b is a hedge $\langle \mathcal{F}, \mathcal{F}' \rangle$ for $P_{\mathbf{x}}(\mathbf{r})$, which can be similarly decomposed into two hedgelets $\mathcal{F}(\{X_1\})$ and $\mathcal{F}(\{X_2\})$.

Now, we will describe a graphical structure relative to the available input distributions entailed by \mathbb{Z} , that precludes the g-identifiability of a causal effect $P_{\mathbf{x}}(\mathbf{y})$ in \mathcal{G} .

Definition 6 (Thicket). Let \mathbf{R} be a non-empty set of variables and \mathbb{Z} be a collection of sets of variables in \mathcal{G} . A thicket $\mathcal{T} \subseteq \mathcal{G}$ is an \mathbf{R} -rooted c-component consisting of a minimal c-component over \mathbf{R} and hedges

$$\mathbb{F}_{\mathcal{T}} = \{\langle \mathcal{F}_{\mathbf{Z}}, \mathcal{T}[\mathbf{R}] \rangle \mid \mathcal{F}_{\mathbf{Z}} \subseteq \mathcal{G} \setminus \mathbf{Z}, \mathbf{Z} \cap \mathbf{R} = \emptyset\}_{\mathbf{Z} \in \mathbb{Z}}.$$

Let \mathbf{X}, \mathbf{Y} be disjoint sets of variables in \mathcal{G} . A thicket \mathcal{T} is said to be formed for $P_{\mathbf{x}}(\mathbf{y})$ in \mathcal{G} with respect to \mathbb{Z} if $\mathbf{R} \subseteq An(\mathbf{Y})_{\mathcal{G}_{\mathbb{Z}}}$ and every hedgelet of each hedge $\langle \mathcal{F}_{\mathbf{Z}}, \mathcal{T}[\mathbf{R}] \rangle$ intersects with \mathbf{X} .

If $\mathbf{Z} \cap \mathbf{R} = \emptyset$ for some $\mathbf{Z} \in \mathbb{Z}$, a thicket can be viewed as a superimposition of hedges where each of them comes from a subgraph of the thicket obtained by excluding an available experiment that was not performed on any of \mathbf{R} . Otherwise if $\mathbf{Z} \cap \mathbf{R} \neq \emptyset$ for every $\mathbf{Z} \in \mathbb{Z}$, that is, every experiment disrupts \mathbf{R} , \mathcal{T} will simply be a spanning tree over \mathbf{R} with bidirected arcs. Whenever this is the case, we call this thicket *degenerate*, which consists of a degenerate hedge with a single degenerate hedgelet. To illustrate see Figs. 2a to 2c. Each causal diagram is a thicket for $P_{\mathbf{x}}(\mathbf{r})$ with respect to $\mathbb{Z} = \{\{X_1\}, \{X_2\}\}$ with two hedges in red and blue where each hedge itself is a hedgelet. We refer readers to the original paper for more examples.

3.2 A Graphical Condition for Non-g-identifiability

We consider constructing two models agreeing in the available distributions but yielding a different result for the causal effect (Lemma 1). The key idea is to parametrize the models such that the distribution of the variables in \mathbf{R} remains the same as long as there is one non-intervened hedgelet. From the definition of Thicket, it can be inferred that no intervention on $\mathbf{Z} \in \mathbb{Z}$ affects all hedgelets, hence they are all affected by the intervention on \mathbf{X} . Nevertheless, this is difficult to achieve for each configuration \mathbf{u} , since each of $R \in \mathbf{R}$ is determined independently (i.e., structural). Instead, our models behave differently with respect to a specific value \mathbf{u} but induce the same distribution over \mathbf{R} when averaged over \mathbf{U} as long as there exists a non-intervened hedgelet. We proved the following theorem by implementing the idea to construct two such models:

