

Transportability of Causal Effects: Completeness Results*

Elias Bareinboim and Judea Pearl

Cognitive Systems Laboratory
Department of Computer Science
University of California, Los Angeles
Los Angeles, CA. 90095
{eb,judea} at cs.ucla.edu

Abstract

The study of transportability aims to identify conditions under which causal information learned from experiments can be reused in a different environment where only passive observations can be collected. The theory introduced in [Pearl and Bareinboim, 2011] (henceforth [PB, 2011]) defines formal conditions for such transfer but falls short of providing an effective procedure for deciding, given assumptions about differences between the source and target domains, whether transportability is feasible. This paper provides such procedure. It establishes a necessary and sufficient condition for deciding when causal effects in the target domain are estimable from both the statistical information available and the causal information transferred from the experiments. The paper further provides a complete algorithm for computing the transport formula, that is, a way of fusing experimental and observational information to synthesize an estimate of the desired causal relation.

Motivation

The problem of transporting knowledge from one environment to another has been pervasive in many data-driven sciences. Invariably, when experiments are performed on a group of subjects, the issue arises whether the conclusions are applicable to a different but somewhat related group. When a robot is trained in a simulated environment, the question arises whether it could put the acquired knowledge into use in a new environment where relationships among agents, objects and features are different.

Surprisingly, the conditions under which this extrapolation can be legitimized were not formally articulated. Although the problem has been discussed in many areas of statistics, economics, and the health sciences, under rubrics such as “external validity” [Campbell and Stanley, 1963; Manski, 2007], “meta-analysis” [Glass, 1976; Hedges and Olkin, 1985; Owen, 2009], “heterogeneity” [Hofler, Gloster, and Hoyer, 2010], “quasi-experiments” [Shadish, Cook, and Campbell, 2002, Ch. 3; Adelman, 1991], these discussions

are limited to verbal narratives in the form of heuristic guidelines for experimental researchers – no formal treatment of the problem has been attempted.

AI is in a unique position to tackle this problem formally. First, the distinction between statistical and causal knowledge has received syntactic representation through causal diagrams [Pearl, 1995; Spirtes, Glymour, and Scheines, 2001; Pearl, 2009; Koller and Friedman, 2009]. Second, graphical models provide a language for representing differences and commonalities among domains, environments, and populations [PB, 2011]. Finally, the inferential machinery provided by the do-calculus [Pearl, 1995; 2009; Koller and Friedman, 2009] is particularly suitable for combining these two features into a coherent framework and developing effective algorithms for knowledge transfer.

Following [PB, 2011], we consider transferring causal knowledge between two environments Π and Π^* . In environment Π , experiments can be performed and causal knowledge gathered. In Π^* , potentially different from Π , only passive observations can be collected but no experiments conducted. The problem is to infer a causal relationship R in Π^* using knowledge obtained in Π . Clearly, if nothing is known about the relationship between Π and Π^* , the problem is unsolvable. Yet the fact that all experiments are conducted with the intent of being used elsewhere (e.g., outside the laboratory) implies that scientific progress relies on the assumption that certain environments share common characteristics and that, owed to these commonalities, causal claims would be valid even where experiments were never performed.

To formally articulate commonalities and differences between environments, a graphical representation named *selection diagrams* was devised in [PB, 2011], which represent differences in the form of unobserved factors capable of causing such differences. Given an arbitrary selection diagram, our challenge is to algorithmically decide whether commonalities override differences to permit the transfer of information across the two environments.

Previous Work and Our Contributions

Consider Fig. 1(a) which concerns the transfer of experimental results between two locations. We first conduct a randomized trial in Los Angeles (LA) and estimate the causal effect of treatment X on outcome Y for every age group

*This research was supported in parts by grants from NIH #1R01 LM009961-01, NSF #IIS-0914211 and #IIS-1018922, and ONR #N000-14-09-1-0665 and #N00014-10-1-0933.
Copyright © 2012, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

$Z = z$, denoted $P(y|do(x), z)$. We now wish to generalize the results to the population of New York City (NYC), but we find the distribution $P(x, y, z)$ in LA to be different from the one in NYC (call the latter $P^*(x, y, z)$). In particular, the average age in NYC is significantly higher than that in LA. How are we to estimate the causal effect of X on Y in NYC, denoted $R = P^*(y|do(x))$?^{1,2}

The selection diagram for this example (Fig. 1(a)) conveys the assumption that the *only* difference between the two population are factors determining age distributions, shown as $S \rightarrow Z$, while age-specific effects $P(y|do(x), Z = z)$ are invariant across cities. Difference-generating factors are represented by a special set of variables called *selection variables* S (or simply S -variables), which are graphically depicted as square nodes (■). From this assumption, the overall causal effect in NYC can be derived as follows³

$$\begin{aligned} R &= \sum_z P^*(y|do(x), z)P^*(z) \\ &= \sum_z P(y|do(x), z)P^*(z) \end{aligned}$$

The last line is the *transport formula* for R . It combines experimental results obtained in LA, $P(y|do(x), z)$, with observational aspects of NYC population, $P^*(z)$, to obtain an experimental claim $P^*(y|do(x))$ about NYC.

