
Nested Counterfactual Identification from Arbitrary Surrogate Experiments

Juan D. Correa
Columbia University
jdcorra@cs.columbia.edu

Sanghack Lee
Seoul National University
sanghack@snu.ac.kr

Elias Bareinboim
Columbia University
eb@cs.columbia.edu

Abstract

The *Ladder of Causation* describes three qualitatively different types of activities an agent may be interested in engaging in, namely, seeing (observational), doing (interventional), and imagining (counterfactual) (Pearl and Mackenzie, 2018). The inferential challenge imposed by the causal hierarchy is that data is collected by an agent observing or intervening in a system (layers 1 and 2), while its goal may be to understand what would have happened had it taken a different course of action, contrary to what factually ended up happening (layer 3). While there exists a solid understanding of the conditions under which cross-layer inferences are allowed from observations to interventions, the results are somewhat scarcer when targeting counterfactual quantities. In this paper, we study the identification of nested counterfactuals from an arbitrary combination of observations and experiments. Specifically, building on a more explicit definition of nested counterfactuals, we prove the counterfactual unnesting theorem (CUT), which allows one to map arbitrary nested counterfactuals to unnested ones. For instance, applications in mediation and fairness analysis usually evoke notions of direct, indirect, and spurious effects, which naturally require nesting. Second, we introduce a sufficient and necessary graphical condition for counterfactual identification from an arbitrary combination of observational and experimental distributions. Lastly, we develop an efficient and complete algorithm for identifying nested counterfactuals; failure of the algorithm returning an expression for a query implies it is not identifiable.

1 Introduction

Counterfactuals provide the basis for notions pervasive throughout human affairs, such as credit assignment, blame and responsibility, and regret. One of the most powerful constructs in human reasoning —“what if?” questions— evokes hypothetical conditions usually contradicting the factual evidence. Judgment and understanding of critical situations found from medicine to psychology to business involve counterfactual reasoning, e.g.: “Joe received the treatment and died, would he be alive had he not received it?,” “Had the candidate been male instead of female, would the decision from the admissions committee be more favorable?,” or “Would the profit this quarter remain within 5% of its value had we increased the price by 2%?”. By and large, counterfactuals are key ingredients that go in the construction of explanations about why things happened as they did [13, 15].

The structural interpretation of causality provides proper semantics for representing counterfactuals [13, Ch. 7]. Specifically, each structural causal model (SCM) \mathcal{M} induces a collection of distributions related to the activities of seeing (called observational), doing (interventional), and imagining (counterfactual), which together were called the *ladder of causation* [15, 1]. The ladder is a containment hierarchy; each type of distribution can be put in increasingly refined layers: observational content goes in layer 1; experimental in layer 2; counterfactual in layer 3 (Fig. 1(a)).




Figure 1: Representation of the ladder of causation and an example of a causal diagram.

It is understood that if we have all the information in the world about layer 1, there are still questions about layers 2 and 3 that are unanswerable, or technically undetermined; further, if we have data from layers 1 and 2, there are still questions in the world about layer 3 that are underdetermined [15, 13, 1].

The inferential challenge in these settings arises since the generating model \mathcal{M} is not fully observed, nor data from all of the layers are necessarily available, perhaps due to the cost or the infeasibility of performing certain interventions. One common task found in the literature is to determine the effect of an intervention of a variable X on an outcome Y , say $P(Y|do(X))$ (layer 2), using data from observations $P(\mathbf{V})$ (layer 1), where \mathbf{V} is the set of observed variables, and possibly other interventions, e.g., $P(\mathbf{V}|do(Z))$. Also, qualitative assumptions about the system are usually articulated in the form of a causal diagram \mathcal{G} . This setting has been studied in the literature under the rubric of non-parametric identification from a combination of observations and experiments. Multiple solutions exist, including Pearl’s celebrated do-calculus [12], and other increasingly refined solutions that are computationally efficient, sufficient, and necessary [21, 5, 16, 22, 17, 7, 2, 9].

There is growing literature on identification with cross-layer inferences from data in layers 1 and 2 to quantities in layer 3. For example, a data scientist may be interested in evaluating the effect of a treatment on the group of subjects that receive it instead of those randomly assigned to treatment. This measure is known as the *effect of treatment on the treated* [6, 13] and there exists a graphical condition for mapping it to a (layer 2) causal effect [19]. Further, there are also results on the identification of *path-specific effects*, which correspond to counterfactuals that isolate specific paths in the graph [14]. In particular, [20] provides a complete algorithm for identification from observational data, and [24] gives identification conditions from observational and experimental data in specific canonical models. Moreover, [18] studied the identification of arbitrary (non-nested) counterfactuals under the assumption that data from experiments in every variable is available. Yet, the problem of identifying such quantities from a subset of the space of all experiments remains open.

For concreteness, consider a counterfactual called *direct effect* in the context of the causal diagram in Fig. 1(b). This quantity quantifies the sensitivity of a variable Y to changes in another variable X while all other factors in the analysis remain fixed. Suppose X is level of exercise, M cholesterol levels, and Y cardiovascular disease. Exercising can improve cholesterol levels, which in turn affect the chances of developing cardiovascular disease. An interesting question is how much exercise prevents the disease by means other than regulating cholesterol. In counterfactual notation, this is to compare Y_{x, M_x} and Y_{x', M_x} where x and x' are different values. The first quantity represents the value of Y when $X=x$ and M varies accordingly. The second expression is the value Y attains if X is held constant at x' while M still follows $X=x$. The difference $E[Y_{x', M_x} - Y_{x, M_x}]$ —known as the natural direct effect (NDE)—is non-zero if there is some direct effect of X on Y . In this instance, this nested counterfactual is identifiable only if observational data and experiments on X are available.

After all, there is no general identification method for this particular counterfactual family (which also includes indirect and spurious effects) and, more broadly, other arbitrary nested counterfactuals that are well-defined in layer 3. Our goal is to understand the non-parametric identification of arbitrary nested and conditional counterfactuals when the input consists of any combination of observational and interventional distributions, whatever is available for the data scientist. More specifically, our contributions are as follows.

1. We look at nested counterfactuals from an SCM perspective and introduce machinery that supports counterfactual reasoning. In particular, we prove the counterfactual unnesting theorem (CUT), which allows one to map any nested counterfactual to an unnested one (Section 2).
2. Building on this new machinery, we derive sufficient and necessary graphical conditions and an algorithm to determine the identifiability of marginal nested counterfactuals from an arbitrary combination of observational and experimental distributions (Section 3).

3. We give a reduction from conditional counterfactuals to marginal ones, and use it to derive a complete algorithm for their identification (Section 4).

See the supplemental material for full proofs of the results in the paper.

