

Generalized Adjustment Under Confounding and Selection Biases

Juan D. Correa

Computer Science Department
Purdue University
correagr@purdue.edu

Jin Tian

Department of Computer Science
Iowa State University
jtian@iastate.edu

Elias Bareinboim

Computer Science Department
Purdue University
eb@purdue.edu

Abstract

Selection and *confounding* biases are the two most common impediments to the applicability of causal inference methods in large-scale settings. We generalize the notion of backdoor adjustment to account for both biases and leverage external data that may be available without selection bias (e.g., data from census). We introduce the notion of adjustment pair and present complete graphical conditions for identifying causal effects by adjustment. We further design an algorithm for listing all admissible adjustment pairs in polynomial delay, which is useful for researchers interested in evaluating certain properties of some admissible pairs but not all (common properties include cost, variance, and feasibility to measure). Finally, we describe a statistical estimation procedure that can be performed once a set is known to be admissible, which entails different challenges in terms of finite samples.

Introduction

A fundamental challenge pervasive throughout science is the study of cause and effect relationships from a combination of non-experimental observations and substantive knowledge about the phenomenon under investigation. Causal relations are deemed more interpretable and robust than their statistical counterparts. They are more amenable to extrapolation to new, unforeseen situations. Understanding the world and constructing explanations are almost invariably accomplished through the presentation of causal knowledge with a coherent articulation of a causal story (Pearl 2000; Spirtes, Glymour, and Scheines 2001; Bareinboim and Pearl 2016; Pearl, Glymour, and Jewell 2016).

Two of the most common obstacles to discovering causal relations appear in the form of two biases – *confounding* and *selection*. The first one may arise from the lack of control over the decision-making process and the selection of actions, possibly due to costs, ethical, or technical considerations. This implies that the data is collected under an observational regime, where the population follows its natural tendency. Our goal, however, is to predict how the population will react when it undergoes a change (intervention), following a new, compulsory decision protocol. For instance, one is usually not interested in estimating the correlation between smoking and cancer (natural), but to establish whether

the incidence of cancer would decrease had smoking been banned in the corresponding population.

The problem of *identifiability* gives formal dressing to this issue (Pearl 2000, Ch. 3). Specifically, it is concerned with determining the effect of a treatment (X) on an outcome (Y), denoted $P(y|do(x))$, based on the observational, non-experimental distribution $P(\mathbf{v})$ (where \mathbf{V} represents observable variables) and causal assumptions commonly represented as a directed acyclic graph. The difference between $P(y|do(x))$ and its probabilistic counterpart, $P(y|x)$, is known as *confounding bias* (Bareinboim and Pearl 2016). For the graph in Fig. 1(a), the probability distribution $P(y|x)$ includes variations due to the “backdoor” path $X \leftarrow Z \rightarrow Y$, while the distribution $P(y|do(x))$ describes a regime (Fig. 1(b)) where the incoming arrows towards X are cut and only causal influence remains. Confounding bias, in this case, appears in the form of extraneous variations of Y that are not legitimately explained by X , but are generated by a third variable, Z in this case.

The problem of confounding has been extensively studied in the literature. A systematic mathematical treatment was given in (Pearl 1995), which included the *do-calculus*. The *do-calculus* was shown complete for non-parametric identifiability (Tian and Pearl 2002; Huang and Valtorta 2006; Shpitser and Pearl 2006; Bareinboim and Pearl 2012a).

Despite the generality of such results, in practice, the most common and pervasive method for controlling confounding bias is known as the backdoor-adjustment (Pearl 1995). The backdoor-adjustment formula dictates that the effect of X on Y can be computed by controlling for a set of covariates Z , i.e., averaging the conditional distribution of outcome Y given treatment X and Z , weighted by the marginal distribution of Z . Pearl provided a formal and graphical justification for under what conditions a set Z could make the adjustment formula valid (for a survey, see (Pearl 1995)).