In this trivial example the transport formula amounts to a simple re-calibration of the age-specific effects to account for the new age distribution. In more elaborate examples, however, the full power of formal analysis would be required. For instance, [PB, 2011] showed that, in the problem depicted in Fig. 1(b), where both the Z -determining mechanism and the U -determining mechanism are suspect of being different, the transport formula for the relation $R = P^*(y|do(x))$ is given by

$$\sum_z P(y|do(x), z) \sum_w P^*(z|w) \sum_t P(w|do(x), t)P^*(t)$$

This formula instructs us to estimate $P(y|do(x), z)$ and $P(w|do(x), t)$ in the experimental domain, then combine them with the estimates of $P^*(z|w)$ and $P^*(t)$ in the target domain.

¹We will use $P_x(y)$ interchangeably with $P(y | do(x))$.

²We use the structural interpretation of causal diagrams. For example, Fig. 1(a) describes the following system of structural equations: $z \leftarrow f_1(s; u_z; u_{zx})$, $x \leftarrow f_2(z; u_x; u_{zx}; u_{xy})$, $y \leftarrow f_3(x; z; u_y; u_{xy})$; each variable in the l.h.s. is assigned a value given by the respective deterministic function on the r.h.s. The hidden (exogenous) variables U are assigned a probability function which induces in turn, a probability distribution on all variables in the model. See Appendix 1 for a gentle introduction to the *do-calculus* and more details on this representation.

³This result can be derived by purely graphical operations if we write $P^*(y|do(x), z)$ as $P(y|do(x), z, s)$, thus attributing the difference between Π and Π^* to a fictitious event $S = s$. The invariance of the age-specific effect then follows from the conditional independence $(S \perp\!\!\!\perp Y|Z, X)_{G_{\overline{X}}}$, which implies $P(y|do(x), z, s) = P(y|do(x), z)$, and licenses the derivation of the transport formula.




Figure 1: Selection diagrams illustrating different facets of the transportability problem. (a) A selection diagram in which transportability is trivial. (b) A selection diagram in which the causal relation R is more involved and shown to be transportable using Theorem 1. (c) A selection diagram in which the procedure given in [PB, 2011] is unable to recognize a transportable relation R .

[PB, 2011] derived this formula using the following theorem, which translates the property of transportability to the existence of a syntactic reduction using a sequence of *do-calculus* operations.

Theorem 1 (Do-calculus characterization [PB, 2011]). *Let D be the selection diagram characterizing Π and Π^* , and \mathbf{S} a set of selection variables in D . The relation $R = P^*(y|do(\mathbf{x}))$ is transportable from Π to Π^* if and only if the expression $P(y|do(\mathbf{x}), \mathbf{s})$ is reducible, using the rules of *do-calculus*, to an expression in which \mathbf{S} appears only as a conditioning variable in *do-free* terms.*

Theorem 1 is declarative but not computationally effective, for it does not specify the sequence of rules leading to the needed reduction, nor does it tell us if such a sequence exists.

To overcome this deficiency, [PB, 2011] proposed a recursive procedure (their Theorem 3) which can handle many examples, among them Fig. 1(b), but is not complete. We will show⁴, for example, that their procedure fails to recognize R as “transportable” in the diagram of Fig. 1(c) whereas the procedures developed in this paper will recognize it as such and will support it with the transport formula:

$$R = \sum_z P(z|do(x)) \sum_w P(w|do(x, z)) \sum_v P^*(v|w)P^*(y|v, w)$$

We summarize our contributions as follows:

- We derive a general graphical condition for deciding transportability of causal effects. We show that transportability is feasible if and only if a certain graph structure does not appear as an edge subgraph of the inputted selection diagram.
- We provide necessary or sufficient graphical conditions for special cases of transportability, for instance, controlled direct effects.
- Finally, we construct a complete algorithm for deciding transportability of joint causal effects and returning the correct transport formula whenever those effects are transportable.

⁴See Corollary 4 in [Bareinboim and Pearl, 2012a].

Preliminary Results

The basic semantical framework in our analysis rests on *probabilistic causal models* as defined in [Pearl, 2000, pp. 205], also called structural causal models or data-generating models. In the structural causal framework [Pearl, 2000, Ch. 7], actions are modifications of functional relationships, and each action $do(\mathbf{x})$ on a causal model M produces a new model $M_{\mathbf{x}} = \langle \mathbf{U}, \mathbf{V}, \mathbf{F}_{\mathbf{x}}, P(\mathbf{U}) \rangle$, where $F_{\mathbf{x}}$ is obtained after replacing $f_X \in \mathbf{F}$ for every $X \in \mathbf{X}$ with a new function that outputs a constant value x given by $do(\mathbf{x})$.

Key to the analysis of transportability is the notion of “identifiability,” defined below, which expresses the requirement that causal effects be computable from a combination of data P and assumptions embodied in a causal graph G .