Theorem 1. *If there exists a thicket \mathcal{T} for $P_{\mathbf{x}}(\mathbf{y})$ in \mathcal{G} with respect to \mathbb{Z} , then, $P_{\mathbf{x}}(\mathbf{y})$ is not g-identifiable in \mathcal{G} .*

4 A Sound and Complete Algorithm for g-Identifiability

Building on the graphical characterization of non-gID, in this section, we develop an algorithm for g-identifiability called gID (Alg. 1). For a given causal query, gID determines whether it's g-identifiable, and if so, it outputs a formula expressing the target effect in terms of the available

Algorithm 1 gID: a complete algorithm for g-identifiability

```
1: function gID( $\mathbf{y}, \mathbf{x}, \mathcal{G}, \mathbb{Z}$ )
   Input:  $\mathbf{y}, \mathbf{x}$ : value assignments,  $\mathcal{G}$ : causal diagram,  $\mathbb{Z}$ : a collection of available experiments
   Output: an estimand computing  $P_{\mathbf{x}}(\mathbf{y})$  with  $\{P_{\mathbf{z}}\}_{\mathbf{z} \in \mathbb{Z}, \mathbf{z} \in \mathbf{x}_{\mathbf{z}}}$ .
2: if  $\exists \mathbf{z} \in \mathbb{Z} \mathbf{X} = \mathbf{Z} \cap \mathbf{V}$  then
   return  $P_{\mathbf{z} \setminus \mathbf{V}, \mathbf{x}}(\mathbf{y})$ 
3: if  $\mathbf{V} \neq \text{An}(\mathbf{Y})_{\mathcal{G}}$  then
   return gID( $\mathbf{y}, \mathbf{x} \cap \text{An}(\mathbf{Y})_{\mathcal{G}}, \mathcal{G}[\text{An}(\mathbf{Y})_{\mathcal{G}}], \mathbb{Z}$ )
4: if  $(\mathbf{W} \leftarrow (\mathbf{V} \setminus \mathbf{X}) \setminus \text{An}(\mathbf{Y})_{\mathcal{G}_{\overline{\mathbf{x}}}}) \neq \emptyset$  then
   return gID( $\mathbf{y}, \mathbf{x} \cup \mathbf{w}, \mathcal{G}, \mathbb{Z}$ )
5: if  $|\mathcal{C}(\mathcal{G} \setminus \mathbf{X})| > 1$  then
   return  $\sum_{\mathbf{v} \setminus (\mathbf{y} \cup \mathbf{x})} \prod_{\mathbf{s} \in \mathcal{C}(\mathcal{G} \setminus \mathbf{X})} \text{gID}(\mathbf{s}, \mathbf{v} \setminus \mathbf{s}, \mathcal{G}, \mathbb{Z})$ 
6: for  $\mathbf{Z} \in \mathbb{Z}$  such that  $\mathbf{Z} \cap \mathbf{V} \subseteq \mathbf{X}$  do
   return SUB-ID( $\mathbf{y}, \mathbf{x} \setminus \mathbf{Z}, P_{(\mathbf{z} \setminus \mathbf{V}), \mathbf{x} \cap \mathbf{Z}}, \mathcal{G} \setminus (\mathbf{Z} \cap \mathbf{X})$ ) if not none
7: throw fail
```

distributions. The design of gID shares some principles established by previous identifiability algorithms. Still, in our case, the identification process is decomposed into two parts: pre- and post-activation of an available distribution, where SUB-ID¹ takes care of a (classic) identification task for each factored query with a fixed distribution treated as observational, relative to the call-specific graph.

