
General Transportability of Soft Interventions: Completeness Results

Juan D. Correa Elias Bareinboim
Causal Artificial Intelligence Laboratory
Columbia University
{jdcorrea,eb}@cs.columbia.edu

Abstract

The challenge of generalizing causal knowledge across different environments is pervasive in scientific explorations, including in AI, ML, and Data Science. Experiments are usually performed in one environment (e.g., in a lab, on Earth) with the intent, almost invariably, of being used elsewhere (e.g., outside the lab, on Mars), where the conditions are likely to be different. In the causal inference literature, this generalization task has been formalized under the rubric of *transportability* (Pearl and Bareinboim, 2011), where a number of criteria and algorithms have been developed for various settings. Despite the generality of such results, transportability theory has been confined to atomic, *do()*-interventions. In practice, many real-world applications require more complex, stochastic interventions; for instance, in reinforcement learning, agents need to continuously adapt to the changing conditions of an uncertain and unknown environment. In this paper, we extend transportability theory to encompass these more complex types of interventions, which are known as “soft,” both relative to the input as well as the target distribution of the analysis. Specifically, we develop a graphical condition that is both necessary and sufficient for deciding soft-transportability. Second, we develop an algorithm to determine whether a non-atomic intervention is computable from a combination of the distributions available across domains. As a corollary, we show that the σ -calculus is complete for the task of soft-transportability.

1 Introduction

Generalizing causal knowledge across disparate domains (i.e., populations, settings, environments) is at the heart of many inferences across the empirical sciences as well as AI [26, 33, 29]. In economics, for example, the Nobel Prize of 2019 was awarded to Duflo, Banerjee, and Kremer “for their experimental approach to alleviating global poverty”. Their work is in fact about systematically performing experiments on the effect of complex policies related to poverty, and then carefully trying to extrapolate the gathered evidence to other populations [1, 12].¹ The same need to generalize across disparate conditions is present in Reinforcement Learning. For instance, consider a rover trained in the California desert for digging rocks. After exhaustive months of training, NASA wants to deploy the vehicle on Mars, where the environment is not exactly the same, but somewhat similar to Earth. The expectation is that the rover will need minimal “experimentation” (i.e., trial-and-error) on Mars by leveraging the knowledge acquired here, operating more surgically and effectively there.

¹They acknowledge and discuss the challenges of pursuing such task [12]: “the number of possible variations on a given program is potentially infinite, and a theoretical framework is definitely needed to understand which variations are important to replicate and which are not.” Some typical question they try to answer include: “If a program worked for poor rural women in Africa, will it work for middle-income urban men in South Asia?”




Figure 1: (a) A causal graph representing the target domain π^* (before intervention). (b) and (c) filled squared nodes encode the differences between π^* and π^1, π^2 . Squared hollow nodes indicate intervened variables due to policies σ^1, σ^2 . (d) represents π^* under the intended policy σ_{X_1, X_2}^* .

The solution to these apparently disparate tasks faced by Duflo, NASA, and so many other scientists, perhaps surprisingly, relies on a common ingredient, namely, the invariances of some causal mechanisms shared across the different settings. The task of leveraging causal invariances so as to extrapolate experimental knowledge across settings has been formally studied in the causal inference literature under the rubric of *transportability* [28, 2, 19, 20, 4, 5]. Recent work by [22] provided complete graphical and algorithmic treatment for this task, bringing under the same umbrella previous identification results on how to combine observations and experiments [36, 34, 31, 15, 3, 21].

Most of this literature, however, focuses on *atomic* interventions, often times represented by the celebrated *do*-operator [25, 26, Ch. 3]. Formally, $do(X=x)$ represents the symbolic operation of replacing the underlying causal mechanism that naturally dictates the behavior of variable X with a constant value x . While the $do(\cdot)$ is a grounding element in causal analysis, many real-world situations involve *soft interventions*, or *policy interventions*, which respond in a stochastic manner to other variables in the system [38, 14, 7, 27]. The *do*-operator was used to support some types of soft interventions [26, Ch. 4], and was expanded towards a broader set of interventions in the context of dynamic plans [30, 18, 13, 35, 11, 32, 8, 9]. Despite the generality of these results, there is still no formal treatment of transportability/generalizability settings in the context of soft interventions.

For concreteness, consider a medical setting (e.g., in the context of HIV, Cancer, and Chronic Diseases [23, 24]), where patient data collected under different treatment protocols and locations need to be combined to generate a policy in a different site. In the causal graph in Fig. 1(a), X_1 and X_2 are treatments applied in sequence, Z is a secondary condition detected after the first treatment (X_1) that affects the application of X_2 , and Y represents survival. With the task of designing a new protocol σ_{X_1, X_2}^* to treat the condition, three sources of data are available:

1. an observational study from the current hospital (π^*), where the protocol is to be implemented;
2. a controlled study from hospital π^1 that is under policy σ^1 , where treatment X_1 is randomized and X_2 administered only if X_1 is given; and
3. another controlled study from hospital π^2 that follows protocol σ^2 , where the treatment X_1 is applied as in π^* but X_2 is determined by X_1 and the secondary condition Z .

Figs. 1(b),(c) represent the *selection diagrams* (to be formally defined in the next section) that show the qualitative differences between π^* and, respectively, π^1 and π^2 . The square nodes pointing to Y and Z indicate that the corresponding distributions are different for the populations of hospitals π^1 and π^2 , respectively. The square nodes annotated with σ^1 and σ^2 point to treatment variables following the protocols described above. Given the causal description of these protocols, the question is then – how to combine these various datasets, collected across disparate conditions, so as to estimate the effect of the new policy in π^* , i.e., $P^*(Y; \sigma_{X_1, X_2}^*)$? It turns out that the new policy can be evaluated from the available data using the following formula:

$$P^*(y; \sigma_{X_1, X_2}^*) = \sum_{x_1, z, x_2} \underbrace{P^2(y | x_1, x_2; \sigma^2)}_{\text{from } \pi^2} \underbrace{P^1(z | x_1; \sigma^1)}_{\text{from } \pi^1} \underbrace{P(x_2 | z, x_1; \sigma^*)}_{\text{defined by } \sigma_{X_1, X_2}^*} P(x_1; \sigma^*) \quad (1)$$

This is a delicate mixture of soft experimental data from π^1 and π^2 , where each of the factors alone are not sufficient, but combined they are necessary for the evaluation of the new policy. In this paper, our goal is to explicate this non-trivial formula, and, more broadly, characterize under what conditions the causal effect of a new policy can be computed from disparate data sources.

Type	Strategy	Function
Idle	$\sigma_X = \emptyset$	$f'_x = f_x$
Atomic/ <i>do</i>	$\sigma_X = x$	$f'_x = x$, for some $x \in \text{Dom}(X)$
Conditional	$\sigma_X = g(\mathbf{Pa}'_x)$	$f'_x = g(\mathbf{Pa}'_x)$
Stochastic/Random	$\sigma_X = P'(X \mathbf{Pa}'_x)$	f'_x s.t. $P'(x \mathbf{Pa}'_x) = \sum_{\mathbf{u}'_x} P(f'_x(\mathbf{Pa}'_x, \mathbf{u}'_x) = x)P(\mathbf{u}'_x)$

Table 1: Summary of the types of interventions considered. Each row contains the type of intervention, its representation using the regime indicator and the way the corresponding replacement function.