1.1 Preliminaries

We denote variables by capital letters, X , and values by small letters, x . Bold letters, \mathbf{X} represent sets of variables and \mathbf{x} sets of 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} and by $\mathbf{x} \cap \mathbf{Z}$ the subset of \mathbf{x} corresponding to variables in \mathbf{Z} . We assume the domain of every variable is finite.

Our analysis relies on causal graphs, which we often assign a calligraphic letter, e.g., \mathcal{G} , \mathcal{H} , etc. We denote by $\mathbf{V}(\mathcal{H})$ the set of vertices (i.e., variables) in a graph \mathcal{H} . Given a graph \mathcal{G} , $\mathcal{G}_{\overline{\mathbf{W}}\mathbf{X}}$ is the result of removing edges coming into variables in \mathbf{W} and going out from variables in \mathbf{X} . $\mathcal{G}[\mathbf{W}]$ denotes a vertex-induced subgraph, which includes \mathbf{W} and the edges among its elements. We use kinship notation for graphical relationships such as parents, children, descendants, and ancestors of a set of variables. For example, the set of parents of \mathbf{X} in \mathcal{G} is denoted by $Pa(\mathbf{X})_{\mathcal{G}} := \mathbf{X} \cup \bigcup_{X \in \mathbf{X}} Pa(X)_{\mathcal{G}}$. Similarly, we define $Ch()$, $De()$, and $An()$.

To articulate and formalize counterfactual questions, we require a framework that allows us to reason about events corresponding to different *alternative worlds* simultaneously. Accordingly, we employ the Structural Causal Model (SCM) paradigm [13, Ch. 7]. An SCM \mathcal{M} is a 4-tuple $\langle \mathbf{U}, \mathbf{V}, \mathcal{F}, P(\mathbf{u}) \rangle$, where \mathbf{U} is a set of exogenous (latent) variables; \mathbf{V} is a set of endogenous (observable) variables; \mathcal{F} is a collection of functions such that each variable $V_i \in \mathbf{V}$ is determined by a function $f_i \in \mathcal{F}$. Each f_i is a mapping from a set of exogenous variables $\mathbf{U}_i \subseteq \mathbf{U}$ and a set of endogenous variables $\mathbf{Pa}_i \subseteq \mathbf{V} \setminus \{V_i\}$ to the domain of V_i . The uncertainty is encoded through a probability distribution over the exogenous variables, $P(\mathbf{U})$. A SCM \mathcal{M} induces a *causal diagram* \mathcal{G} where every \mathbf{V} is the set of vertices, there is a directed edge $(V_j \rightarrow V_i)$ for every $V_i \in \mathbf{V}$ and $V_j \in \mathbf{Pa}_i$, and a bidirected edge $(V_i \leftrightarrow V_j)$ for every pair $V_i, V_j \in \mathbf{V}$ such that $U_i \cap U_j \neq \emptyset$ (V_i and V_j have a common exogenous parent).

We assume that the underlying model is recursive. That is, there are no cyclic dependencies among the variables. Equivalently, that is to say, that the corresponding causal diagram is acyclic.

The set \mathbf{V} can be partitioned into subsets called *c-components* [23] according to a diagram \mathcal{G} such that two variables belong to the same c-component if they are connected in \mathcal{G} by a path made entirely of bidirected edges.

2 SCMs and Nested Counterfactuals

Intervening on a system represented by an SCM \mathcal{M} results in a new model differing from \mathcal{M} only on the mechanisms associated with the intervened variables [11, 3, 4]. If the intervention consists on fixing the value of a variable X to a constant $x \in \mathfrak{X}_X$, it induces a *submodel*, denoted as \mathcal{M}_x [13, Def. 7.1.2]. To formally study nested counterfactuals, we extend this notion to account for models derived from interventions that replace functions from the original SCM with other, not necessarily constant, functions.

Definition 1 (Derived Model). Let \mathcal{M} be an SCM, $\widehat{\mathbf{U}} \subseteq \mathbf{U}$, $X \in \mathbf{V}$, and $\widehat{X} : \widehat{\mathbf{U}} \rightarrow \mathfrak{X}_X$ a function. Then, $\mathcal{M}_{\widehat{X}}$, called the *derived model* of \mathcal{M} according to \widehat{X} , is identical to \mathcal{M} , except that the function f_X is replaced with a function \widehat{f}_X identical to \widehat{X} .

This definition is easily extendable to models derived from an intervention on a set \mathbf{X} instead of a singleton. When $\widehat{\mathbf{X}}$ is a collection of functions $\{\widehat{X} : \widehat{\mathbf{U}}_X \rightarrow \mathfrak{X}_X\}_{X \in \mathbf{X}}$, the derived model $\mathcal{M}_{\widehat{\mathbf{X}}}$ is obtained by replacing each f_X with \widehat{X} for $X \in \mathbf{X}$. Next, we discuss the concept of *potential response* [13, Def. 7.4.1] with respect to derived models.

Definition 2 (Potential Response). Let $\mathbf{X}, \mathbf{Y} \subseteq \mathbf{V}$ be subsets of observable variables, let \mathbf{u} be a unit, and let $\widehat{\mathbf{X}}(\mathbf{u})$ be a set of functions from $\widehat{\mathbf{U}}_X \rightarrow \mathfrak{X}_X$, for $X \in \mathbf{X}$ where $\widehat{\mathbf{U}}_X \subseteq \mathbf{U}$. Then, $\mathbf{Y}_{\mathbf{X}=\widehat{\mathbf{X}}}(\mathbf{u})$

(or $Y_{\widehat{\mathbf{X}}}(\mathbf{u})$, for short) is called the *potential response* of \mathbf{Y} to $\mathbf{X} = \widehat{\mathbf{X}}$, and is defined as the solution of \mathbf{Y} , for a particular \mathbf{u} , in the derived model $\mathcal{M}_{\widehat{\mathbf{X}}}$.

A potential response $Y_{\widehat{\mathbf{X}}}(\mathbf{u})$ describes the value that variable Y would attain for a unit (or individual) \mathbf{u} if the intervention $\widehat{\mathbf{X}}$ is performed. This concept is tightly related to that of *potential outcome*, but the former explicitly allows for interventions that do not necessarily fix the variables in \mathbf{X} to a constant value. Averaging over the space of \mathbf{U} , a potential response $Y_{\widehat{\mathbf{X}}}(\mathbf{u})$ induces a random variable that we will denote simply as $Y_{\widehat{\mathbf{X}}}$. If the intervention replaces a function f_X with a potential response of X in \mathcal{M} , we say the intervention is *natural*.