Definition 1 (Causal Effects Identifiability [Pearl, 2000, pp. 77]). *The causal effect of an action $do(\mathbf{x})$ on a set of variables \mathbf{Y} such that $\mathbf{Y} \cap \mathbf{X} = \emptyset$ is said to be identifiable from P in G if $P_{\mathbf{x}}(\mathbf{y})$ is uniquely computable from $P(V)$ in any model that induces G .*

Causal models and their induced graphs are normally associated with one particular domain (also called setting, study, population, environment). In the transportability case, we extend this representation to capture properties of several domains simultaneously. This is made possible if we assume that there are no structural changes between the domains, that is, all structural equations share the same set of arguments, though the functional forms of the equations may vary arbitrarily.^{5, 6}

Definition 2 (Selection Diagram). *Let $\langle M, M^* \rangle$ be a pair of structural causal models [Pearl, 2000, pp. 205] relative to domains $\langle \Pi, \Pi^* \rangle$, sharing a causal diagram G . $\langle M, M^* \rangle$ is said to induce a selection diagram D if D is constructed as follows:*

1. Every edge in G is also an edge in D ;
2. D contains an extra edge $S_i \rightarrow V_i$ whenever there might exist a discrepancy $f_i \neq f_i^*$ or $P(U_i) \neq P^*(U_i)$ between M and M^* .

In words, the S -variables locate the *mechanisms* where structural discrepancies between the two domains are suspected to take place.⁷ Alternatively, one can see a selection diagram as a carrier of invariance claims between the mechanisms of both domains – the absence of a selection node pointing to a variable represents the assumption that the mechanism responsible for assigning value to that variable is the same in the two domains.

⁵This definition was left implicit in [PB, 2011].

⁶The assumption that there are no structural changes between domains can be relaxed starting with $D = G^*$ and adding S -nodes following the same procedure as in Def. 2, while enforcing acyclicity.

⁷Transportability assumes that enough structural knowledge about both domains is known in order to substantiate the production of their respective causal diagrams. In the absence of such knowledge, *causal discovery* algorithms can be used to infer the diagrams from data [Pearl and Verma, 1991; Pearl, 2000; Spirtes, Glymour, and Scheines, 2001].

Armed with a selection diagram and the concept of identifiability, transportability of causal effects (or transportability, for short) can be defined as follows:

Definition 3 (Causal Effects Transportability). *Let D be a selection diagram relative to domains $\langle \Pi, \Pi^* \rangle$. Let $\langle P, I \rangle$ be the pair of observational and interventional distributions of Π , and P^* be the observational distribution of Π^* . The causal effect $R = P_{\mathbf{x}}^*(\mathbf{y})$ is said to be transportable from Π to Π^* in D if $P_{\mathbf{x}}^*(\mathbf{y})$ is uniquely computable from P, P^*, I in any model that induces D .*

The problem of transportability generalizes the problem of identifiability, to witness note that all identifiable causal relations in (G^*, P^*) are also transportable, because they can be computed directly from Π^* and require no experimental information from Π . This observation engender the following definition of *trivial transportability*.

Definition 4. (Trivial Transportability) *A causal relation R is said to be trivially transportable from Π to Π^* , if $R(\Pi^*)$ is identifiable from (G^*, P^*) .*

The following observation establishes another connection between identifiability and transportability. For a given causal diagram G , one can produce a selection diagram D such that identifiability in G is equivalent to transportability in D . First set $D = G$, and then add selection nodes pointing to all variables in D , which represents that the target domain does not share any commonality with its pair – this is equivalent to the problem of identifiability because the only way to achieve transportability is to identify R from scratch in the target domain.

Another special case of transportability occurs when a causal relation has identical form in both domains – no recalibration is needed. This is captured by the following definition.

Definition 5. (Direct Transportability) *A causal relation R is said to be directly transportable from Π to Π^* , if $R(\Pi^*) = R(\Pi)$.*

A graphical test for direct transportability of $R = P^*(y|do(x), z)$ follows from do-calculus and reads: $(S \perp\!\!\!\perp Y|X, Z)_{G_{\overline{X}}}$; in words, X blocks all paths from S to Y once we remove all arrows pointing to X and condition on Z . As a concrete example, the z -specific effects in Fig. 1(a) is the same in both domains, hence, it is directly transportable.

These two cases will act as a basis to decompose the problem of transportability into smaller and more manageable subproblems (to be shown later on).

The following lemma provides an auxiliary tool to prove non-transportability and is based on refuting the uniqueness property required by Definition 3.

Lemma 1. *Let \mathbf{X}, \mathbf{Y} be two sets of disjoint variables, in population Π and Π^* , and let D be the selection diagram. $P_{\mathbf{x}}^*(\mathbf{y})$ is not transportable from Π to Π^* if there exist two causal models M^1 and M^2 compatible with D such that $P_1(\mathbf{V}) = P_2(\mathbf{V}), P_1^*(\mathbf{V}) = P_2^*(\mathbf{V}), P_1(\mathbf{V} \setminus \mathbf{W}|do(\mathbf{W})) = P_2(\mathbf{V} \setminus \mathbf{W}|do(\mathbf{W}))$, for any set \mathbf{W} , all families have positive distribution, and $P_1^*(\mathbf{y}|do(\mathbf{x})) \neq P_2^*(\mathbf{y}|do(\mathbf{x}))$.*

Proof. Let I be the set of interventional distributions $P(\mathbf{V} \setminus \mathbf{W} | do(\mathbf{W}))$, for any set \mathbf{W} . The latter inequality rules out the existence of a function from P, P^*, I to $P_x^*(y)$. \square

While Lemma 1 may sound cumbersome, it is nothing more than a formalization of the statement: “query Q cannot be computed from information set IS alone.” Naturally, when IS has three components, $\langle P^*, P, I \rangle$, the lemma becomes lengthy. But conceptually it is nothing more than the definition of non-identifiability which is standard in computer science, statistics and econometrics.