More specifically, our contributions are as follow:

1. **Reduction to atomic case.** We reduce the problem of transporting $P^*(\mathbf{y}; \sigma_{\mathbf{X}})$ to that of transporting the effect of an atomic-intervention. We then prove the tightness of the reduction.
2. **Algorithmic solution.** We design an efficient algorithm to determine the existence of an estimand for the effect of a non-atomic intervention as a function of the available distributions.
3. **Symbolic characterization.** We prove that the σ -calculus is necessary and sufficient for the task of transportability when both the input and the output distributions involve soft interventions.
4. **Graphical characterization.** We derive a complete graphical condition for this task.

Preliminaries. Random variables are represented with uppercase letters (e.g, C) and their realizations with lowercase ones (e.g, c). Letters in bold (e.g, $\mathbf{C}$) represent sets of variables, and lowercase-bold letters (e.g., $\mathbf{c}$) a joint value assignment for the variables in the set. 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 $\mathbf{X}$. Also, $\mathcal{G}_{[\mathbf{C}]}$ is the subgraph of $\mathcal{G}$ made only of nodes in $\mathbf{C} \subseteq \mathbf{V}$ and the edges between them. We define $Pa(\mathbf{C})$ and $An(\mathbf{C})$, as the union of $\mathbf{C} \subseteq \mathbf{V}$ with its parents and ancestors in the graph, respectively.

We use the Structural Causal Model (SCM) paradigm [26, Ch. 7] to represent the causal structure of the system and elucidate the effects of changing it. 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 $\{f_i\}_{V_i \in \mathbf{V}}$. 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, $P(\mathbf{U})$. Each SCM $\mathcal{M}$ induces a *causal diagram* where every $V_i \in \mathbf{V}$ is a vertex, 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 ($V_i \longleftrightarrow V_j$) for every pair $V_i, V_j \in \mathbf{V}$ such that $\mathbf{U}_i \cap \mathbf{U}_j \neq \emptyset$ (V_i and V_j have a common exogenous parent). We assume the underlying model is recursive, that is, there are no cyclic dependencies among the variables.

An intervention on a variable X replaces f_x with a new function f'_x of some $\mathbf{Pa}'_x \subseteq \mathbf{V} \setminus \{X\}$ and variables $\mathbf{U}'_x$. $\mathbf{Pa}'_x$ could differ from $\mathbf{Pa}_x$, and $\mathbf{U}'_x \cap \mathbf{U} = \emptyset$. In particular, $\mathbf{Pa}'_x$ need not be a subset of $\mathbf{Pa}_x$ as long as it does not include any descendant of X . We consider four types of interventions summarized in Table 1. An *idle* intervention leaves the function as it is, so we often omit $\sigma_X = \emptyset$ in any expression. *Atomic* interventions fix the intervened variable to a constant value. *Conditional* and *stochastic* interventions allow the intervened variable to change as a deterministic function or a conditional probability distribution of a set of observable parents. For interventions on a set of variables $\mathbf{X} \subseteq \mathbf{V}$ let $\sigma_{\mathbf{X}} = \{\sigma_X\}_{X \in \mathbf{X}}$, that is, the result of applying one intervention after the other. Each σ_X affects a different variable, so the order in which they are considered is not important. Given an intervention $\sigma_{\mathbf{X}}$ a new model can be defined as $\mathcal{M}_{\sigma_{\mathbf{X}}} = \langle \mathbf{V}, \mathbf{U} \cup \mathbf{U}'_{\mathbf{X}}, \mathcal{F}', P(\mathbf{U}, \mathbf{U}'_{\mathbf{X}}) \rangle$, where $\mathbf{U}'_{\mathbf{X}} = \bigcup_{X \in \mathbf{X}} \mathbf{U}'_X$ and $\mathcal{F}' = (\mathcal{F} \setminus \{f_x\}_{X \in \mathbf{X}}) \cup \{f'_x\}_{X \in \mathbf{X}}$. Hereafter, we will only consider interventions that result in a recursive $\mathcal{M}_{\sigma_{\mathbf{X}}}$. The model $\mathcal{M}_{\sigma_{\mathbf{X}}}$ induces a probability distribution

$$P(\mathbf{v}; \sigma_{\mathbf{X}}) = \sum_{\mathbf{u}, \mathbf{u}'_{\mathbf{X}}} \prod_{\{i|V_i \in \mathbf{X}\}} P(v_i | \mathbf{pa}_i, \mathbf{u}_i, \mathbf{u}'_i; \sigma_{\mathbf{X}}) P(\mathbf{u}'_{\mathbf{X}}; \sigma_{\mathbf{X}}) \prod_{\{i|V_i \in \mathbf{V} \setminus \mathbf{X}\}} P(v_i | \mathbf{pa}_i, \mathbf{u}_i) P(\mathbf{u}), \quad (2)$$

and a causal graph $\mathcal{G}_{\sigma_{\mathbf{X}}}$. Here the notation $P(\cdot; \sigma_{\mathbf{X}})$ is used to refer to the probability distribution under $\sigma_{\mathbf{X}}$ similar to the way $P(\cdot | do(\mathbf{x}))$ is used in the atomic case. Following the representation in [8] we annotate $\mathcal{G}_{\sigma_{\mathbf{X}}}$ with extra nodes called *regime indicators* [10] to mark the variables in $\mathcal{M}_{\sigma_{\mathbf{X}}}$ whose functions have been replaced relative to $\mathcal{M}$. Specifically, $\mathcal{G}_{\sigma_{\mathbf{X}}}$ contains a node σ_{X_i} for every $X_i \in \mathbf{X}$ together with an edge ($\sigma_{X_i} \rightarrow X_i$).

The proofs are provided in the Appendix (Supplemental Material).




Figure 2: Diagram (a) represents the natural regime in π^* and (b) the regime after the target intervention. The selection diagram (c) compares π^* with π^1 and π^2 , while (d) and (e) are selection diagrams specific to domains π^1 and π^2 under interventions σ_Z^1 and σ_W^2 , respectively.

2 Transporting Causal Relations across Domains

Our goal is to predict the effect of a soft intervention σ_X using assumptions encoded in the form of a graph and different observational and/or experimental distributions arising from different domains. Let $\Pi = \{\pi^*, \pi^1, \pi^2, \dots\}$ be the set of domains/populations involved in the analysis where π^* has a special status as the *target* domain where the query is to be inferred.

To better understand this setting, let us consider an example of a hypothetical study of government-backed loan programs and successful payment of family house purchases (inspired on studies such as [6, 17]). In city π^* , the percentage of the value of the property that can be borrowed (X) depends mainly on the credit history of the applicant (W) and his current employment conditions. The amount borrowed (or principal), Z , depends on the allowed percentage and the characteristics (R) of the property — mainly cost and location — as well as financial aspects of the borrower. The goal of the policy is to increase the number of families purchasing their own homes and paying their mortgages, so Y represents an indicator of the progress of the payment after a certain amount of years. Fig. 2(a) represents this situation in π^* . Let π^1 be another city where the distributions of credit history W is different than in π^* , and π^2 another where the distribution of R is the one that differs. These differences are called “domain discrepancies”, formally defined next.