When variables are enumerated as W_1, W_2, \dots , we may add square brackets around the part of the subscript denoting interventions. We use \mathbf{W}_* to denote sets of arbitrary counterfactual variables. Let $\mathbf{W}_* = \{W_{1[\widehat{\mathbf{T}}_1]}, W_{2[\widehat{\mathbf{T}}_2]}, \dots\}$ represent a set of counterfactual variables such that $W_i \in \mathbf{V}$ and $\mathbf{T}_i \subseteq \mathbf{V}$ for $i = 1, \dots, l$. Define $\mathbf{V}(\mathbf{W}_*) = \{W \in \mathbf{V} \mid W_{\widehat{\mathbf{T}}} \in \mathbf{W}_*\}$, that is, the set of observables that appear in \mathbf{W}_* . Let \mathbf{w}_* represent a vector of values, one for each variable in \mathbf{W}_* and define $\mathbf{w}_*(\mathbf{X}_*)$ as the subset of \mathbf{w}_* corresponding to \mathbf{X}_* for any $\mathbf{X}_* \subseteq \mathbf{W}_*$.

The probability of any counterfactual event is given by

$$P(\mathbf{Y}_* = \mathbf{y}_*) = \sum_{\{\mathbf{u} \mid \mathbf{Y}_*(\mathbf{u}) = \mathbf{y}_*\}} P(\mathbf{u}), \quad (1)$$

where the predicate $\mathbf{Y}_*(\mathbf{u}) = \mathbf{y}_*$ means $\bigwedge_{\{Y_{\widehat{\mathbf{X}}} \in \mathbf{Y}_*\}} Y_{\widehat{\mathbf{X}}}(\mathbf{u}) = y$.

When all of the variables in the expression have the same subscript, that is, they belong to the same submodel; we will often denote it as $P_{\mathbf{x}}(W_1, W_2, \dots)$.

For most real-world scenarios, having access to a fully specified SCM of the underlying system is unfeasible. Nevertheless, our analysis does not rely on such privileged access but the aspects of the model captured by the causal graph and data samples generated by the unobserved model.

2.1 Nested Counterfactuals

Potential responses can be compounded based on natural interventions. For instance, the counterfactual $Y_{Z_x}(\mathbf{u})$ ($Y_{Z=Z_x}(\mathbf{u})$) can be seen as the potential response of Y to an intervention that makes \widehat{Z} equal to Z_x . Notice that $Z_x(\mathbf{u})$ is in itself a potential response, but from a different (nested) model. Hence we call Y_{Z_x} a *nested counterfactual*.

Recall the causal diagram in [Fig. 1\(b\)](#) and consider once again the NDE as

$$NDE_{x \rightarrow x', Z}(Y) = E[Y_{x'Z_x}] - E[Y_x]. \quad (2)$$

The second term is also equal to Y_{xZ_x} as Z_x is consistent with $X = x$, so it is the value Y listens to in \mathcal{M}_x . Meanwhile, the first one is indeed related to $P(Y_{x'Z_x})$, the probability of a nested counterfactual.

The following result shows how nested counterfactuals can be written in terms of non-nested ones.

Theorem 1 (Counterfactual Unnesting Theorem (CUT)). *Let $\widehat{\mathbf{X}}, \widehat{\mathbf{Z}}$ be any natural interventions on disjoint sets $\mathbf{X}, \mathbf{Z} \subseteq \mathbf{V}$. Then for $\mathbf{Y} \subseteq \mathbf{V}$ disjoint from \mathbf{X} and \mathbf{Z} , we have*

$$P(\mathbf{Y}_{\widehat{\mathbf{Z}}, \widehat{\mathbf{X}}} = \mathbf{y}) = \sum_{\mathbf{x} \in \mathfrak{X}_{\mathbf{X}}} P(\mathbf{Y}_{\widehat{\mathbf{Z}}, \mathbf{x}} = \mathbf{y}, \widehat{\mathbf{X}} = \mathbf{x}). \quad (3)$$

Proof outline. Based on [Eq. \(1\)](#), $P(\mathbf{Y}_{\widehat{\mathbf{Z}}, \widehat{\mathbf{X}}} = \mathbf{y})$ can be seen as a sum of the probabilities $P(\mathbf{u})$ for the \mathbf{u} that induce the event $(\mathbf{Y}_{\widehat{\mathbf{Z}}, \widehat{\mathbf{X}}} = \mathbf{y})$. Such set of \mathbf{u} can be partitioned based on the values $\mathbf{x} = \widehat{\mathbf{X}}(\mathbf{u})$, $\mathbf{x} \in \mathfrak{X}_{\mathbf{X}}$ they induce, which are the same that induce the event $\mathbf{Y}_{\widehat{\mathbf{Z}}, \mathbf{x}} = \mathbf{y}, \widehat{\mathbf{X}} = \mathbf{x}$. Then, the sum over $P(\mathbf{u})$ for each subset is equal to the value of the original nested counterfactual. \square

For instance, for the model in [Fig. 1\(b\)](#) we can write

$$P(Y_{x'Z_x} = y) = \sum_z P(Y_{x'z} = y, Z_x = z). \quad (4)$$

As [Theorem 1](#) allows us to re-write any nested counterfactual in terms of non-nested counterfactuals, we focus on the latter and assume that any given counterfactual is already unnested.




Figure 2: Two causal diagrams and the subgraphs considered when finding sets of ancestors for a counterfactual variable.

2.2 Tools for Counterfactual Reasoning

Before characterizing the identification of counterfactuals from observational and experimental data, we develop from first principles a canonical representation of any such query. First, we extend the notion of ancestors for counterfactual variables, which subsumes the usual one described before.

Definition 3 (Ancestors, of a counterfactual). Let $Y_{\mathbf{x}}$ be such that $Y \in \mathbf{V}$, $\mathbf{X} \subseteq \mathbf{V}$. Then, the set of (counterfactual) ancestors of $Y_{\mathbf{x}}$, denoted $An(Y_{\mathbf{x}})$, consist of each $W_{\mathbf{z}}$, such that $W \in An(Y)_{\mathcal{G}_{\mathbf{x}}}$ (which includes Y itself), and $\mathbf{z} = \mathbf{x} \cap An(W)_{\mathcal{G}_{\mathbf{x}}}$.

For a set of variables \mathbf{W}_* , we define $An(\mathbf{W}_*)$ as the union of the ancestors of each variable in the set. That is, $An(\mathbf{W}_*) = \bigcup_{W_{\mathbf{t}} \in \mathbf{W}_*} An(W_{\mathbf{t}})$. For instance, in Fig. 2(a), $An(Y_x) = \{Y_x, Z\}$, $An(X_{y_z}) = \{X_z\}$ and $An(Y_z) = \{Y_z, X_z\}$ (depicted in Fig. 2(b)). In Fig. 2(c), $An(Z, Y_z) = \{Y_z, X_z, Z, W\}$ and $An(Y_x) = \{Y_x\}$ (represented in Fig. 2(d)).