Even though the problems of identifiability and transportability are related, Lemma 1 indicates that proofs of non-transportability are more involved than those of non-identifiability. Indeed, to prove non-transportability requires the construction of two models agreeing on $\langle P, I, P^* \rangle$, while non-identifiability requires the two models to agree solely on the observational distribution P .

The simplest non-transportable structure is an extension of the famous ‘bow arc’ graph named here ‘s-bow arc’, see Fig. 2(a). The s-bow arc has two endogenous nodes: X , and its child Y , sharing a hidden exogenous parent U , and a S -node pointing to Y . This and similar structures that prevent transportability will be useful in our proof of completeness, which requires a demonstration that whenever the algorithm fails to transport a causal relation, the relation is indeed non-transportable.

Theorem 2. $P_x^*(y)$ is not transportable in the s-bow arc graph.

Proof. The proof will show a counter-example to the transportability of $P_x^*(Y)$ through two models M_1 and M_2 that agree in $\langle P, P^*, I \rangle$ and disagree in $P_x^*(y)$.

Assume that all variables are binary. Let the model M_1 be defined by the following system of structural equations: $X_1 = U, Y_1 = ((X \oplus U) \oplus S)$, $P_1(U) = 1/2$, and M_2 by the following one: $X_2 = U, Y_2 = S \vee (X \oplus U)$, $P_2(U) = 1/2$, where \oplus represents the *exclusive or* function.

Lemma 2. The two models agree in the distributions $\langle P, P^*, I \rangle$.

Proof. We show that the following equations must hold for M_1 and M_2 :

$$\begin{cases} P_1(X|S) = P_2(X|S), & S = \{0, 1\} \\ P_1(Y|X, S) = P_2(Y|X, S), & S = \{0, 1\} \\ P_1(Y|do(X), S = 0) = P_2(Y|do(X), S = 0) \end{cases}$$

for all values of X, Y . The equality between $P_i(X|S)$ is obvious since $(S \perp\!\!\!\perp X)$ and X has the same structural form in both models. Second, let us construct the truth table for Y :

X	S	U	Y_1	Y_2
0	0	0	0	0
0	0	1	1	1
0	1	0	1	1
0	1	1	0	1
1	0	0	1	1
1	0	1	0	0
1	1	0	0	1
1	1	1	1	1




Figure 2: (a) Smallest selection diagram in which $P(y|do(x))$ is not transportable (s-bow graph). (b) A selection diagram in which even though there is no S -node pointing to Y , the effect of X on Y is still not transportable due to the presence of a sC -tree (see Corollary 2).

To show that the equality between $P_i(Y = 1|X, S = 0), X = \{0, 1\}$ holds, we rewrite it as follows:

$$P_i(Y = 1|X, S = 0) = \frac{P_i(Y = 1|X, S = 0, U = 1)P_i(X|U = 1)P_i(U = 1)}{P_i(X)} + \frac{P_i(Y = 1|X, S = 0, U = 0)P_i(X|U = 0)P_i(U = 0)}{P_i(X)} \quad (1)$$

In eq. (1), the expressions for $X = \{0, 1\}$ are functions of the tuples $\{(X = 1, S = 0, U = 1), (X = 0, S = 0, U = 0)\}$, which evaluate to the same value in both models. Similarly, the expressions $P_i(Y = 1|X, S = 1)$ for $X = \{0, 1\}$ are functions of the tuples $\{(X = 1, S = 1, U = 1), (X = 0, S = 1, U = 0)\}$, which also evaluate to the same value in both models.

We further assert the equality between the interventional distributions in Π , which can be written using the do-calculus as

$$P_i(Y = 1|do(X), S = 0) = \sum_U P_i(Y|do(X), S = 0, U)P_i(U|do(X), S = 0) = P_i(Y = 1|X, S = 0, U = 1)P_i(U = 1) + P_i(Y = 1|X, S = 0, U = 0)P_i(U = 0), \quad X = \{0, 1\} \quad (2)$$

Evaluating this expression points to the tuples $\{(X = 1, S = 0, U = 1), (X = 1, S = 0, U = 0)$ and $(X = 0, S = 0, U = 1), (X = 0, S = 0, U = 0)\}$, which map to the same value in both models. \square

Lemma 3. There exist values of X, Y such that $P_1(Y|do(X), S = 1) \neq P_2(Y|do(X), S = 1)$.