The second bias, *selection*, may appear because of preferential exclusion of units from the sample. For instance, in a typical study of the effect of grades on college admission, subjects with higher achievement tend to report their scores more frequently than those who scored lower. In this case, the data-gathering process will reflect a distortion in the sample’s proportions and, since the data is no longer a faithful representation of the underlying population, biased estimates will be produced regardless of the number of sam-

ples collected (even if the treatment is controlled).

The problem of selection bias can also be modeled graphically. In Fig. 1(c), for example, S represents a binary indicator of entry into the data pool, such that $S=1$ if a unit is included in the sample and $S=0$ otherwise (Bareinboim and Pearl 2012b). In this case, the selection is affected by the treatment as represented by the arrow $X \rightarrow S$ (e.g., people with higher grades have a higher chance of reporting their scores). Clearly, when the sampling process is completely random, S is independent of all variables in the analysis. When samples are collected preferentially, the causal effects not only need to be identified, but also *recovered* (Bareinboim and Pearl 2012b) from the distribution $P(\mathbf{v}|S=1)$ instead of $P(\mathbf{v})$.

Selection bias has been studied in a wide range of subjects and contexts, including different tasks in AI (Cooper 1995; Elkan 2001; Zadrozny 2004; Cortes et al. 2008), statistics (Whittemore 1978; Little and Rubin 1987; Robinson and Jewell 1991; Kuroki and Cai 2006; Evans and Didelez 2015), throughout the empirical sciences (e.g., genetics (Pirinen, Donnelly, and Spencer 2012; Mefford and Witte 2012), economics (Heckman 1979; Angrist 1997), and epidemiology (Robins 2001; Glymour and Greenland 2008)).

The backdoor-adjustment was not used to control for selection bias until recently. Bareinboim, Tian, and Pearl (2014) provided a sufficient condition, formally showing that adjustment could be used to control for both confounding and selection biases. Later on, Correa and Bareinboim (2017) studied how externally available, unbiased data over the covariates could be leveraged to further the reach of this technique. For instance, the effect $P(y|do(x))$ for the model in Fig. 1(c) can be identified and recovered only if external data over Z , (i.e. $P(z)$) is available. In this case, the adjustment averages the biased conditional distribution of outcome Y given the treatment X ($P(y|x, z, S=1)$) weighted by the unbiased distribution of Z ($P(z)$).

There are still simple (but subtle) situations that remain unsolved by these previous results. To witness, consider the model in Fig. 1(d), where X represents whether a patient took or not a drug, Y indicates whether the patient recovered or not from the disease, Z_1 and Z_2 represent if the patient has a certain genetic condition and has severe headaches, respectively. The arrow from Z_2 to the selection mechanism (S) encodes the fact that patients with a headache are more likely to seek help in the hospital and, therefore, are more likely of being sampled.

The previous methods require one to collect unbiased (external) data on Z_1 , which may prove too costly. In fact, this would require performing a genetic test on a significant amount of patients, which is not routinely done in the hospital when the person reports a headache. Nevertheless, unbiased measurements of Z_2 may be obtained from test reports conducted over the whole population, given that headache is a pervasive side effect and vast amounts of demographic information are available about it. It will be shown later on that the adjustment technique can be extended to combine the partial unbiased data with biased data to produce an estimand of the causal effect in the overall population.

The goal of this paper is to explain the general principle




Figure 1: (a, b) represent the pre- and post-interventional models where X, Y, Z are, respectively, the treatment, outcome, and set of confounders. (c) Setting where both selection and confounding biases are present. (d) Settings where confounding or selection can be controlled, but not both, unless external data on $P(Z_1), P(Z_2)$ or both are available.

that licenses this extrapolation to take place. We will characterize the use of covariate adjustment for causal effect identification under selection bias for arbitrary causal graphs when a combination of biased and unbiased datasets are available. Specifically, the contributions of our paper are:

1. **(Graphical Characterization)** We introduce a generalized notion of adjustment formula to produce an estimand that combines biased and unbiased datasets. We then prove a necessary and sufficient graphical condition for the admissibility of a set of covariates for adjustment.
2. **(Algorithmic Characterization)** We construct a complete algorithm that efficiently finds *all* sets that are admissible for generalized adjustment. The algorithm runs with polynomial delay and is useful for identifying admissible sets with certain special properties (e.g., low measurement cost, higher statistical precision).
3. **(Statistical Procedure)** We demonstrate a general statistical procedure based on *inverse probability weighting (IPW)* to estimate the adjustment formula from data.