Definition 1 (Domain Discrepancy). Let π^a and π^b be domains associated, respectively, with SCMs $\mathcal{M}^a$ and $\mathcal{M}^b$ conforming to a causal diagram $\mathcal{G}$. We denote by $\Delta^{a,b} \subseteq \mathbf{V}$ a set of variables such that, for every $V_i \in \Delta^{a,b}$, there might exist a discrepancy if $f_i^a \neq f_i^b$ or $P^a(\mathbf{U}_i) \neq P^b(\mathbf{U}_i)$.

We will write $\Delta^{*,i}$ simply as Δ^i to represent the differences between the target and each source domain, with $\Delta^* = \emptyset$. In this example, $\Delta^1 = \{W\}$ and $\Delta^2 = \{R\}$, which can be encoded in a *selection diagram* $\mathcal{G}^\Delta$ shown in Fig. 2(c), which according to the following definition:

Definition 2 (Selection Diagram). Given a causal diagram $\mathcal{G}^i = \langle \mathbf{V}, \mathbf{E} \rangle$ and domain discrepancies Δ^i , let $\mathbf{S} = \{S_v \mid \exists_{i=1}^n V \in \Delta^i\}$ be called *selection variables*. Then, a *selection diagram* $\mathcal{G}^{\Delta^i}$ is defined as a graph $\langle \mathbf{V} \cup \mathbf{S}, \mathbf{E} \cup \{S_v \rightarrow V\}_{S_v \in \mathbf{S}} \rangle$.

Policy makers in city π^* are planning to increase the percentage X by offering loaners insurance from default as long as the percentage ought is above the regular threshold, for properties within certain locations (accounted in R). We denote this policy as σ_X^* and represent the resulting regime in Fig. 2(b). To assess the impact of such policy, data of the current mortgage payments in π^* have been collected together with data from other two cities, π^1 and π^2 . The administration of π^1 allocated loans where the amount Z was determined as a function of the allowed percentage (X) and the property (R), as depicted in Fig. 2(d). In city π^2 , a study selected borrowers at random and had loaners process their applications with randomized credit history W , to observe its effect on Y (Fig. 2(e)). Data available in each domain is specified by $\mathbb{Z} = \{\mathbb{Z}^i \mid \pi^i \in \Pi\}$, where each $\mathbb{Z}^i = \{\sigma_{Z_1}, \sigma_{Z_2}, \dots\}$, $\mathbb{Z}_j \subseteq \mathbf{V}$, corresponds to domain π^i . This means that distributions

$\{P^i(\mathbf{V}; \sigma_{\mathbf{Z}_j}) \mid \mathbf{Z}_j \in \mathbb{Z}^i\}_{\mathbb{Z}^i \in \mathbb{Z}}$ are assumed to be available. Notice that $P^i(\mathbf{V}; \sigma_\emptyset) = P^i(\mathbf{V})$ is a valid distribution and describes the observational (non-interventional) distribution in domain π^i . In this example, $\mathbb{Z} = \{\mathbb{Z}^* = \{\sigma_\emptyset\}, \mathbb{Z}^1 = \{\sigma_Z\}, \mathbb{Z}^2 = \{\sigma_W\}\}$.

The effect of the new policy σ_X^* can be measured by comparing $E[Y; \sigma_X^*]$ with $E[Y]$ for π^* . While the latter is readily estimable via $P^*(Y)$, which is part of the input; the challenge is to assess $P^*(Y; \sigma_X^*)$. The ability to infer the target effect from the input distributions is formalized next.

Definition 3 (Effect Transportability). Let $\mathbf{Y}, \mathbf{X}, \mathbf{W} \subset \mathbf{V}$ with $\mathbf{W} \cap \mathbf{Y} = \emptyset$. The (conditional) effect intervention σ_X on a set of outcome variables $\mathbf{Y}$, conditional on $\mathbf{W}$, $P^*(\mathbf{y}|\mathbf{w}; \sigma_X)$, in a target environment π^* , is said to be transportable from $\langle \mathcal{G}^\Delta, \mathbb{Z} \rangle$, if it is uniquely computable from the set of distributions $\mathbb{Z}$ for every assignment $(\mathbf{y}, \mathbf{w})$ and every set of models $\{\mathcal{M}^i\}_{\pi^i \in \Pi}$ inducing $\mathcal{G}^\Delta$ and $\mathbb{Z}$.

To solve this particular transportability instance, we use σ -calculus [9] and standard probability axioms.² For simplicity $\sigma_X = \sigma_X^*$, $\sigma_Z = \sigma_Z^1$ and $\sigma_W = \sigma_W^2$ are simply written as σ_X, σ_Z and σ_W .

$$P^*(y; \sigma_X) = \sum_{r, x, z} P^*(y|z, x, r; \sigma_X) P^*(z|x, r; \sigma_X) P^*(x|r; \sigma_X) P^*(r; \sigma_X). \quad (3)$$

By rule 3 and the separation $(R \perp\!\!\!\perp X)$ in $\mathcal{G}_{\sigma_X^* \bar{X}}$ and $\mathcal{G}_{\bar{X}}$, we have $P^*(r; \sigma_X) = P^*(r)$, which is estimable from the input distribution $P^*(\mathbf{V})$. The factor $P^*(x|r; \sigma_X)$ is determined by σ_X^* (and the policy's specification). For the second factor, $P^*(z|x, r; \sigma_X) = P^*(z|x, r; \sigma_X, \sigma_W)$ by rule 3 and $(Z \perp\!\!\!\perp W \mid X, R)$ in $\mathcal{G}_{\sigma_X \sigma_W \bar{W}}$ and $\mathcal{G}_{\sigma_X \bar{W}}$. Then, by rule 2 and $(Z \perp\!\!\!\perp X \mid R)$ in $\mathcal{G}_{\sigma_X \sigma_W \underline{X}}$ and $\mathcal{G}_{\sigma_W \underline{X}}$, we obtain $P^*(z|x, r; \sigma_W)$. From the graph $\mathcal{G}_{\sigma_W}^{\Delta_2}$ (Fig. 2(e)), $(Z \perp\!\!\!\perp S_r \mid R, X)$ hence $P^*(z|x, r; \sigma_W) = P^2(z|x, r; \sigma_W)$ estimable from the given $P^2(\mathbf{V}; \sigma_W = \sigma_W^2)$. The first factor is equal to $P^*(y|z, x, r; \sigma_X, \sigma_Z = z)$ by rule 2 and $(Y \perp\!\!\!\perp Z \mid X, R)$ in $\mathcal{G}_{\sigma_X \sigma_Z = z \underline{Z}}$ and $\mathcal{G}_{\sigma_X \underline{Z}}$. We remove the observed X by rule 1 and $(Y \perp\!\!\!\perp X \mid Z, R)$ in $\mathcal{G}_{\sigma_X \sigma_Z = z}$. Then, by rule 3 and $(Y \perp\!\!\!\perp X \mid Z, R)$ in $\mathcal{G}_{\sigma_X \sigma_Z = z \bar{X}}$ and $\mathcal{G}_{\sigma_Z = z \bar{X}}$, this is equal to $P^*(y|z, r; \sigma_Z = z)$ and transportable from π^1 since $(Y \perp\!\!\!\perp S_w \mid Z, R)$. We sum over X and use rule 2 with $(Y \perp\!\!\!\perp Z \mid X, R)$ in $\mathcal{G}_{\sigma_Z = z \underline{Z}}$ and $\mathcal{G}_{\sigma_Z \underline{Z}}$ to exchange $\sigma_Z = z$ with $\sigma_Z = \sigma_Z^1$ (see Appendix B.1 for the detailed derivation). Then

$$P^*(y; \sigma_X) = \sum_{r, x, z} \underbrace{\left(\sum_{x'} P^1(y|z, x', r; \sigma_Z) P^1(x'|r; \sigma_Z) \right)}_{\text{from } \sigma_Z^1 \text{ in } \pi^1} \underbrace{P^2(z|x, r; \sigma_W)}_{\text{from } \sigma_W^2 \text{ in } \pi^2} \underbrace{P^*(x|r; \sigma_X^*)}_{\text{def. } \sigma_X^*} \underbrace{P^*(r)}_{\text{from } \pi^*}. \quad (4)$$