Proof. Fix $X = 1, Y = 1$, and let us rewrite the desired quantity $R_i = P_i(Y = 1|do(X = 1), S = 1)$ in Π^* as

$$R_i = P_i(Y = 1|X = 1, S = 1, U = 1)P_i(U = 1) + P_i(Y = 1|X = 1, S = 1, U = 0)P_i(U = 0) \quad (3)$$

Since R_i is a function of the tuples $(X = 1, S = 1, U = 1), (X = 1, S = 1, U = 0)$, it evaluates in M_1 to $\{1, 1\}$ and in M_2 to $\{1, 0\}$.

Hence, together with the uniformity of $P(U)$, it follows that $R_1 = 1$ and $R_2 = 1/2$, which finishes the proof. \square

By Lemma 1, Lemmas 2 and 3 prove Theorem 2. \square

Characterizing Transportable Relations

The concept of confounded components (or C -components) was introduced in [Tian and Pearl, 2002] to represent clusters of variables connected through bidirected edges, and was instrumental in establishing a number of conditions for ordinary identification (Def. 1). If G is not a C -component itself, it can be uniquely partitioned into a set $\mathcal{C}(G)$ of C -components. We now recast C -components in the context of transportability.⁸

Definition 6 (sC-component). *Let G be a selection diagram such that a subset of its bidirected arcs forms a spanning tree over all vertices in G . Then G is a sC-component (selection confounded component).*

A special subset of C -components that embraces the ancestral set of Y was noted by [Shpitser and Pearl, 2006] to play an important role in deciding identifiability – this observation can also be applied to transportability, as formulated in the next definition.

Definition 7 (sC-tree). *Let G be a selection diagram such that $\mathcal{C}(G) = \{G\}$, all observable nodes have at most one child, there is a node Y , which is a descendent of all nodes, and there is a selection node pointing to Y . Then G is called a Y -rooted sC-tree (selection confounded tree).*

The presence of this structure (and generalizations) will prove to be an obstacle to transportability of causal effects. For instance, the s-bow arc in Fig. 2(a) is a Y -rooted sC-tree where we know $P_x^*(y)$ is non-transportable.

In certain classes of problems, the absence of such structures will prove sufficient for transportability. One such class is explored below, and consists of models in which the set X coincides with the parents of Y .

Theorem 3. *Let G be a selection diagram. Then for any node Y , the causal effects $P_{Pa(Y)}^*(y)$ is transportable if there is no subgraph of G which forms a Y -rooted sC-tree.*

See the full version [Bareinboim and Pearl, 2012a] for the complete proofs.

Theorem 3 provides a tractable transportability condition for the Controlled Direct Effect (CDE) – a key concept in modern mediation analysis, which permits the decomposition of effects into their direct and indirect components [Pearl, 2001; 2012]. CDE is defined as the effect of X on Y when all other parents of Y are held constant, and it is identifiable if and only if $P_{Pa(Y)}^*(y)$ is identifiable [Pearl, 2009, pp. 128].

The selection diagram in Fig. 1(a) does not contain any Y -rooted sC-trees as subgraphs, and therefore the direct effects (causal effects of Y 's parents on Y) is indeed transportable. In fact, the transportability of CDE can be determined by a more visible criterion:

Corollary 1. *Let G be a selection diagram. Then for any node Y , the direct effect $P_{Pa(Y)}^*(y)$ is transportable if there is no S node pointing to Y .*

⁸ C -components can itself be seen as an extension of the more elementary notion of *inducing path*, which was introduced much earlier in [Verma and Pearl, 1990].




Figure 3: Example of a selection diagram in which $P(Y|do(X))$ is not transportable, there is no sC-tree but there is a s^* -tree.

Generalizing to arbitrary effects, the following result provides a necessary condition for transportability whenever the whole graph is a sC-tree.

Theorem 4. *Let G be a Y -rooted sC-tree. Then the effects of any set of nodes in G on Y are not transportable.*

The next corollary demonstrates that sC-trees are obstacles to the transportability of $P_x(y)$ even when they do not involve Y , i.e., transportability is not a local problem – if there exists a node W that is an ancestor of Y but not necessarily “near” it, transportability is still prohibited (see Fig. 2(b)). This fact anticipates that transporting causal effects of singleton Y is not necessarily easier than the general problem of transportability.

Corollary 2. *Let G be a selection diagram, and \mathbf{X} and \mathbf{Y} a set of variables. If there exists a node W that is an ancestor of some node $Y \in \mathbf{Y}$ such that there exists a W -rooted sC-tree which contains any variables in \mathbf{X} , then $P_x(\mathbf{y})$ is not transportable.*

We now generalize the definition of sC-trees (and Theorem 4) in two ways: first, Y is augmented and can be a set of variables; second, S -nodes can point to any variable within the sC-component, not necessarily to root nodes. For instance, consider the graph G in Fig. 3. Note that there is no Y -rooted sC-tree nor W -rooted sC-tree in G (where W is an ancestor of Y), and so the previous results cannot be applied even though the effect of X on Y is not transportable in G – still, there exists a Y -rooted s^* -tree in G , which will prevent the transportability of the causal effect.