All proofs can be found in the sup. material (Correa, Tian, and Bareinboim 2017).

Definitions and Related Work

The systematic analysis of confounding and selection biases requires a formal language where the characterization of the underlying data-generating model can be encoded explicitly. We use the language of Structural Causal Models (SCM) (Pearl 2000, pp. 204-207). Formally, a SCM M is a 4-tuple $\langle U, V, F, P(u) \rangle$, where U is a set of exogenous (latent) variables and V is a set of endogenous (measured) variables. F represents a collection of functions $F = \{f_i\}$ such that each endogenous variable $V_i \in V$ is determined by a function $f_i \in F$, where f_i is a mapping from the respective domain of $U_i \cup PA_i$ to V_i , $U_i \subseteq U$, $PA_i \subseteq V \setminus V_i$, and the entire set F forms a mapping from U to V . The uncertainty is encoded through a probability distribution over the exogenous variables, $P(u)$. Within the structural semantics, performing an action $X = x$ is represented through the do-operator, $do(X = x)$, which encodes the operation of replacing the

original equation of X by the constant x and induces a sub-model M_x . For a detailed discussion on the properties of structural models, we refer readers to (Pearl 2000, Ch. 7).

In this paper, bold capital letters denote sets of variables, while bold lower-case letters stand for particular assignments to those variables. The family relationships in the graph are written as $An(X)$, $Pa(X)$, $De(X)$, which stand for the set of ancestors, parents, and descendants, respectively, of a given variable X . Variables are assumed to be ancestors and descendants of themselves. The letter G is used to refer to the causal graph, $G_{\overline{\mathbf{X}}}$ the graph resulting from the removal of all incoming edges to \mathbf{X} in G , and $G_{\underline{\mathbf{X}}}$ the graph resulting from removing all outgoing edges from \mathbf{X} .

Adjustment for Confounding Bias

We discuss in this section the notion of adjustment and how it is used to control for confounding bias. It also provides a basic survey of the most significant results in the literature.

Definition 1 (Adjustment (Pearl 2000)). *Given a causal diagram G containing a set of variables \mathbf{V} and pairwise disjoint sets \mathbf{X} , \mathbf{Y} , $\mathbf{Z} \subseteq \mathbf{V}$, the set \mathbf{Z} is called covariate adjustment for estimating the causal effect of \mathbf{X} on \mathbf{Y} (or usually just adjustment), if for every distribution $P(\mathbf{v})$ compatible with G it holds that*

$$P(\mathbf{y} \mid do(\mathbf{x})) = \sum_{\mathbf{z}} P(\mathbf{y} \mid \mathbf{x}, \mathbf{z})P(\mathbf{z}). \quad (1)$$

Finding an adjustment set relative to \mathbf{X} and \mathbf{Y} enables the identification of the corresponding causal effect. Several criteria have been developed to determine whether a set \mathbf{Z} is valid for adjustment. The most representative result for controlling for confounding bias by adjustment is known as the ‘‘Backdoor criterion’’ (Pearl 1993; 2000), defined below:

Definition 2 (Backdoor Criterion). *A set of variables \mathbf{Z} satisfies the Backdoor Criterion relative to a pair of variables (X, Y) in a directed acyclic graph G if:*

- (i) *No node in \mathbf{Z} is a descendant of X .*
- (ii) *\mathbf{Z} blocks every path between X and Y that contains an arrow into X .*

Intuitively, the backdoor criterion identifies the sets that block the non-causal paths (paths with arrows incoming towards X) while leaving the causal paths undisturbed.