To solve this kind of transportability instances, we can leverage the transportability theory for atomic interventions. Accordingly, we establish a crisp relationship between any instance with soft interventions and its atomic counterpart.

Theorem 1. Let $\mathbf{Y}, \mathbf{X} \subseteq \mathbf{V}$ be any two sets of variables, and let σ_X^* be an atomic, conditional or stochastic intervention. Then, the effect of σ_X^* on $\mathbf{Y}$ can be written as

$$P^*(\mathbf{y}; \sigma_X = \sigma_X^*) = \sum_{\mathbf{d} \setminus \mathbf{Y}} P^*(\mathbf{d} \setminus \mathbf{x}; \sigma_X = \mathbf{x}) \prod_{X \in \mathbf{X} \cap \mathbf{D}} P^*(x \mid \mathbf{pa}_x; \sigma_X = \sigma_X^*). \quad (5)$$

Moreover, this effect is transportable from $\langle \mathcal{G}^\Delta, \mathbb{Z} \rangle$ if and only if $P^*(\mathbf{d} \setminus \mathbf{x}; \sigma_X = \mathbf{x})$ is transportable from $\langle \mathcal{G}^\Delta, \mathbb{Z} \rangle$, where $\mathbf{D} = \text{An}(\mathbf{Y})_{\mathcal{G}_{\sigma_X^*}}$.

For the example just discussed,

$$P^*(y; \sigma_X) = \sum_{r, x, z} P^*(y, z, r; \sigma_X = x) P^*(x \mid r; \sigma_X = \sigma_X^*), \quad (6)$$

where the first factor can be transported with a similar derivation to the one described above.

Given the tightness of the reduction provided by Thm. 1, one may surmise that it is possible to blindly use existing transportability algorithms (e.g., GTR [22]) to solve for soft interventions. In particular, if one pretends the input consist of $do()$ -experiments, as expected by known algorithms, the resulting expression cannot be mapped directly in terms of the original soft-experiments (for further details, see Appendix B.3). This motivates the clean algorithmic approach undertaken in the next section.

²The rules of σ -calculus are provided in Appendix A.

3 A Complete Algorithm for Soft Transportability

Both input and output of the transportability task refer to probability distributions within different domains and for different interventions. A key building block of our algorithm is the ability to decompose such distributions in factors with the finest granularity possible. We use the concept of *C-factors* and *C-components* developed by Tian and Pearl in [36, 37]. The set $\mathbf{V}$ can be partitioned into *C-components* 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. For instance the graph in Fig. 1(a) induces two C-components $\{X_1, Z\}$ and $\{X_2, Y\}$.

For any $\mathbf{C} \subseteq \mathbf{V}$, the quantity $Q^k[\mathbf{C}; \sigma_{\mathbf{X}}](\mathbf{v})$ is the *C-factor* of $\mathbf{C}$ under intervention $\sigma_{\mathbf{X}}$ in domain π^k , and denotes the following function

$$Q^k[\mathbf{C}; \sigma_{\mathbf{X}}](\mathbf{v}) = \sum_{\mathbf{u}(\mathbf{C})} \prod_{\{i|V_i \in \mathbf{C}\}} P^k(v_i | \mathbf{pa}_i, \mathbf{u}_i; \sigma_{\mathbf{X}}) P^k(\mathbf{u}(\mathbf{C}); \sigma_{\mathbf{X}}), \quad (7)$$

where $\mathbf{U}(\mathbf{C}) = \bigcup_{V_i \in \mathbf{C}} \mathbf{U}_i$. When $\mathbf{C} = \mathbf{V}$, $Q^k[\mathbf{V}; \sigma_{\mathbf{X}}](\mathbf{v}) = P^k(\mathbf{v}; \sigma_{\mathbf{X}})$. For simplicity we write $Q^k[\mathbf{C}; \sigma_{\mathbf{X}}](\mathbf{v})$ as $Q^k[\mathbf{C}; \sigma_{\mathbf{X}}]$, $Q^k[\mathbf{C}; \sigma_{\mathbf{X}}]$ as ${}^kQ[\mathbf{C}]$ if $\sigma_{\mathbf{X}} = \emptyset$, $Q^k[\{V_i\}; \sigma_{\mathbf{X}}]$ just as $Q^k[V_i; \sigma_{\mathbf{X}}]$, and $Q[\mathbf{C}]$ when talking about C-factors in general, for any domain and intervention.

A C-factor $Q[\mathbf{C}]$ can be further factorized according to the C-component structure of $\mathcal{G}_{[\mathbf{C}]}$, the subgraph with only variables in $\mathbf{C}$, as stated in the following lemma from [37].

Lemma 1. [*C-component decomposition*] Let $\mathbf{C} \subseteq \mathbf{V}$, $\mathbf{C}_1, \dots, \mathbf{C}_l$ the C-components of $\mathcal{G}_{[\mathbf{C}]}$. Then $Q[\mathbf{C}] = \prod_j Q[\mathbf{C}_j]$, and for any topological order $C_1 < C_2 < \dots < C_n$ of the variables in $\mathbf{C}$

$$Q[\mathbf{C}_j] = \prod_{\{C_i \in \mathbf{C}_j\}} \frac{Q[C_1, \dots, C_i]}{Q[C_1, \dots, C_{i-1}]}, \text{ where } Q[C_1, \dots, C_i] = \sum_{c_{i+1}, \dots, c_n} Q[\mathbf{C}]. \quad (8)$$

This result allow us to decompose C-factors as products of minimal C-factors. Once both input and output distributions factorized that way, the task is solvable if every factor of the output can mapped to factors in the input distributions. First, we discuss how to match C-factors across different regimes:

Lemma 2. Let $\mathbf{X}, \mathbf{Z} \subset \mathbf{V}$ be disjoint sets of variables, $\sigma_{\mathbf{X}}$ and $\sigma_{\mathbf{Z}}$ be any two interventions, and $\mathbf{C} \subseteq \mathbf{V}$. Then, $Q[\mathbf{C}; \sigma_{\mathbf{X}}, \sigma_{\mathbf{Z}}] = Q[\mathbf{C}; \sigma_{\mathbf{X}}]$ whenever $\mathbf{C} \cap \mathbf{Z} = \emptyset$.