Given a counterfactual variable $Y_{\mathbf{x}}$, it could be the case that some values in \mathbf{x} become causally irrelevant to Y after the rest of \mathbf{x} has been fixed. Formally,

Lemma 1. Let $\|Y_{\mathbf{x}}\| := Y_{\mathbf{z}}$ where $\mathbf{Z} = \mathbf{X} \cap An(Y)_{\mathcal{G}_{\mathbf{x}}}$ and \mathbf{z} is consistent with \mathbf{x} . Then, $Y_{\mathbf{x}} = \|Y_{\mathbf{x}}\|$.

Moreover, such simplification may reveal counterfactual expressions with equivalent or contradicting events. In Fig. 2(c), $(Y_{xz} = y, Y_{xz'} = y') = (Y_x = y, Y_x = y')$ which has probability 0 if $y \neq y'$, or $(Y_{xz} = y, Y_{xz'} = y) = (Y_x = y)$. Similarly, the probabilities of counterfactual events of the form $P(X_x = x')$, $x \neq x'$, and $P(X_x = x)$ are trivially 0 and 1 respectively.

For a set of counterfactual variables \mathbf{Y}_* let $\|\mathbf{Y}_*\| = \bigcup_{Y_{\mathbf{x}} \in \mathbf{Y}_*} \|Y_{\mathbf{x}}\|$. Notice that each variable in the ancestral set is ‘‘interventionally minimal’’ in the sense of Lemma 1.

Probabilistic and causal inference with graphical models exploits local structure among variables, specifically parent-child relationships, to infer and even estimate probabilities. In particular, Tian [23] introduced c -factors which have proven instrumental in solving many problems in causal inference. We naturally generalize this notion to the counterfactual setting with the following definition.

Definition 4 (Counterfactual Factor (ctf-factor)). A counterfactual factor is a distribution of the form

$$P(W_{1[\mathbf{pa}_1]} = w_1, W_{2[\mathbf{pa}_2]} = w_2, \dots, W_{l[\mathbf{pa}_l]} = w_l), \quad (5)$$

where each $W_i \in \mathbf{V}$ and there could be $W_i = W_j$ for some $i, j \in \{1, \dots, l\}$.

For example, for Fig. 2(c) $P(Y_x = y, Y_{x'} = y')$, $P(Y_x = y, X_z = x)$ are ctf-factors but $P(Y_z = y, Z_w = z)$ is not. Using the notion of ancestrality introduced in Definition 3, we can factorize counterfactual probabilities as ctf-factors.

Theorem 2 (Ancestral set factorization). Let \mathbf{W}_* be an ancestral set, that is, $An(\mathbf{W}_*) = \mathbf{W}_*$, and let \mathbf{w}_* be a vector with a value for each variable in \mathbf{W}_* . Then,

$$P(\mathbf{W}_* = \mathbf{w}_*) = P\left(\bigwedge_{W_{\mathbf{t}} \in \mathbf{W}_*} W_{\mathbf{pa}_w} = w\right), \quad (6)$$

where each w is taken from \mathbf{w}_* and \mathbf{pa}_w is determined for each $W_{\mathbf{t}} \in \mathbf{W}_*$ as follows:

- (i) the values for variables in $\mathbf{Pa}_w \cap \mathbf{T}$ are the same as in \mathbf{t} , and
- (ii) the values for variables in $\mathbf{Pa}_w \setminus \mathbf{T}$ are taken from \mathbf{w}_* corresponding to the parents of $W_{\mathbf{t}}$.

Proof outline. Following a reverse topological order in \mathcal{G} , look at each $W_{\mathbf{t}_i} \in \mathbf{W}_*$. Since any parent of $W_{\mathbf{t}_i}$ not in \mathbf{T}_i must appear in \mathbf{W}_* , the composition axiom [13, 7.3.1] licenses adding them to




Figure 3: Three causal diagrams representing plausible structures in mediation analysis.

the subscript. Then, by exclusion restrictions [12], any intervention not involving $Pa(W_i)$ can be removed to obtain the form in Eq. (6) \square

For example, consider the diagram in Fig. 2(c) and the counterfactual $P(Y_x = y | X = x')$ known as the *effect of the treatment on the treated* (ETT) [6, 13]. First note that $P(Y_x = y | X = x') = P(Y_x = y, X = x') / P(X = x')$ and that $An(Y_x, X) = \{Y_x, X, Z, W\}$, then

$$P(Y_x = y, X = x') = \sum_{z,w} P(Y_x = y, X = x', Z = z, W = w). \quad (7)$$

Then, by Theorem 2 we can write

$$P(Y_x = y, X = x') = \sum_{z,w} P(Y_x = y, X_z = x', Z_w = z, W = w). \quad (8)$$

Moreover, the following result describes a factorization of ctf-factors based on the c-component structure of the graph, which will prove instrumental in the next section.

Theorem 3 (Counterfactual factorization). *Let $P(\mathbf{W}_* = \mathbf{w}_*)$ be a ctf-factor, let $W_1 < W_2 < \dots$ be a topological order over the variables in $\mathcal{G}[\mathbf{V}(\mathbf{W}_*)]$, and let $\mathbf{C}_1, \dots, \mathbf{C}_k$ be the c-components of the same graph. Define $\mathbf{C}_{j*} = \{W_{\text{pa}_w} \in \mathbf{W}_* \mid W \in \mathbf{C}_j\}$ and \mathbf{c}_{j*} as the values in \mathbf{w}_* corresponding to \mathbf{C}_{j*} , then $P(\mathbf{W}_* = \mathbf{w}_*)$ decomposes as*

$$P(\mathbf{W}_* = \mathbf{w}) = \prod_j P(\mathbf{C}_{j*} = \mathbf{c}_{j*}). \quad (9)$$

Furthermore, each factor can be computed from $P(\mathbf{W}_* = \mathbf{w})$ as

$$P(\mathbf{C}_{j*} = \mathbf{c}_{j*}) = \prod_{\{W_i \in \mathbf{C}_j\}} \frac{\sum_{\{w \mid W_{\text{pa}_w} \in \mathbf{W}_*, W_i < W\}} P(\mathbf{W}_* = \mathbf{w}_*)}{\sum_{\{w \mid W_{\text{pa}_w} \in \mathbf{W}_*, W_{i-1} < W\}} P(\mathbf{W}_* = \mathbf{w}_*)}. \quad (10)$$

Armed with these results, we consider the identification problem in the next section.