It was further noted that certain descendants of X could be included in the adjustment set without sacrificing its validity (Pearl and Paz 2010). When selection bias is not present, Shpitser, VanderWeele, and Robins (2010) further showed that adjustment is complete if the non-proper causal paths are blocked (while the proper ones are left undisturbed), namely:

Definition 3 (Proper Causal Path). *Let \mathbf{X} and \mathbf{Y} be sets of nodes. A causal path from a node in \mathbf{X} to a node in \mathbf{Y} is called proper if it does not intersect \mathbf{X} except at the starting point.*

van der Zander, Liskiewicz, and Textor (2014) proposed an alternative complete formulation of adjustment called ‘‘Constructive Backdoor’’, which led to an efficient algorithmic treatment of the problem (without selection bias). This characterization follows a graph transformation that will prove useful in the context of selection bias:

Definition 4 (Proper Backdoor Graph). *Let G be a causal diagram, and \mathbf{X}, \mathbf{Y} be disjoint subsets of variables. The proper backdoor graph, denoted as $G_{\mathbf{X}\mathbf{Y}}^{pbd}$, is obtained from G by removing the first edge of every proper causal path from \mathbf{X} to \mathbf{Y} .*

Adjustment for Confounding and Selection Bias

Formally, the task of estimating a probabilistic quantity from a selection-biased distribution is known as *recovering* from selection bias (Bareinboim, Tian, and Pearl 2014). From now on, we assume that the set \mathbf{V} stands for all the observed variables measured under selection bias (not including the selection mechanism S). In this context, the input usually consists of a distribution collected under selection bias, $P(\mathbf{v} \mid S=1)$. The probability of selection $P(s)$ is assumed to be unknown. The goal of the analysis is to determine the unbiased causal distribution $P(\mathbf{y} \mid do(\mathbf{x}))$. In practical applications, unbiased observations of a subset of the variables may be available for use (e.g., the age and gender distributions). We’ll show this data can be leveraged to help the recoverability of causal effects by adjustment. We use \mathbf{T} to denote the set of externally and unbiasedly measured variables, and consider $P(\mathbf{t})$ as an input to the analysis.

Bareinboim, Tian, and Pearl (2014) studied the use of adjustment for simultaneously controlling for both confounding and selection biases. In particular, they introduced the *Selection-Backdoor* criterion (called s-backdoor), which is a sufficient condition for recovering causal effects from a biased distribution $P(\mathbf{v} \mid S=1)$ and externally unbiased data $P(\mathbf{t})$. Building on these results, Correa and Bareinboim (2017) devised a set of complete conditions for when none of the covariates are measured (Def. 5) externally ($\mathbf{Z} \cap \mathbf{T} = \emptyset$), and the case when all of them are (Def. 6) measured without selection bias ($\mathbf{Z} \subseteq \mathbf{T}$). Let $\mathbf{X}, \mathbf{Y}, \mathbf{Z}$ be disjoint sets of variables and G a causal diagram augmented with a variable S , then the criteria are shown next:

Definition 5 (Generalized Adjustment Criterion Type 1 (GACT1)). *\mathbf{Z} satisfies the criterion w.r.t. \mathbf{X}, \mathbf{Y} in G if:*

- (a) *No element of \mathbf{Z} is a descendant in $G_{\overline{\mathbf{X}}}$ of any $W \notin \mathbf{X}$ which lies on a proper causal path from \mathbf{X} to \mathbf{Y} .*
- (b) *All non-causal paths between \mathbf{X} and \mathbf{Y} in G are blocked by \mathbf{Z} and S .*
- (c) *\mathbf{Y} is d -separated from S given \mathbf{X} under the intervention $do(\mathbf{x})$, i.e., $(\mathbf{Y} \perp\!\!\!\perp S \mid \mathbf{X})_{G_{\overline{\mathbf{X}}}}$.*
- (d) *Every $X \in \mathbf{X}$ is either a non-ancestor of S or it is independent of \mathbf{Y} in $G_{\underline{\mathbf{X}}}$, i.e., $\forall X \in \mathbf{X} \cap An(S) (X \perp\!\!\!\perp \mathbf{Y})_{G_{\underline{\mathbf{X}}}}$.*