For instance, the C-factor $Q^*[Z, Y]$ in Fig. 1(a) is equal to $Q^*[Z, Y; \sigma_{X_1, X_2}]$ because $\{Z, Y\}$ does not intersect $\{X_1, X_2\}$. Similarly, in Fig. 2(d), (e), $Q^1[R, Y; \sigma_Z^1] = Q^1[R, Y; \sigma_Z^1, \sigma_W^2] = Q^1[R, Y]$ because $\{R, Y\}$ is not affected by σ_Z^1 or σ_W^2 . Next, we consider matching C-factors across domains:

Lemma 3. Let $\mathcal{G}^\Delta$ be a selection diagram for $\langle M^k, M^l \rangle$, then $Q^k[\mathbf{C}; \sigma_{\mathbf{X}}] = Q^l[\mathbf{C}; \sigma_{\mathbf{X}}]$ if $\mathcal{G}^\Delta$ does not contain selection nodes S_{v_i} pointing to any variable in $V_i \in \mathbf{C}$, that is, $V_i \notin \Delta^{k,l}$.

From the selection diagram in Fig. 2(c), we can infer $Q^*[X, Z, Y] = Q^1[X, Z, Y] = Q^2[X, Z, Y]$, and $Q^*[W] = Q^2[W]$, but not $Q^*[W] = Q^1[W]$, for example. Next, a query of interest can be written in terms of C-factors based on Thm. 1:

Corollary 1. Let $\mathbf{Y}, \mathbf{X} \subseteq \mathbf{V}$ be any two sets of variables, and let $\sigma_{\mathbf{X}} = \sigma_{\mathbf{X}}^*$ be an atomic, conditional or stochastic intervention. The effect of $\sigma_{\mathbf{X}}$ on $\mathbf{Y}$ is given by

$$P^*(\mathbf{y}; \sigma_{\mathbf{X}} = \sigma_{\mathbf{X}}^*) = \sum_{\mathbf{d} \setminus \mathbf{y}} Q^*[\mathbf{X} \cap \mathbf{D}; \sigma_{\mathbf{X}} = \sigma_{\mathbf{X}}^*] Q^*[\mathbf{D} \setminus \mathbf{X}], \quad (9)$$

where $\mathbf{D} = An(\mathbf{Y})_{\mathcal{G}_{\sigma_{\mathbf{X}}}}$. Furthermore, this effect is transportable from $\langle \mathcal{G}^\Delta, \mathbb{Z} \rangle$ if and only if $Q^*[\mathbf{D} \setminus \mathbf{X}]$ is transportable from $\langle \mathcal{G}^\Delta, \mathbb{Z} \rangle$.

The effect of an intervention conditioned on some evidence $\mathbf{W} = \mathbf{w}$, $P(\mathbf{y} | \mathbf{w}; \sigma_{\mathbf{X}})$, may involve less C-factors than $P(\mathbf{y}, \mathbf{w}; \sigma_{\mathbf{X}})$ depending on the topology of the graph. The exact factorization of the conditional query is given by the following theorem.

Theorem 2. Let $\mathbf{Y}, \mathbf{X}, \mathbf{W} \subset \mathbf{V}$, $\mathbf{W} \cap \mathbf{Y} = \emptyset$, $\sigma_{\mathbf{X}}$ be any intervention, and $\mathcal{G}_{\sigma_{\mathbf{X}}}$ the corresponding interventional causal graph. Then, the effect of $\sigma_{\mathbf{X}}$ on $\mathbf{Y}$ conditioned on $\mathbf{W}$ is given by

$$P(\mathbf{y} | \mathbf{w}; \sigma_{\mathbf{X}}) = P(\mathbf{y} | \mathbf{w}_{\mathbf{y}}; \sigma_{\mathbf{X}}, \sigma_{\mathbf{W}_{\bar{\mathbf{y}}} = \mathbf{w}_{\bar{\mathbf{y}}}}) = \sum_{\mathbf{a} \setminus (\mathbf{y} \cup \mathbf{w}_{\mathbf{y}})} Q[\mathbf{A}; \sigma_{\mathbf{X}}] / \sum_{\mathbf{a} \setminus \mathbf{w}_{\mathbf{y}}} Q[\mathbf{A}; \sigma_{\mathbf{X}}], \quad (10)$$

where $\mathbf{W}_{\mathbf{y}} \subseteq \mathbf{W}$ is the set of variables in $\mathbf{W}$ connected to any $Y \in \mathbf{Y}$ by any path (regardless of the directionality) in $\mathcal{G}_{\sigma_{\mathbf{X}}[\mathbf{D}]\mathbf{W}}$, with $\mathbf{D} = An(\mathbf{Y} \cup \mathbf{W})_{\mathcal{G}_{\sigma_{\mathbf{X}}}}$, $\mathbf{W}_{\bar{\mathbf{y}}} = \mathbf{W} \setminus \mathbf{W}_{\mathbf{y}}$, and $\mathbf{A} = An(\mathbf{Y} \cup \mathbf{W}_{\mathbf{y}})_{\mathcal{G}_{\sigma_{\mathbf{X}}\mathbf{W}}}$. Furthermore, this effect is transportable from $\langle \mathcal{G}^\Delta, \mathbb{Z} \rangle$ iff $Q[\mathbf{A}; \sigma_{\mathbf{X}}]$ is transportable from $\langle \mathcal{G}^\Delta, \mathbb{Z} \rangle$.

Algorithm 1 σ -TR($\mathbf{Y}, \mathbf{W}, \sigma_{\mathbf{X}}, \mathbb{Z}, \mathcal{G}^{\Delta}$)

Input: $\mathcal{G}^{\Delta}$ selection diagrams over variables $\mathbf{V}$ for domains Π ; $\mathbf{Y}, \mathbf{W} \subseteq \mathbf{V}$ disjoint subsets of variables; an intervention $\sigma_{\mathbf{X}}^*$ defined over a set $\mathbf{X} \subseteq \mathbf{V}$; and available distribution specification $\mathbb{Z}$.