Definition 8 (s^* -tree). *Let G be a selection diagram, where \mathbf{Y} is the maximal root set. Then G is a \mathbf{Y} -rooted s^* -tree if G is a sC-component, all observable nodes have at most one child, and there is a selection node pointing to some vertex of G (not necessarily in \mathbf{Y}).*

We next conveniently introduce a structure that witnesses non-transportability characterized by a pair of s^* -trees, which extends Shpitser’s hedge used for ordinary identifiability. Transportability will be shown impossible whenever such structure exists as an edge subgraph of the given selection diagram.

Definition 9 (s-hedge). *Let \mathbf{X}, \mathbf{Y} be set of variables in G . Let F, F' be \mathbf{R} -rooted s^* -trees such that $F \cap \mathbf{X} \neq \emptyset$, $F' \cap \mathbf{X} = \emptyset$, $F' \subseteq F$, $\mathbf{R} \subset An(\mathbf{Y})_{G_{\overline{\mathbf{X}}}}$. Then F and F' form a s-hedge for $P_x^*(\mathbf{Y})$ in G .*

For instance, in Fig. 3, the s^* -trees $F' = \{C, Y\}$, and $F = F' \cup \{X, A, B\}$ form a s-hedge to $P_x^*(y)$.

We state below the formal connection between s-edges and non-transportability.

Theorem 5. Assume there exist F, F' that form a s -hedge for $P_x^*(y)$ in Π and Π^* . Then $P_x^*(y)$ is not transportable from Π to Π^* .

To prove that the s -hedges characterize non-transportability in selection diagrams, we construct in the next section an algorithm which transport any causal effects that do not contain a s -hedge.

A Complete Algorithm For Transportability of Joint Effects

The algorithm proposed to solve transportability is called **sID** (see Fig. 4) and extends previous analysis of identifiability given in [Pearl, 1995; Kuroki and Miyakawa, 1999; Tian and Pearl, 2002; Shpitser and Pearl, 2006; Huang and Valtorta, 2006]. We choose to modify the version provided by Shpitser since the hedge structure is explicitly employed, which will show to be instrumental to prove completeness. We build on two observations developed along the paper:

- (i) Transportability: Causal relations can be partitioned into trivially and directly transportable.
- (ii) Non-transportability: The existence of a s -hedge as an edge subgraph of the inputted selection diagram can be used to prove non-transportability.

The algorithm **sID** first applies the typical c -component decomposition on top of the inputted selection diagram D , partitioning the original problem into smaller blocks (call these blocks sc -factors) until either the entire expression is transportable, or it runs into the problematic s -hedge structure.

More specifically, for each sc -factor Q , **sID** tries to directly transport Q . If it fails, **sID** tries to trivially transport Q , which is equivalent to solving an ordinary identification problem. **sID** alternates between these two types of transportability, and whenever it exhausts the possibility of applying these operations, it exits with failure with a counterexample for transportability – that is, the graph local to the faulty call witnesses the non-transportability of the causal query since it contains a s -hedge as edge subgraph.

Before showing the more formal properties of **sID**, we demonstrate how **sID** works through the transportability of $Q = P^*(y|do(x))$ in the graph in Fig. 1(c).

Since $D = An(Y)$ and $\mathcal{C}(D \setminus \{X\}) = (C_0, C_1, C_2)$, where $C_0 = D(\{Z\})$, $C_1 = D(\{W\})$, and $C_2 = D(\{V, Y\})$, we invoke line 4 and try to transport respectively $Q_0 = P_{x,w,v,y}^*(z)$, $Q_1 = P_{x,z,v,y}^*(w)$, and $Q_2 = P_{x,z,w}^*(v, y)$. Thus the original problem reduces to transporting $\sum_{z,w,v} P_{x,w,v,y}^*(z) P_{x,z,v,y}^*(w) P_{x,z,w}^*(v, y)$.

Evaluating the first expression, we trigger line 2, noting that nodes that are not ancestors of Z can be ignored. This implies that $P_{x,w,v,y}^*(z) = P_x^*(z)$ with induced subgraph $G_0 = \{X \rightarrow Z, X \leftarrow U_{xz} \rightarrow Z\}$, where U_{xz} stands for the hidden variable between X and Z . **sID** goes to line 5, in which in the local call $\mathcal{C}(D \setminus \{X\}) = \{G_0\}$. Note that in the ordinary identifiability problem the procedure would fail at this point, but **sID** proceeds to line 6 testing whether $(S \perp\!\!\!\perp Z|X)_{D_{\overline{X}}}$. The test comes true, which makes **sID** directly transport Q_0 with data from the experimental domain Π , i.e., $P_x^*(z) = P_x(z)$.

function **sID**(y, x, P^*, I, D)

INPUT: x, y value assignments, P^* observational distribution in Π^* , I set of interventional distributions in Π , D a selection diagram, S set of selection nodes.