Definition 6 (Generalized Adjustment Criterion Type 2 (GACT2)). *\mathbf{Z} satisfies the criterion w.r.t. \mathbf{X}, \mathbf{Y} in G if:*

- (a) *No element of \mathbf{Z} is a descendant in $G_{\overline{\mathbf{X}}}$ of any $W \notin \mathbf{X}$ which lies on a proper causal path from \mathbf{X} to \mathbf{Y} .*
- (b) *All non-causal paths between \mathbf{X} and \mathbf{Y} in G are blocked by \mathbf{Z} .*
- (c) *\mathbf{Y} is d -separated from the selection mechanism S given \mathbf{Z} and \mathbf{X} , i.e., $(\mathbf{Y} \perp\!\!\!\perp S \mid \mathbf{X}, \mathbf{Z})$.*

It was shown that a set \mathbf{Z} satisfies the criterion in Def. 5 if and only if:

$$P(y | do(\mathbf{x})) = \sum_{\mathbf{z}} P(y | \mathbf{x}, \mathbf{z}, S=1)P(\mathbf{z} | S=1), \quad (2)$$

and \mathbf{Z} satisfies the criterion in Def. 6 if and only if:

$$P(y | do(\mathbf{x})) = \sum_{\mathbf{z}} P(y | \mathbf{x}, \mathbf{z}, S=1)P(\mathbf{z}). \quad (3)$$

Generalized Adjustment with Partial External Data

The criteria discussed above (Defs. 5 and 6) are complete to decide whether an adjustment set \mathbf{Z} is valid to identify the effect $P(y|do(\mathbf{x}))$ from the inputs $\{P(\mathbf{v} | S=1)\}$ and $\{P(\mathbf{v} | S=1), P(\mathbf{t}); \text{ with } \mathbf{T} \supseteq \mathbf{Z}\}$, respectively. We note that these tasks represent two extremes over the spectrum of how much unbiased data may be available for the researcher – the former assume that no external data is available, while the latter that all covariates are externally available.

A natural question to ask is whether it is possible to find valid adjustment sets within this spectrum, that is, to perform adjustment when only a subset \mathbf{Z}^T of the covariates \mathbf{Z} require external measurements (i.e., $\mathbf{Z}^T \subseteq \mathbf{Z} \cap \mathbf{T}$).

The possibility of using different subsets of covariates for adjustment has practical implications in the design of the study and the feasibility of estimating causal effects by adjustment. To witness, consider the causal diagram in Fig. 1(d), and note that there is no set \mathbf{Z} satisfying Def. 5, and that sets $\mathbf{Z}=\{Z_1\}$ and $\mathbf{Z}=\{Z_1, Z_2\}$ are valid for adjustment with respect to Def. 6, if the distribution $P(\mathbf{z})$ is available in addition to $P(x, y, z_1, z_2 | S=1)$. As discussed previously, in practical terms, unbiased measurements of Z_2 may be obtainable from test reports of the drug, but getting unbiased data for Z_1 may prove very challenging. Interestingly enough, the causal effect can still be identified by adjustment on the set $\mathbf{Z}=\{Z_1, Z_2\}$, if $P(z_2)$ is available even if $P(z_1)$ is not, which is shown in the expression below:

$$P(y|do(\mathbf{x})) = \sum_{\mathbf{z}} P(y|\mathbf{x}, \mathbf{z}, S=1)P(z_1|z_2, S=1)P(z_2) \quad (4)$$