Output: $P^*(\mathbf{y}|\mathbf{w}; \sigma_{\mathbf{X}})$ in terms of available distributions or FAIL if not transportable from $\langle \mathcal{G}^{\Delta}, \mathbb{Z} \rangle$.

```
1: let  $\mathbf{A}$  be defined as in Thm. 2, and let  $\mathbf{A}_1, \dots, \mathbf{A}_n$  be the C-components of  $\mathcal{G}_{\sigma_{\mathbf{X}}[\mathbf{A}]}$ .
2: for each  $\mathbf{A}_i$  s.t.  $\mathbf{A}_i \cap \mathbf{X} = \emptyset$ ,  $\sigma_{\mathbf{Z}} \in \mathbb{Z}^k \in \mathbb{Z}$  s.t.  $\mathbf{A}_i \cap \mathbf{Z} = \emptyset$  and  $\mathbf{A}_i \cap \Delta^k = \emptyset$  do
3:   let  $\mathbf{B}_i$  be the C-component of  $\mathcal{G}_{\sigma_{\mathbf{Z}}}$  such that  $\mathbf{A}_i \subseteq \mathbf{B}_i$ , compute  $Q^k[\mathbf{B}_i; \sigma_{\mathbf{Z}}]$  from  $Q^k[\mathbf{V}; \sigma_{\mathbf{Z}}]$ .
4:   if IDENTIFY( $\mathbf{A}_i, \mathbf{B}_i, Q[\mathbf{B}_i; \sigma_{\mathbf{Z}}], \mathcal{G}_{\sigma_{\mathbf{Z}}}$ ) does not FAIL then
5:     let  $Q^*[\mathbf{A}_i; \sigma_{\mathbf{X}}] = \text{IDENTIFY}(\mathbf{A}_i, \mathbf{B}_i, Q[\mathbf{B}_i; \sigma_{\mathbf{Z}}], \mathcal{G}_{\sigma_{\mathbf{Z}}})$ .
6:     move to the next  $\mathbf{A}_i$ .
7:   end if
8: end for
9: for each  $\mathbf{A}_i$  containing variables in  $\mathbf{X}$  let  $Q^*[\mathbf{A}_i; \sigma_{\mathbf{X}}] = \text{REPLACE}(\mathbf{A}_i, \sigma_{\mathbf{X}})$ .
10: if any  $Q^*[\mathbf{A}_i]$  was not assigned then return FAIL.
11: let  $Q^*[\mathbf{A}; \sigma_{\mathbf{X}}] = \prod_i Q^*[\mathbf{A}_i; \sigma_{\mathbf{X}}]$ .
12: return  $\sum_{\mathbf{a} \setminus (\mathbf{Y} \cup \mathbf{W})} Q^*[\mathbf{A}; \sigma_{\mathbf{X}}] / \sum_{\mathbf{a} \setminus \mathbf{W}} Q^*[\mathbf{A}; \sigma_{\mathbf{X}}]$ .
```

For instance, the effect $P^*(y | r, z; \sigma_X = \sigma_X^*)$ in the example of Fig. 2(b) can be transported as

$$P^*(y | r, z; \sigma_X) = P^*(y | r; \sigma_X, \sigma_Z = z) = Q^*[R, Y; \sigma_X] / \sum_y Q^*[R, Y; \sigma_X]. \quad (11)$$

Using Lemma 3 we have $Q^*[R, Y; \sigma_X] = Q^1[R, Y; \sigma_X]$ and by Lemma 2 equal to $Q^1[R, Y; \sigma_Z]$. From $Q^1[\mathbf{V}; \sigma_Z] = P^1(\mathbf{v}; \sigma_Z)$ we obtain $Q^1[R, Y, W, X; \sigma_Z]$, using Lemma 1, as $Q^1[\mathbf{V}; \sigma_Z] / Q^1[\mathbf{Z}; \sigma_Z] = P^1(\mathbf{v}; \sigma_Z) / P^1(z | w, r, x; \sigma_Z) = P^1(y | z, x, r, w; \sigma_Z) P^1(x, r, w; \sigma_Z)$. Then, since $R < Y < W < X$ is a valid topological order in $\mathcal{G}_{[R, Y, W, X]}$, Eq. (8) leads to $Q^1[R, Y; \sigma_Z] = \sum_{x, w} Q^1[R, Y, W, X] = \sum_{x, w} P^1(y | z, x, r, w; \sigma_Z) P^1(x, r, w; \sigma_Z)$. Replacing in Eq. (11) and simplifying results in:

$$P^*(y | r, z; \sigma_X) = \sum_{x, w} P^1(y | z, x, r, w; \sigma_Z) P^1(x, w | r; \sigma_Z). \quad (12)$$

Building on the observations and results we have so far, we design the algorithm σ -TR (Alg. 1) that takes as input the variables defining a query, the specification of $\sigma_{\mathbf{X}}$ (i.e., what type of intervention is being applied and its arguments), a set of available distributions ($\mathbb{Z}$), and the selection diagrams. Internally σ -TR uses the subroutine IDENTIFY from [36] that applies Lemma 1 systematically to obtain a C-factor $Q[\mathbf{A}]$ from $Q[\mathbf{B}]$, where $\mathbf{A} \subseteq \mathbf{B}$, and the subroutine ‘REPLACE’ to determine the factors of intervened variables according to the particular type of intervention.

Recall the example from the introduction of the paper where the target effect is $P^*(y; \sigma_{X_1, X_2} = \sigma^*)$ corresponding to the regime shown in Fig. 1(d). With $\mathbf{Y} = \{Y\}$, $\mathbf{W} = \emptyset$, $\mathbb{Z} = \{\{\sigma_{\emptyset}\}, \{\sigma_{X_1, X_2}^1\}, \{\sigma_{X_2}^2\}\}$ and $\mathcal{G}^{\Delta}$ as in Figs. 1(b), (c); σ -TR is invoked so that $\mathbf{A} = \{X_1, Z, X_2, Y\}$ with C-components $\mathbf{A}_1 = \{X_1\}$, $\mathbf{A}_2 = \{Z\}$, $\mathbf{A}_3 = \{X_2\}$ and $\mathbf{A}_4 = \{Y\}$ ($\mathcal{G}_{\sigma^*[\mathbf{A}]} = \mathcal{G}_{\sigma^*}$). The loop in line 2 will take $\mathbf{A}_2$ for domain π^k , $k=1$ and $\sigma_{\mathbf{Z}} = \sigma^1$. $\mathbf{B}_i = \{Z\}$ is the C-component of $\mathcal{G}_{\sigma^1}$ (Fig. 1(b)) that contains $\mathbf{A}_2$, and $Q^1[\mathbf{Z}; \sigma^1] = P^1(z | x_1; \sigma^1)$ by Lemma 1. Then IDENTIFY is called and returns $Q^*[\mathbf{A}_2] = P^1(z | x_1; \sigma^1)$. Similarly, $Q^*[\mathbf{A}_4]$ is obtained from π^2 with $\mathbf{B}_i = \{Y\}$ as $Q^2[Y; \sigma^2] = P^2(y | x_1, x_2; \sigma^2)$. Next, $Q^*[X_1; \sigma^*]$ and $Q^*[X_2; \sigma^*]$ are given by policy σ^* as $P^*(x_1; \sigma^*)$ and $P^*(x_2 | z, x_1; \sigma^*)$. Finally, line 12 returns an expression equivalent to Eq. (1).

Determining non-transportability is non-trivial task, however, σ -TR is both sufficient and necessary:

Theorem 3. [σ -TR Completeness] *The effect $P(\mathbf{y} | \mathbf{w}; \sigma_{\mathbf{X}})$ is transportable from $\mathbb{Z}$ in $\mathcal{G}^{\Delta}$ if and only if the algorithm σ -TR (Alg. 1) outputs an estimand for it.*

σ -TR finds an estimand for the query or determines non-transportability efficiently, in $O(n^4 z)$ time, where $n = |\mathbf{V}|$ and $z = \sum_{\pi^i} |\mathbb{Z}^i|$ (see Appendix B.1). Moreover, every step in σ -TR can be mapped to a derivation with σ -calculus and probability axioms like the one given in the previous section.

Corollary 2. [σ -calculus Completeness] *The σ -calculus together with standard probability axioms is complete for the task of transportability with soft interventions.*

4 Graphical Characterization

We proved that the procedure σ -TR, discussed in the previous section, is both sufficient and necessary for transportability task at hand. In this section, we investigate a sufficient and necessary graphical conditions to decide the transportability of a query by graphical inspection.