OUTPUT: Expression for $P_x^*(y)$ in terms of P^*, I or $FAIL(F, F')$.

```

1  if  $x = \emptyset$ , return  $\sum_{V \setminus Y} P^*(V)$ 
2  if  $V \setminus An(Y)_D \neq \emptyset$ ,
    return sID( $y, x \cap An(Y)_D, \sum_{V \setminus An(Y)_D} P^*, An(Y)_D$ )
3  Set  $W = (V \setminus X) \setminus An(Y)_{D_{\overline{X}}}$ .
    if  $W \neq \emptyset$ , return sID( $y, x \cup w, P^*, D$ )
4  if  $\mathcal{C}(D \setminus X) = \{C_0, C_1, \dots, C_k\}$ ,
    return  $\sum_{V \setminus \{Y, X\}} \prod_i \mathbf{sID}(c_i, V \setminus c_i, P^*, D)$ 
5  if  $\mathcal{C}(D \setminus X) = \{C_0\}$ 
    if  $(S \perp\!\!\!\perp Y | X)_{D_{\overline{X}}}$ , return  $P(y|do(x))$ 
    if  $\mathcal{C}(D) = \{D\}$ , FAIL( $D, C_0$ )
6  if  $C_0 \in \mathcal{C}(D)$ , return  $\sum_{s \setminus Y} \prod_{i|V_i \in S} P^*(v_i|V_D^{(i-1)})$ 
7  if  $(\exists C') C_0 \subset C' \in \mathcal{C}(D)$ , return sID( $y, x \cap C'$ ,
     $\prod_{i|V_i \in C'} P^*(V_i|V_D^{(i-1)} \cap C', v_D^{(i-1)} \setminus C')$ ,  $C'$ ).

```

Figure 4: Modified version of identification algorithm capable of recognizing transportable relations.

Evaluating the second expression, we again trigger line 2, which implies that $P_{x,z,v,y}^*(w) = P_{x,z}^*(w)$ with induced subgraph $G_1 = \{X \rightarrow Z, Z \rightarrow W, X \leftarrow U_{xz} \rightarrow Z\}$. **sID** goes to line 5, in which in the local call $\mathcal{C}(D \setminus \{X\}) = \{G_1\}$. Thus it proceeds to line 6 testing whether $(S \perp\!\!\!\perp W|X, Z)_{D_{\overline{X}, \overline{Z}}}$. The test comes true again, which makes **sID** directly transport Q_1 with data from the experimental domain Π , i.e., $P_{x,z}^*(w) = P_{x,z}(w)$.

Evaluating the third expression, **sID** goes to line 5 in which $\mathcal{C}(D \setminus \{X, Z, W\}) = \{G_2\}$, where $G_2 = \{V \rightarrow Y, S \rightarrow V, V \leftarrow U_{vy} \rightarrow Y\}$. It proceeds to line 6 testing whether $(S \perp\!\!\!\perp W|X, Z)_{D_{\overline{X}, \overline{Z}}}$, which is false in this case. It tests the other conditions until it reaches line 9, in which $C' = G_0 \cup G_2 \cup \{X \leftarrow U_{xy} \rightarrow Y\}$. Thus it tries to transport $Q_2 = P_{x,z}^*(v, y)$ over the induced graph C' , which stands for ordinary identification, and trivially yields (after simplification) $\sum_v P^*(v|w) P^*(y|v, w)$. The return of these calls composed indeed coincide with the expression provided in the first section.

We state next soundness and completeness of **sID**, see the proofs in [Bareinboim and Pearl, 2012].

Theorem 6 (soundness). Whenever **sID** returns an expression for $P_x^*(y)$, it is correct.

Theorem 7. Assume **sID** fails to transport $P_x^*(y)$ (executes line 7). Then there exists $X' \subseteq X, Y' \subseteq Y$, such that the graph pair D, C_0 returned by the fail condition of **sID** contain as edge subgraphs s^* -trees F, F' that form a s -hedge for $P_{x'}^*(y')$.

Corollary 3 (completeness). **sID** is complete.

Conclusions

We formally study the problem of “transportability,” or “external validity,” which we view as a license to transfer causal information learned in experimental studies to a different environment, in which only observational studies can be conducted.⁹

More specifically, we provide a complete (necessary and sufficient) graphical condition for deciding when the causal effect of one set of variables on another can be transported from experimental to non-experimental environment. We further provide a complete algorithm for computing the correct transport formula whenever this graphical condition holds.

Appendix 1

The *do*-calculus [Pearl, 1995] consists of three rules that permit us to transform expressions involving *do*-operators into other expressions of this type, whenever certain conditions hold in the causal diagram G .

We consider a DAG G in which each child-parent family represents a deterministic function $x_i = f_i(pa_i, \epsilon_i)$, $i = 1, \dots, n$, where pa_i are the parents of variables X_i in G ; and ϵ_i , $i = 1, \dots, n$ are arbitrarily distributed random disturbances, representing background factors that the investigator chooses not to include in the analysis.

Let X , Y , and Z be arbitrary disjoint sets of nodes in a causal DAG G . An expression of the type $E = P(y|do(x), z)$ is said to be compatible with G if the interventional distribution described by E can be generated by parameterizing the graph with a set of functions f_i and a set of distributions of ϵ_i , $i = 1, \dots, n$.

We denote by $G_{\overline{X}}$ the graph obtained by deleting from G all arrows pointing to nodes in X . Likewise, we denote by $G_{\underline{X}}$ the graph obtained by deleting from G all arrows emerging from nodes in X . To represent the deletion of both incoming and outgoing arrows, we use the notation $G_{\overline{X}\underline{Z}}$.