3 Counterfactual Identification from Observations and Experiments

In this section, we consider the identification of a counterfactual probability from a collection of observational and experimental distributions. This task can be seen as a generalization of that in [9] where the available data is the same, but the query is a causal effect $P_{\mathbf{x}}(\mathbf{Y})$. Let $\mathbb{Z} = \{\mathbf{Z}_1, \mathbf{Z}_2, \dots\}$, $\mathbf{Z}_j \subseteq \mathbf{V}$, and assume that all of $\{P_{\mathbf{z}_j}(\mathbf{V})\}_{\mathbf{z}_j \in \mathbb{Z}}$ are available. Notice that $\mathbf{Z}_j = \emptyset$ is a valid choice corresponding to $P(\mathbf{V})$ the observational (non-interventional) distribution.

Definition 5 (Counterfactual Identification). A query $P(\mathbf{Y}_* = \mathbf{y}_*)$ is said to be identifiable from \mathbb{Z} in \mathcal{G} , if $P(\mathbf{Y}_* = \mathbf{y}_*)$ is uniquely computable from the distributions $\{P_{\mathbf{z}_j}(\mathbf{V})\}_{\mathbf{z}_j \in \mathbb{Z}}$ in any causal model which induces \mathcal{G} .

Given an arbitrary query $P(\mathbf{Y}_* = \mathbf{y}_*)$, we could express it in terms of ctf-factors by writing $P(\mathbf{Y}_* = \mathbf{y}_*) = \sum_{\mathbf{d}_* \setminus \mathbf{y}_*} P(\mathbf{D}_* = \mathbf{d}_*)$ where $\mathbf{D}_* = An(\mathbf{Y}_*)$ and then using Theorem 2 to write $P(\mathbf{D}_* = \mathbf{d}_*)$ as a ctf-factor. For instance, the ancestral set $\mathbf{W}_* = \{Y_x, X, Z, W\}$ with $\mathbf{w} = \{y, x', z, w\}$ in Eq. (8) can be written in terms of ctf-factors as

$$P(Y_x = y, X_z = x', Z_w = z, W = w) = P(Y_x = y, X_z = x', W = w)P(Z_w = z). \quad (11)$$

The following lemma characterizes the relationship between the identifiability of $P(\mathbf{Y}_* = \mathbf{y}_*)$ and $P(\mathbf{D}_* = \mathbf{d}_*)$.




Figure 4: Examples of causal diagrams and inconsistent ctf-factors derived from them.

Lemma 2. Let $P(\mathbf{W}_* = \mathbf{w}_*)$ be a ctf-factor and let $\mathbf{Y}_* \subseteq \mathbf{W}_*$ be such that $\mathbf{W}_* = \text{An}(\mathbf{Y}_*)$. Then, $\sum_{\mathbf{w}_* \setminus \mathbf{y}_*} P(\mathbf{W}_* = \mathbf{w}_*)$ is identifiable from \mathbb{Z} if and only if $P(\mathbf{W}_* = \mathbf{w}_*)$ is identifiable from \mathbb{Z} .

Once the query of interest is in ctf-factor-form, the identification task reduces to identifying smaller ctf-factors according to the c-components of \mathcal{G} . In this respect, [Theorem 3](#) implies the following

Corollary 1. Let $P(\mathbf{W}_* = \mathbf{w}_*)$ be a ctf-factor and \mathbf{C}_j be a c-component of $\mathcal{G}[\mathbf{V}(\mathbf{W}_*)]$. Then, if $P(\mathbf{C}_{j*} = \mathbf{c}_{j*})$ is not identifiable, $P(\mathbf{W}_* = \mathbf{w}_*)$ is also not identifiable.

Proof. Assume for the sake of contradiction that $P(\mathbf{C}_{j*} = \mathbf{c}_{j*})$ is not identifiable but $P(\mathbf{W}_* = \mathbf{w}_*)$ is. Then, by [Theorem 3](#) the former is identifiable from the latter, a contradiction. \square

Let us consider the causal diagrams in [Fig. 3](#) and the counterfactual $Y_{x_1, W_{x_0}} = y, X = x$, with $x_0, x_1, x \in \mathfrak{X}_X$, used to define quantities for fairness analysis in [\[24\]](#) (e.g., $Y_{x_1, W_{x_0}} = y | X = x$):

$$\begin{aligned} P(Y_{x_1, W_{x_0}} = y, X = x) & \\ &= \sum_w P(Y_{x_1, w} = y, W_{x_0} = w, X = x) && \text{Unnesting} \quad (12) \end{aligned}$$

$$= \sum_{w, z} P(Y_{x_1, w} = y, W_{x_0} = w, X = x, Z = z) \quad \text{Complete ancestral set} \quad (13)$$

$$= \sum_{w, z} P(Y_{x_1, w, z} = y, W_{x_0} = w, X_z = x, Z = z) \quad \text{Write in ctf-factor-form} \quad (14)$$

Due to the particular c-component structure of each model, we can factorize $P(Y_{x_1, w, z} = y, W_{x_0} = w, X_z = x, Z = z)$ according to each model as:

$$P(Y_{x_1, w, z} = y)P(W_{x_0} = w)P(X_z = x)P(Z = z), \quad (15)$$

$$P(Y_{x_1, w, z} = y, Z = z)P(W_{x_0} = w)P(X_z = x), \text{ and} \quad (16)$$

$$P(Y_{x_1, w, z} = y)P(W_{x_0} = w, X_z = x)P(Z = z). \quad (17)$$

The question then becomes, whether ctf-factors corresponding to individual c-components can be identified from the available input. In this example, all factors in [Eq. \(15\)](#) and [Eq. \(16\)](#) are identifiable from $P(\mathbf{V})$. For [Eq. \(16\)](#) in particular, they are given by

$$P(Y = y, Z = z | W = w, X = x_1)P(W = w | X = x_0)P(X = x | Z = z). \quad (18)$$

In contrast, the factor $P(W_{x_0} = w, X_z = x)$ in [Eq. \(17\)](#) (model [Fig. 3\(c\)](#)) is only identifiable if $x = x_0$. The following definition and theorem characterize the factors that can be identified from \mathbb{Z} and \mathcal{G} .

Definition 6 (Inconsistent ctf-factor). $P(\mathbf{W}_* = \mathbf{w}_*)$ is an inconsistent ctf-factor if it is a ctf-factor, $\mathcal{G}[\mathbf{V}(\mathbf{W}_*)]$ has a single c-component, and one of the following situations hold:

- (i) there exist $W_t \in \mathbf{W}_*, Z \in \mathbf{T} \cap \mathbf{V}(\mathbf{W}_*)$ such that $z \in \mathbf{t}, z' \in \mathbf{w}_*$ and $z \neq z'$, or
- (ii) there exists $W_{i[t_i]}, W_{j[t_j]} \in \mathbf{W}_*$ and $T \in \mathbf{T}_i \cap \mathbf{T}_j$ such that $t \in \mathbf{t}_1, t' \in \mathbf{t}_2$ and $t \neq t'$.