Consider once more the loan program example in Section 2. If we learn that R also differs from those in π^* (i.e., R is also in Δ^1), then the target query is not transportable from the given input distributions. We can examine the reason σ -TR outputs FAIL for the query $P^*(y; \sigma_X)$. On line 1, since $\mathbf{W}=\emptyset$, $\mathbf{A}=\{R, X, Z\}$ with C-factors $\mathbf{A}_1=\{R, Y\}$, $\mathbf{A}_2=\{Z\}$, and $\mathbf{A}_3=\{X\}$. There are three sources of data in $\mathbb{Z}$, namely, $\mathbb{Z}^*=\{\sigma_\emptyset\}$, $\mathbb{Z}^1=\{\sigma_Z^1\}$ and $\mathbb{Z}^2=\{\sigma_W^2\}$. If $R \in \Delta^1, \Delta^2$, no C-factor including R can be transported from $\mathbb{Z}^1$ or $\mathbb{Z}^2$. Then, the loop in line 2 only runs for $\sigma_\emptyset \in \mathbb{Z}^*$ such that $\mathbf{B}_i=\{W, R, X, Z, Y\}$, and IDENTIFY fails to obtain $Q^*[R, Y]$ from $Q^*[\mathbf{B}_i]$. At line 10, σ -TR detects that $Q^*[R, Y]$ was not transported and outputs FAIL.




Figure 3: Selection diagram depicting intervention σ_X^* .

Overall, σ -TR fails when there exists some C-component $\mathbf{A}_i$ of $\mathcal{G}_{\sigma_X[\mathbf{A}]}$, a C-component $\mathbf{C}_i$ of $\mathcal{G}_{\Delta_{[An(\mathbf{A}_i)]}}$, and for every $\sigma_Z \in \mathbb{Z}^k \in \mathbb{Z}$ at least one of the following three conditions occur:

- (i) $\mathbf{A}_i \cap \Delta^k \neq \emptyset$, that is, at least one variable in $\mathbf{A}_i$ has a different mechanism in π^k , or
- (ii) $\mathbf{A}_i \cap \mathbb{Z} \neq \emptyset$, meaning at least one of the variables in $\mathbf{A}_i$ have been intervened by σ_Z in π^k , or
- (iii) there exists some $\mathbf{T}_i$ s.t. $\mathbf{A}_i \subset \mathbf{T}_i \subseteq \mathbf{C}_i$, $\mathcal{G}_{\sigma_Z[\mathbf{T}_i]}$ has a single C-component.

Condition (iii) occurs whenever IDENTIFY (in line 4) fails to obtain $Q^*[\mathbf{A}_i]=Q^k[\mathbf{A}_i; \sigma_Z]$ from $Q^k[\mathbf{B}_i; \sigma_Z]$ in $\mathcal{G}_{\sigma_Z}$ [16, Thm. 3]. The following result follows immediately from Thm. 3:

Corollary 3. *Given query $P^*(y; \sigma_X=\sigma_X^*)$, selection diagram $\mathcal{G}^\Delta$, and the distribution specified by $\mathbb{Z}$, let $\mathbf{A}$ be defined as in Thm. 2. Then, the query is not transportable from $\langle \mathcal{G}^\Delta, \mathbb{Z} \rangle$ if and only if there exists a C-component $\mathbf{A}_i$ of $\mathcal{G}_{\sigma_X^*[\mathbf{A}]}$ and a C-component $\mathbf{C}_i$ of $\mathcal{G}_{\Delta_{[An(\mathbf{A}_i)]}}$ such that for every $\sigma_Z \in \mathbb{Z}^k \in \mathbb{Z}$, $\mathbf{A}_i$ satisfies conditions (i), (ii), or (iii).*

For instance consider $\mathcal{G}^\Delta$ in Fig. 3 with $\mathbb{Z}=\{\{\sigma_\emptyset\}, \{\sigma_X=x\}, \{\sigma_W=g(R), \sigma_R=P'(R)\}\}$ and intervention $\sigma_X=\sigma_X^*$, where σ_X^* enforces some pre-specified $P^*(X|R, W; \sigma_X^*)$. The edge $R \rightarrow X$ is added due to σ_X^* while the two bidirected edges in grey are removed by it. Let $\mathbf{A}_i=\{Z, R\}$ and $\mathbf{C}_i=\{R, X, W, Z\}$. We can verify that condition (iii) is satisfied for $\sigma_\emptyset$ from π^* and σ_W from π^2 . From π^1 , $\sigma_X=x$ cannot be used due to condition (i) and σ_R in π^3 satisfies condition (ii). Although for this instance there is no *s-Thicket* [22] —a graphical structure precluding transportability in the atomic case— for $P(y \mid do(x))$, the effect of σ_X^* is not transportable. Instead, the condition in Cor. 3 can be proved equivalent to the presence of an *s-Thicket* for the effect $P(\mathbf{a} \mid do(\mathbf{x}, an(\mathbf{a}) \setminus \mathbf{a}))$ (see Lemma 10 in Appendix D).

5 Conclusions

We studied the problem of transporting effects of soft interventions from knowledge encoded in the form of a selection diagram and a combination of observational and experimental data from multiple, different domains. We showed how the problem can be solved by transporting the effect of an atomic intervention from the same input (Thm. 1). Similarly, we proved that a conditional effect is transportable only if the marginal effect derived through Thm. 2 is transportable. We then developed an efficient algorithm called σ -TR (Alg. 1) that is both sufficient and necessary for finding a function of the available data that matches the target effect (whenever such a function exists). As a corollary, we conclude that σ -calculus together with basic probability axioms are complete for the soft transportability task (Cor. 2). Finally, we described a complete graphical condition to determine the transportability of any transportability instance (Cor. 3). We hope that this series of results related to soft interventions, and knowledge of its relationship with atomic interventions, can help data scientists to apply causal inference in broader and more realistic scenarios.

Broader Impact

Our work investigates the formal conditions under which knowledge acquired in one domain (e.g., setting, population, environment) can be generalized to a different one that may be related, but is unlikely to be the same. This is known in the causal inference literature as the problem of "transportability." As alluded to in the introduction, issues of transportability are pervasive throughout the empirical sciences as well as AI and ML. We believe that having a more foundational tool that allows the empirical investigator to determine whether (and how) her/his understanding of the underlying system is sufficient to support the generalization of an empirical claim is a critical addition to the scientific toolbox. Not having such a tool, on the other hand, may lead researchers to operate on a more heuristical basis, which may lead to a lack of understanding of when things can go wrong and how to fix them. For instance, public policies that will fail or have unintended consequences, potentially harming people, or spending unnecessary societal resources. In the context of automated decision-making in AI, we could have systems following policies that harm the users, can act unfairly, or discriminate against certain groups (e.g., the policy was trained in Scandinavia and moved to the US). By and large, we believe this research on the theory of generalization of policies based on soft interventions can benefit a large group of individuals, including empirical scientists, policy-makers, AI researchers, and society in general.

Acknowledgments and Disclosure of Funding

This research was supported by grants from IBM Research, Adobe Research, NSF IIS-1704352, and IIS-1750807 (CAREER).