The algorithm takes a query $P_{\mathbf{x}}(\mathbf{y})$, the causal graph \mathcal{G} , and available experiments \mathbb{Z} as inputs. During the process, the query and the causal graph may be transformed when necessary, and broken down into smaller sub-problems. Accordingly, the parameters \mathbf{y} , \mathbf{x} , and \mathcal{G} are local to each call, while \mathbb{Z} is preserved throughout recursive calls. The given \mathcal{G} is modified only through Line 3, since experiments on variables that are not ancestors of \mathbf{Y} have no effect on it, we only need to pay attention to experiments on ancestors of \mathbf{Y} . Line 2 utilizes any matching experiment whenever possible. As mentioned above, \mathbf{Z} outside the current scope can be of any value. Lines 4 and 5 modify and factorize the given query, respectively. At Line 6, given a factorized query, the algorithm examines whether an available distribution might be useful to estimate it, and delegates the identification to a subroutine. In the original paper, we proved its completeness for g-identifiability by showing the existence of a thicket whenever the algorithm fails on a given query. Further, as a corollary, we showed that do-calculus is complete for the same problem by translating the algorithm to a series of application of do-calculus.

5 Conclusion

We studied the identification of causal effects from arbitrary combinations of observational and experimental distributions. This problem has been called g-identifiability, or gID for short. We developed a general algorithm for solving gID and proved its completeness. We introduced new machinery to better understand and more precisely character-

¹It is a shortened ID algorithm (Shpitser and Pearl 2006), which returns none instead of throwing failure when a delegated query is not identified.

ize non-trivial forbidden structures that preclude gID, which can be seen as instances of hedgelets and thickets. Finally, as a corollary of these results, we proved that do-calculus is complete for the task of g-identifiability.

References

- Bareinboim, E., and Pearl, J. 2012. Causal inference by surrogate experiments: z -identifiability. In *Proceedings of the Twenty-Eighth Conference on Uncertainty in Artificial Intelligence*, 113–120. Corvallis, OR: AUAI Press.
- Bareinboim, E., and Pearl, J. 2016. Causal inference and the data-fusion problem. *Proceedings of the National Academy of Sciences* 113:7345–7352.
- Ferrannini, E., and Cushman, W. C. 2012. Diabetes and hypertension: the bad companions. *The Lancet* 380(9841):601–610.
- Fisher, R. 1951. *The Design of Experiments*. Edinburgh: Oliver and Boyd, 6th edition.
- Galles, D., and Pearl, J. 1995. Testing identifiability of causal effects. In *Proceedings of the Eleventh Conference on Uncertainty in Artificial Intelligence*. San Francisco: Morgan Kaufmann. 185–195.
- Huang, Y., and Valorta, M. 2006. Identifiability in causal bayesian networks: A sound and complete algorithm. In *Proceedings of the Twenty-First National Conference on Artificial Intelligence*. Menlo Park, CA: AAAI Press. 1149–1156.
- Lee, S.; Correa, J. D.; and Bareinboim, E. 2019. General identifiability with arbitrary surrogate experiments. In *Proceedings of the 35th Conference on Uncertainty in Artificial Intelligence*.
- Pearl, J., and Mackenzie, D. 2018. *The Book of Why: The New Science of Cause and Effect*. Basic Books.
- Pearl, J., and Robins, J. 1995. Probabilistic evaluation of sequential plans from causal models with hidden variables. In Besnard, P., and Hanks, S., eds., *Proceedings of the Eleventh Conference on Uncertainty in Artificial Intelligence*. San Francisco: Morgan Kaufmann. 444–453.
- Pearl, J. 1995. Causal diagrams for empirical research. *Biometrika* 82(4):669–688.
- Pearl, J. 2000. *Causality: Models, Reasoning, and Inference*. New York: Cambridge University Press.
- Shpitser, I., and Pearl, J. 2006. Identification of joint interventional distributions in recursive semi-Markovian causal models. In *Proceedings of The Twenty-First National Conference on Artificial Intelligence*, 1219–1226. AAAI Press.
- Spirites, P.; Glymour, C.; and Scheines, R. 2001. *Causation, Prediction, and Search*. Cambridge, MA: MIT Press, 2nd edition.
- Tian, J., and Pearl, J. 2002. A general identification condition for causal effects. In *Proceedings of the Eighteenth National Conference on Artificial Intelligence*, 567–573. Menlo Park, CA: AAAI Press/The MIT Press.