Theorem 4 (Ctf-factor identifiability). A ctf-factor $P(\mathbf{W}_* = \mathbf{w})$ is identifiable from \mathbb{Z} if and only if it is consistent. If consistent, let $\mathbf{W} = \mathbf{V}(\mathbf{W}_*)$ and $\mathbf{W}' = \mathbf{V} \setminus \mathbf{W}$; then $P(\mathbf{W}_* = \mathbf{w}_*)$ is equal to $P_{\mathbf{w}'}(\mathbf{w})$ where \mathbf{w} and \mathbf{w}' are consistent with $\mathbf{w}_* \cup \bigcup_{\{W_{\mathbf{pa}_w} \in \mathbf{W}_*\}} \mathbf{pa}_w$.

Consider the $NDE_{x \rightarrow x', Z}(Y)$ in [Fig. 4\(a\)](#), we can write

$$P(Y_{x'Z_x} = y) = \sum_z P(Y_{x'z} = y, Z_x = z) = \sum_z P(Y_{x'z} = y)P(Z_x = z). \quad (19)$$

While the factor $P(Y_{x'z} = y)$ is identifiable from $P(\mathbf{V})$ as $P(Y = y | X = x', Z = z)$, the second factor is identifiable only if experimental data on X is available, as $P_x(z)$.

Algorithm 1 CTFID($\mathbf{Y}_*, \mathbf{y}_*, \mathbb{Z}, \mathcal{G}$)

Input: \mathcal{G} causal diagram over variables \mathbf{V} ; \mathbf{Y}_* a set of counterfactual variables in \mathbf{V} ; \mathbf{y}_* a set of values for \mathbf{Y}_* ; and available distribution specification \mathbb{Z} .

Output: $P(\mathbf{Y}_* = \mathbf{y}_*)$ in terms of available distributions or FAIL if not identifiable from $\langle \mathcal{G}, \mathbb{Z} \rangle$.

- 1: let $\mathbf{Y}_* \leftarrow \|\mathbf{Y}_*\|$.
 - 2: **if** there exists $Y_x \in \mathbf{Y}_*$ with two or more different values in $\mathbf{y}_*(Y_x)$ or $Y_y \in \mathbf{Y}_*$ with $\mathbf{y}_*(Y_y) \neq y$ **then return 0**.
 - 3: **if** there exists $Y_x \in \mathbf{Y}_*$ with two consistent values in $\mathbf{y}_*(Y_x)$ or $Y_y \in \mathbf{Y}_*$ with $\mathbf{y}_*(Y_y) = y$ **then** remove repeated variables from \mathbf{Y}_* and values \mathbf{y}_* .
 - 4: let $\mathbf{W}_* \leftarrow An(\mathbf{Y}_*)$, and let $\mathbf{C}_1, \dots, \mathbf{C}_n$ be the c-components of $\mathcal{G}[\mathbf{V}(\mathbf{W}_*)]$.
 - 5: **for each** \mathbf{C}_i s.t. $(\mathbf{C}_{i*} = \mathbf{c}_{i*})$ is not inconsistent, $\mathbf{Z} \in \mathbb{Z}$ s.t. $\mathbf{C}_i \cap \mathbf{Z} = \emptyset$ **do**
 - 6: let \mathbf{B}_i be the c-component of $\mathcal{G}_{\mathbb{Z}}$ such that $\mathbf{C}_i \subseteq \mathbf{B}_i$, compute $P_{\mathbf{V} \setminus \mathbf{B}_i}(\mathbf{B}_i)$ from $P_{\mathbf{Z}}(\mathbf{V})$.
 - 7: **if** IDENTIFY($\mathbf{C}_i, \mathbf{B}_i, P_{\mathbf{V} \setminus \mathbf{B}_i}(\mathbf{B}_i), \mathcal{G}$) does not FAIL **then**
 - 8: let $P_{\mathbf{V} \setminus \mathbf{C}_i}(\mathbf{C}_i) \leftarrow$ IDENTIFY($\mathbf{C}_i, \mathbf{B}_i, P_{\mathbf{V} \setminus \mathbf{B}_i}(\mathbf{B}_i), \mathcal{G}$).
 - 9: let $P(\mathbf{C}_{i*} = \mathbf{c}_{i*}) \leftarrow [P_{\mathbf{V} \setminus \mathbf{C}_i}(\mathbf{C}_i)]_{(\mathbf{c}_{i*} \cup \cup_{\mathbf{C}_t \in \mathbf{C}_{i*}} \mathbf{pa}_{\mathbf{C}_t})}$.
 - 10: move to the next \mathbf{C}_i .
 - 11: **end if**
 - 12: **end for**
 - 13: **if** any $P(\mathbf{C}_{i*} = \mathbf{c}_{i*})$ is inconsistent or was not identified from \mathbb{Z} **then return FAIL**.
 - 14: **return** $P(\mathbf{Y}_* = \mathbf{y}_*) \leftarrow \sum_{\mathbf{w}_* \setminus \mathbf{y}_*} \prod_i P(\mathbf{C}_{i*} = \mathbf{c}_{i*})$.
-

We can also verify that the factor $P(Y_x = y, X = x')$ in Fig. 4(b) is inconsistent. For another example consider the ETT-like expression $P(Y_{x,z} = y, X = x', Z = z')$ in Fig. 4(c), we have

$$P(Y_{xz} = y, X = x', Z = z') = \sum_{w, w'} P(Y_{xz} = y, X = x', Z = z', W_x = w, W = w') \quad (20)$$

$$= \sum_{w, w'} P(Y_{xz} = y, X = x', Z_{w'} = z', W_x = w, W_{x'} = w') \quad (21)$$

$$= \sum_{w, w'} P(Y_{xz} = y)P(X = x')P(Z_{w'} = z')P(W_x = w, W_{x'} = w'), \quad (22)$$

where the factor $P(W_x = w, W_{x'} = w')$ is inconsistent.

Using the results in this section, we propose the algorithm CTFID (Algorithm 1) which given a set of counterfactual variables \mathbf{Y}_* , corresponding values \mathbf{y}_* , a collection of observational and experimental distributions \mathbb{Z} , and a causal diagram \mathcal{G} ; outputs an expression for $P(\mathbf{Y}_* = \mathbf{y}_*)$ in terms of the specified distributions or FAIL if the query is not identifiable from such input in \mathcal{G} . Line 1 removes irrelevant subscripts from the query by virtue of Lemma 1. Then, lines 2 and 3 look for inconsistent events and redundant events, respectively. Line 4 finds the relevant ctf-factors consisting of a single c-component, as licensed by Theorem 2 and Theorem 3. As long as the factors are consistent, and allowed by Theorem 4, lines 6-11 carry out identification of the causal effect $P_{\mathbf{V} \setminus \mathbf{C}_i}(\mathbf{C}_i)$ from the available distributions employing the algorithm IDENTIFY [22] as a subroutine.¹ The procedure fails if any of the factors $P(\mathbf{C}_{i*} = \mathbf{c}_{i*})$ is inconsistent or not identifiable from \mathbb{Z} . Otherwise, it returns the corresponding expression.