References

- [1] A. V. Banerjee, S. Cole, E. Duflo, and L. Linden. Remedying Education: Evidence from Two Randomized Experiments in India*. *The Quarterly Journal of Economics*, 122(3):1235–1264, 2007.
- [2] E. Bareinboim and J. Pearl. Transportability of causal effects: Completeness Results. In *Proceedings of the 26th AAAI Conference on Artificial Intelligence*, CA, 2012. Department of Computer Science, University of California, Los Angeles.
- [3] E. Bareinboim and J. Pearl. Causal Inference by Surrogate Experiments: z-Identifiability. In N. d. F. Murphy and Kevin, editors, *Proceedings of the Twenty-Eighth Conference on Uncertainty in Artificial Intelligence*, pages 113–120. AUAI Press, 2012.
- [4] E. Bareinboim and J. Pearl. A general algorithm for deciding transportability of experimental results. *Journal of Causal Inference*, 1(1):107–134, 2013.
- [5] E. Bareinboim and J. Pearl. Transportability from Multiple Environments with Limited Experiments: Completeness Results. In Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 27*, pages 280–288. Curran Associates, Inc., 2014.
- [6] M. Bertrand, D. Karlan, S. Mullainathan, E. Shafir, and J. Zinman. What’s Advertising Content Worth? Evidence from a Consumer Credit Marketing Field Experiment*. *The Quarterly Journal of Economics*, 125(1):263–306, 2 2010.
- [7] N. Cartwright. *Hunting Causes and Using Them: {A}pproaches in Philosophy and Economics*. Cambridge University Press, New York, NY, 2007.
- [8] J. D. Correa and E. Bareinboim. From Statistical Transportability to Estimating the Effects of Stochastic Interventions. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI-19*, pages 1661–1667. International Joint Conferences on Artificial Intelligence Organization, 2019.
- [9] J. D. Correa and E. Bareinboim. A Calculus For Stochastic Interventions: Causal Effect Identification and Surrogate Experiments. In *Proceedings of the 34th AAAI Conference on Artificial Intelligence*. AAAI Press, 2020.
- [10] A. P. Dawid. Influence diagrams for causal modelling and inference. *International Statistical Review*, 70:161–189, 2002.

- [11] A. P. Dawid, V. Didelez, and others. Identifying the consequences of dynamic treatment strategies: A decision-theoretic overview. *Statistics Surveys*, 4:184–231, 2010.
- [12] E. Duflo, R. Glennerster, and M. Kremer. Using Randomization in Development Economics Research: A Toolkit. In T. P. Schultz and J. A. Strauss, editors, *Handbook of Development Economics*, volume 4, chapter 61, pages 3895–3962. Elsevier, 2007.
- [13] F. Eberhardt and R. Scheines. Interventions and causal inference. *Philosophy of science*, 74(5): 981–995, 2007.
- [14] J. J. Heckman. The Scientific Model of Causality. *Sociological Methodology*, 35:1–97, 2005.
- [15] Y. Huang and M. Valtorta. 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, 2006.
- [16] Y. Huang and M. Valtorta. On the completeness of an identifiability algorithm for semi-Markovian models. *Annals of Mathematics and Artificial Intelligence*, 54(4):363–408, 2008.
- [17] D. Karlan and J. Zinman. Expanding credit access: Using randomized supply decisions to estimate the impacts. *The Review of Financial Studies*, 23(1):433–464, 2010.
- [18] K. Korb, L. Hope, A. Nicholson, and K. Axnick. Varieties of Causal Intervention. In *Lecture Notes in Artificial Intelligence (Subseries of Lecture Notes in Computer Science)*, volume 3157, pages 322–331, 2004.
- [19] S. Lee and V. Honavar. Causal Transportability of Experiments on Controllable Subsets of Variables: z-Transportability. In A. Nicholson and P. Smyth, editors, *Proceedings of the Twenty-Ninth Conference on Uncertainty in Artificial Intelligence (UAI)*, pages 361–370. AUAI Press, 2013.
- [20] S. Lee and V. Honavar. m-Transportability: Transportability of a Causal Effect from Multiple Environments. In M. desJardins and M. Littman, editors, *Proceedings of the Twenty-Seventh National Conference on Artificial Intelligence*, pages 583–590, Menlo Park, CA, 2013. AAAI Press.
- [21] S. Lee, J. D. Correa, and E. Bareinboim. General Identifiability with Arbitrary Surrogate Experiments. In *Proceedings of the Thirty-Fifth Conference Annual Conference on Uncertainty in Artificial Intelligence*, Corvallis, OR, 2019. AUAI Press, in press.
- [22] S. Lee, J. Correa, and E. Bareinboim. Generalized Transportability: Synthesis of Experiments from Heterogeneous Domains. In *Proceedings of the 34th AAAI Conference on Artificial Intelligence*, Menlo Park, CA, 2020. AAAI Press.
- [23] J. K. Lunceford, M. Davidian, and A. A. Tsiatis. Estimation of Survival Distributions of Treatment Policies in Two-Stage Randomization Designs in Clinical Trials. *Biometrics*, 58(1): 48–57, 2002.
- [24] S. A. Murphy. An experimental design for the development of adaptive treatment strategies. *Statistics in Medicine*, 24(10):1455–1481, 2005.
- [25] J. Pearl. Causal diagrams for empirical research. *Biometrika*, 82(4):669–688, 1995.
- [26] J. Pearl. *Causality: Models, Reasoning, and Inference*. Cambridge University Press, New York, NY, USA, 2nd edition, 2000.
- [27] J. Pearl. Review of N. Cartwright ‘Hunting Causes and Using Them’. *Economics and Philosophy*, 26:69–77, 2010.
- [28] J. Pearl and E. Bareinboim. Transportability of Causal and Statistical Relations: A Formal Approach. In *Proceedings of the Twenty-Fifth Conference on Artificial Intelligence (AAAI-11)*, pages 247–254, Menlo Park, CA, 8 2011.
- [29] J. Pearl and D. Mackenzie. *The Book of Why*. Basic Books, New York, 2018.
- [30] J. Pearl and J. M. Robins. Probabilistic evaluation of sequential plans from causal models with hidden variables. In P. Besnard and S. Hanks, editors, *Uncertainty in Artificial Intelligence 11*, pages 444–453. Morgan Kaufmann, San Francisco, 1995.
- [31] I. Shpitser and J. Pearl. 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, 2006.

- [32] I. Shpitser and E. Sherman. Identification of Personalized Effects Associated With Causal Pathways. In *Proceedings of the 34th Conference on Uncertainty in Artificial Intelligence*, pages 530–539, 2018.
- [33] P. Spirtes, C. N. Glymour, and R. Scheines. *Causation, Prediction, and Search*. MIT Press, 2nd edition, 2001.
- [34] J. Tian. Identifying Conditional Causal Effects. In *Proceedings of the 20th Conference on Uncertainty in Artificial Intelligence*, UAI '04, pages 561–568, Arlington, Virginia, United States, 2004. AUAI Press.
- [35] J. Tian. Identifying Dynamic Sequential Plans. In *In Proceedings of the Twenty-Fourth Conference Annual Conference on Uncertainty in Artificial Intelligence (UAI-08)*, page 554–561, Corvallis, Oregon, 2008. AUAI Press.
- [36] J. Tian and J. Pearl. 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, 2002. AAAI Press/The MIT Press.
- [37] J. Tian and J. Pearl. 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, 2002.
- [38] J. Woodward. *Making Things Happen*. Oxford University Press, New York, NY, 2003.