The following three rules are valid for every interventional distribution compatible with G :

Rule 1: $P_x(y|z, w) = P_x(y|w)$ if $(Y \perp\!\!\!\perp Z|X, W)_{G_{\overline{X}}}$.

Rule 2: $P_{x,z}(y|w) = P_x(y|z, w)$ if $(Y \perp\!\!\!\perp Z|X, W)_{G_{\overline{X}\underline{Z}}}$.

Rule 3: $P_{x,z}(y|w) = P_x(y|w)$ if $(Y \perp\!\!\!\perp Z|X, W)_{G_{\overline{X}\underline{Z}^*}}$, where $Z^* = Z \setminus Anc(W)_{G_{\overline{X}}}$.

References

Adelman, L. 1991. Experiments, quasi-experiments, and case studies: A review of empirical methods for evaluating decision support systems. *Systems, Man and Cybernetics, IEEE Transactions on* 21(2):293–301.

Bareinboim, E., and Pearl, J. 2012a. Transportability of causal effects: Completeness results. Technical Report TR-390, Cognitive Systems Laboratory, Department of Computer Science, UCLA.

Bareinboim, E., and Pearl, J. 2012b. Controlling selection bias in causal inference. In Girolami, M., and Lawrence, N., eds., *Proceedings of The Fifteenth International Conference on Artificial Intelligence and Statistics (AISTATS 2012), JMLR (22)*, 100–108.

⁹In practice, the estimates produced may still suffer from misspecification bias, finite-sample bias, and sample-selection bias (see [Bareinboim and Pearl, 2012b]).

Campbell, D., and Stanley, J. 1963. *Experimental and Quasi-Experimental Designs for Research*. Chicago: Wadsworth Publishing.

Glass, G. V. 1976. Primary, secondary, and meta-analysis of research. *Educational Researcher* 5(10):pp. 3–8.

Hedges, L. V., and Olkin, I. 1985. *Statistical Methods for Meta-Analysis*. Academic Press.

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*. Menlo Park, CA: AAAI Press. 1149–1156.

Koller, D., and Friedman, N. 2009. *Probabilistic Graphical Models: Principles and Techniques*. MIT Press.

Kuroki, M., and Miyakawa, M. 1999. Identifiability criteria for causal effects of joint interventions. *Journal of the Royal Statistical Society* 29:105–117.

Manski, C. 2007. *Identification for Prediction and Decision*. Cambridge, Massachusetts: Harvard University Press.

Owen, A. B. 2009. Karl pearsons meta-analysis revisited. *Annals of Statistics* 37(6B):3867–3892.

Pearl, J., and Bareinboim, E. 2011. Transportability of causal and statistical relations: A formal approach. In *Proceedings of the Twenty-Fifth National Conference on Artificial Intelligence*. Menlo Park, CA: AAAI Press. 247–254.

Pearl, J., and Verma, T. 1991. A theory of inferred causation. In Allen, J.; Fikes, R.; and Sandewall, E., eds., *Principles of Knowledge Representation and Reasoning: Proceedings of the Second International Conference*. San Mateo, CA: Morgan Kaufmann. 441–452.

Pearl, J. 1995. Causal diagrams for empirical research. *Biometrika* 82(4):669–710.

Pearl, J. 2000. *Causality: Models, Reasoning, and Inference*. New York: Cambridge University Press. Second ed., 2009.

Pearl, J. 2001. Direct and indirect effects. In *Proceedings of the Seventeenth Conference on Uncertainty in Artificial Intelligence*. San Francisco, CA: Morgan Kaufmann. 411–420.

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

Pearl, J. 2012. The mediation formula: A guide to the assessment of causal pathways in nonlinear models. In Berzuini, C.; Dawid, P.; and Bernardinell, L., eds., *Causality: Statistical Perspectives and Applications*. New York: Wiley. Chapter 12.

Shadish, W.; Cook, T.; and Campbell, D. 2002. *Experimental and Quasi-Experimental Designs for Generalized Causal Inference*. Boston: Houghton-Mifflin, second edition.

Shpitser, I., and Pearl, J. 2006. Identification of joint interventional distributions in recursive semi-Markovian causal models. In *Proceedings of the Twenty-First National Conference on Artificial Intelligence*. Menlo Park, CA: AAAI Press. 1219–1226.

Spirtes, P.; Glymour, C.; and Scheines, R. 2001. *Causation, Prediction, and Search*. Cambridge, MA: MIT Press, 2nd edition.

Tian, J., and Pearl, J. 2002. A general identification condition for causal effects. In *Proceedings of the Eighteenth National Conference on Artificial Intelligence*. Menlo Park, CA: AAAI Press/The MIT Press. 567–573.

Verma, T., and Pearl, J. 1990. Equivalence and synthesis of causal models. In *Proceedings of the Sixth Conference on Uncertainty in Artificial Intelligence*, 220–227. Also in P. Bonissone, M. Henrion, L.N. Kanal and J.F. Lemmer (Eds.), *Uncertainty in Artificial Intelligence 6*, Elsevier Science Publishers, B.V., 255–268, 1991.