Theorem 5 (CTFID completeness). *A counterfactual probability $P(\mathbf{Y}_* = \mathbf{y}_*)$ is identifiable from \mathbb{Z} and \mathcal{G} if and only if CTFID returns an expression for it.*

4 Identification of Conditional Counterfactuals

In this section we consider counterfactual quantities of the form $P(\mathbf{Y}_* = \mathbf{y}_* \mid \mathbf{X}_* = \mathbf{x}_*)$. It is immediate to write such a query as $P(\mathbf{W}_* = \mathbf{w}_*)$ with $\mathbf{W}_* = \mathbf{Y}_* \cup \mathbf{X}_*$, and try to identify it using CTFID. Nevertheless, depending on the graphical structure, the original query may be identifiable even if the latter is not. To witness, consider the causal diagram in Fig. 5(a) and the counterfactual $P(Y_x = y \mid Z_x = z, X = x')$, which can be written as $P(Y_x = y, Z_x = z, X =$

¹For a running example of the CTRID and details on how to use IDENTIFY, see Appendix E.



(a) While $P(Y_x=y \mid Z_x=z, X=x')$ is identifiable from \mathbb{Z} , $P(Y_x=y, Z_x=z, X=x')$ is not. (b) $P(Y_x=y \mid Z_x=z, X=x')$ is not identifiable from \mathbb{Z} because of the factor $P(Y_{xz}=y, X=x')$.

Figure 5: Examples of conditional queries.

Algorithm 2 COND-CTFID($\mathbf{Y}_*, \mathbf{y}_*, \mathbf{X}_*, \mathbf{x}_*, \mathbb{Z}, \mathcal{G}$)

Input: \mathcal{G} causal diagram over variables \mathbf{V} ; $\mathbf{Y}_*, \mathbf{X}_*$ a set of counterfactual variables in \mathbf{V} ; $\mathbf{y}_*, \mathbf{x}_*$ a set of values for \mathbf{Y}_* and \mathbf{X}_* ; and available distribution specification \mathbb{Z} .

Output: $P(\mathbf{Y}_*=\mathbf{y}_* \mid \mathbf{X}_*=\mathbf{x}_*)$ in terms of available distributions or FAIL if non-ID from $\langle \mathcal{G}, \mathbb{Z} \rangle$.

- 1: Let $\mathbf{A}_{1*}, \mathbf{A}_{2*}, \dots$ be the ancestral components of $\mathbf{Y}_* \cup \mathbf{X}_*$ given \mathbf{X}_* .
 - 2: Let \mathbf{D}_* be the union of the ancestral components containing a variable in \mathbf{Y}_* and \mathbf{d}_* the corresponding set of values.
 - 3: let $Q \leftarrow \text{CTFID}(\bigcup_{D_t \in \mathbf{D}_*} \mathbf{D}_{\text{pa}_d}, \mathbf{d}_*, \mathbb{Z}, \mathcal{G})$.
 - 4: **return** $\sum_{\mathbf{d}_* \setminus (\mathbf{y}_* \cup \mathbf{x}_*)} Q / \sum_{\mathbf{d}_* \setminus \mathbf{x}_*} Q$.
-

$x') / \sum_y P(Y_x = y, Z_x = z, X = x')$. Following the strategy explained so far, the numerator is equal to $P(Y_z = y)P(Z_x = z, X = x')$, where the second ctf-factor is inconsistent, and therefore not identifiable from \mathbb{Z} . Nevertheless, the conditional query is identifiable as

$$\frac{P(Y_{xz} = y)P(Z_x = z, X = x')}{P(Z_x = z, X = x') \sum_y P(Y_{xz} = y)} = P(Y_{xz} = y) = P(Y = y \mid Z = z, X = x). \quad (23)$$

To characterize such simplifications of the query, we look at the causal diagram paying special attention to variables after the conditioning bar that are also ancestors of those before. Let $\mathbf{X}_*(W_t) = \mathbf{V}(\|\mathbf{X}_* \parallel \cap \text{An}(W_t))$, that is, the primitive variables in \mathbf{X}_* that are also ancestors of W_t .

Definition 7 (Ancestral components). Let \mathbf{W}_* be a set of counterfactual variables, $\mathbf{X}_* \subseteq \mathbf{W}_*$, and \mathcal{G} be a causal diagram. Then the ancestral components induced by \mathbf{W}_* , given \mathbf{X}_* , are sets $\mathbf{A}_{1*}, \mathbf{A}_{2*}, \dots$ that form a partition over $\text{An}(\mathbf{W}_*)$, made of unions of ancestral the sets $\text{An}(W_t)_{\mathcal{G}_{\mathbf{X}_*(W_t)}}$, $W_t \in \mathbf{W}_*$. Sets $\text{An}(W_{1[t_1]})_{\mathcal{G}_{\mathbf{X}_*(W_{1[t_1]})}}$ and $\text{An}(W_{2[t_2]})_{\mathcal{G}_{\mathbf{X}_*(W_{2[t_2]})}}$ are put together if they are not disjoint or there exists a bidirected arrow in \mathcal{G} connecting variables in those sets.

Lemma 3. Let $\mathbf{Y}_*, \mathbf{X}_*$ be two sets of counterfactual variables and let \mathbf{D}_* be the set of variables in the same ancestral component, given \mathbf{X}_* , as any variable in \mathbf{Y}_* , then

$$P(\mathbf{Y}_* = \mathbf{y}_* \mid \mathbf{X}_* = \mathbf{x}_*) = \frac{\sum_{\mathbf{d}_* \setminus (\mathbf{y}_* \cup \mathbf{x}_*)} P(\bigwedge_{D_t \in \mathbf{D}_*} \mathbf{D}_{\text{pa}_d} = d)}{\sum_{\mathbf{d}_* \setminus \mathbf{x}_*} P(\bigwedge_{D_t \in \mathbf{D}_*} \mathbf{D}_{\text{pa}_d} = d)}, \quad (24)$$

where pa_d is consistent with \mathbf{t} and \mathbf{d}_* , for each $D_t \in \mathbf{D}_*$. Moreover, $P(\mathbf{Y}_* = \mathbf{y}_* \mid \mathbf{X}_* = \mathbf{x}_*)$ is identifiable from \mathbb{Z} if and only if $P(\bigwedge_{D_t \in \mathbf{D}_*} \mathbf{D}_{\text{pa}_d} = d)$ is identifiable from \mathbb{Z} .

Using the notion of ancestral components and [Lemma 3](#), we propose a conditional version of CTFID ([Algorithm 2](#)). Due to [Lemma 3](#), it is easy to see that COND-CTFID is complete.

Theorem 6 (COND-CTFID completeness). A counterfactual probability $P(\mathbf{Y}_* = \mathbf{y}_* \mid \mathbf{X}_* = \mathbf{x}_*)$ is identifiable from \mathbb{Z} and \mathcal{G} if and only if COND-CTFID returns an expression for it.

5 Conclusions

We examined the identification of nested and non-nested counterfactuals from an arbitrary combination of observational and experimental distributions. We study nested counterfactuals within the SCM framework and prove several properties of counterfactual distributions ([Theorem 2](#), [3](#)) together with the counterfactual unnesting theorem ([Theorem 1](#)). Moreover, we developed a graphical condition ([Definition 6](#), [Theorem 4](#)) and an efficient algorithm for identifying marginal counterfactuals ([Algorithm 1](#)) and proved their sufficiency and necessity ([Theorem 5](#)). Lastly, we reduce the identification of conditional counterfactuals to that of marginal ones ([Lemma 3](#)) and give a corresponding complete algorithm ([Algorithm 2](#), [Theorem 6](#)) for this task. These results advance the state of the art by allowing for nested counterfactuals and relaxing the requirements on the data available to the analyst.

References

- [1] Bareinboim, E., Correa, J. D., Ibeling, D., and Icard, T. (2020). On Pearl’s Hierarchy and the Foundations of Causal Inference. Technical Report R-60, Causal Artificial Intelligence Lab, Columbia University. In: “Probabilistic and Causal Inference: The Works of Judea Pearl”, ACM Books, in press.
- [2] Bareinboim, E. and Pearl, J. (2012). Causal Inference by Surrogate Experiments: z-Identifiability. In Murphy, N. d. F. and Kevin, editors, *Proceedings of the Twenty-Eighth Conference on Uncertainty in Artificial Intelligence*, pages 113–120. AUAI Press.
- [3] Dawid, A. P. (2002). Influence diagrams for causal modelling and inference. *International Statistical Review*, 70:161–189.
- [4] Dawid, A. P. (2015). Statistical Causality from a Decision-Theoretic Perspective. *Annual Review of Statistics and Its Application*, 2(1):273–303.
- [5] Galles, D. and Pearl, J. (1995). Testing identifiability of causal effects. In Besnard, P. and Hanks, S., editors, *Uncertainty in Artificial Intelligence 11*, pages 185–195. Morgan Kaufmann, San Francisco.
- [6] Heckman, J. J. (1992). Randomization and Social Policy Evaluation. In Manski, C. and Garfinkle, I., editors, *Evaluations: Welfare and Training Programs*, pages 201–230. Harvard University Press, Cambridge, MA.
- [7] Huang, Y. and Valtorta, M. (2006). Identifiability in Causal Bayesian Networks: A Sound and Complete Algorithm. In *Proceedings of the Twenty-First National Conference on Artificial Intelligence (AAAI 2006)*, pages 1149–1156. AAAI Press, Menlo Park, CA.
- [8] Huang, Y. and Valtorta, M. (2008). On the completeness of an identifiability algorithm for semi-Markovian models. *Annals of Mathematics and Artificial Intelligence*, 54(4):363–408.
- [9] Lee, S., Correa, J. D., and Bareinboim, E. (2019). General Identifiability with Arbitrary Surrogate Experiments. In *Proceedings of the Thirty-Fifth Conference Annual Conference on Uncertainty in Artificial Intelligence*, Corvallis, OR. AUAI Press.
- [10] Lee, S., Correa, J. D., and Bareinboim, E. (2020). Generalized Transportability: Synthesis of Experiments from Heterogeneous Domains. In *Proceedings of the 34th AAAI Conference on Artificial Intelligence*, Menlo Park, CA. AAAI Press.
- [11] Pearl, J. (1994). A probabilistic calculus of actions. In de Mantaras, R. L. and D. Poole, editors, *Uncertainty in Artificial Intelligence 10*, pages 454–462. Morgan Kaufmann, San Mateo, CA.
- [12] Pearl, J. (1995). Causal diagrams for empirical research. *Biometrika*, 82(4):669–688.
- [13] Pearl, J. (2000). *Causality: Models, Reasoning, and Inference*. Cambridge University Press, New York, NY, USA, 2nd edition.
- [14] Pearl, J. (2001). Direct and indirect effects. In *Proceedings of the Seventeenth Conference on Uncertainty in Artificial Intelligence*, pages 411–420. Morgan Kaufmann, San Francisco, CA.
- [15] Pearl, J. and Mackenzie, D. (2018). *The Book of Why*. Basic Books, New York.
- [16] Pearl, J. and Robins, J. M. (1995). Probabilistic evaluation of sequential plans from causal models with hidden variables. In Besnard, P. and Hanks, S., editors, *Uncertainty in Artificial Intelligence 11*, pages 444–453. Morgan Kaufmann, San Francisco.
- [17] Shpitser, I. and Pearl, J. (2006). Identification of Joint Interventional Distributions in Recursive semi-Markovian Causal Models. In *Proceedings of the Twenty-First AAAI Conference on Artificial Intelligence*, volume 2, pages 1219–1226.
- [18] Shpitser, I. and Pearl, J. (2007). What Counterfactuals Can Be Tested. In *Proceedings of the Twenty-Third Conference on Uncertainty in Artificial Intelligence*, pages 352–359. AUAI Press, Vancouver, BC, Canada.
- [19] Shpitser, I. and Pearl, J. (2009). Effects of Treatment on the Treated: Identification and Generalization. In *Proceedings of the Twenty-Fifth Conference on Uncertainty in Artificial Intelligence*, Montreal, Quebec. AUAI Press.
- [20] Shpitser, I. and Sherman, E. (2018). Identification of Personalized Effects Associated With Causal Pathways. In *Proceedings of the 34th Conference on Uncertainty in Artificial Intelligence*, pages 530–539.

- [21] Spirtes, P., Glymour, C. N., and Scheines, R. (2001). *Causation, Prediction, and Search*. MIT Press, 2nd edition.
- [22] Tian, J. and Pearl, J. (2002a). A General Identification Condition for Causal Effects. In *Proceedings of the Eighteenth National Conference on Artificial Intelligence (AAAI 2002)*, pages 567–573, Menlo Park, CA. AAAI Press/The MIT Press.
- [23] Tian, J. and Pearl, J. (2002b). On the Testable Implications of Causal Models with Hidden Variables. *Proceedings of the Eighteenth Conference on Uncertainty in Artificial Intelligence (UAI-02)*, pages 519–527.
- [24] Zhang, J. and Bareinboim, E. (2018). Fairness in Decision-Making-The Causal Explanation Formula. In *AAAI Conference on Artificial Intelligence*.