Even when unbiased data over all the candidate covariates is available, it may be the case that no valid adjustment in the form given by Eqs. (2) and (3) exists, while it is still possible to adjust by a subset of the covariates. To witness, consider the model shown in Fig. 2 and note that while neither criteria is applicable, the effect of X on Y is estimable by adjustment using external data on Z_2 , i.e.:

$$P(y|do(\mathbf{x})) = \sum_{\mathbf{z}} P(y|\mathbf{x}, \mathbf{z}, S=1)P(z_1, z_3|z_2, S=1)P(z_2) \quad (5)$$

This adjustment requires Z_1 to be used, but only its biased measurements. If external measurements on Z_1 are included the adjustment is no longer valid. This may be surprising since a biasing path between S and Y is opened when X_1 or Z_1 are observed. In fact, selection bias can be controlled by external measurements over Z_2 alone (refer to the appendix for the detailed derivation of (4) and (5)). The following definition extends the notion of adjustment to account for selection bias and external data:

Definition 7 (Adjustment Pair). *Given a causal diagram G augmented with selection variable S , disjoint sets of variables $\mathbf{X}, \mathbf{Y}, \mathbf{Z}$, and a set $\mathbf{Z}^T \subseteq \mathbf{Z}$, $(\mathbf{Z}, \mathbf{Z}^T)$ is said to be an*




Figure 2: Instance of adjustment with partial external data.

adjustment pair for recovering the causal effect of \mathbf{X} on \mathbf{Y} if for every model compatible with G it holds that:

$$P(y|do(\mathbf{x})) = \sum_{\mathbf{z}} P(y|\mathbf{x}, \mathbf{z}, S=1)P(\mathbf{z} \setminus \mathbf{z}^T | \mathbf{z}^T, S=1)P(\mathbf{z}^T) \quad (6)$$

Remark. *The expression given in Eq. (6) is a natural extension of Eq. (1) and it captures the orthogonal nature of confounding and selection biases while allowing for the use of unbiased data over a subset of the covariates. Furthermore, Eqs. (2) and (3) are special cases of (6) corresponding, respectively, to the pairs (\mathbf{Z}, \emptyset) and (\mathbf{Z}, \mathbf{Z}) .*

The following criterion determines whether a pair $(\mathbf{Z}, \mathbf{Z}^T)$ yields a valid adjustment:

Definition 8 (Generalized Adjustment Criterion Type 3 (GACT3)). *Given a causal diagram G augmented with selection variable S , disjoint sets of variables $\mathbf{X}, \mathbf{Y}, \mathbf{Z}$ and a set $\mathbf{Z}^T \subseteq \mathbf{Z}$; $(\mathbf{Z}, \mathbf{Z}^T)$ is an admissible pair relative to \mathbf{X}, \mathbf{Y} in G if:*

- (a) *No element in \mathbf{Z} is a descendant in $G_{\overline{\mathbf{X}}}$ of any $W \notin \mathbf{X}$ lying on a proper causal path from \mathbf{X} to \mathbf{Y} .*
- (b) *All non-causal paths in G from \mathbf{X} to \mathbf{Y} are blocked by \mathbf{Z} and S .*
- (c) *\mathbf{Z}^T d -separates \mathbf{Y} from S in the proper backdoor graph, i.e. $(\mathbf{Y} \perp\!\!\!\perp S | \mathbf{Z}^T)_{G_{\overline{\mathbf{X}}}}^{pbd}$.*

In other words, cond. (a) prevents causal paths to be compromised by conditioning on an element in \mathbf{Z} , (b) requires all non-causal paths to be blocked by \mathbf{Z} , and (c) ensures that the influence of the selection mechanism on the outcome is nullified by \mathbf{Z}^T .

The following theorem states that the pairs *admissible* by the graphical criterion in Def. 8 are exactly those that constitute *adjustment pairs* as in Def. 7.

Theorem 1 (Admissible Pairs are Adjustment Pairs). *$(\mathbf{Z}, \mathbf{Z}^T)$ is an adjustment pair for \mathbf{X}, \mathbf{Y} in G if and only if it is admissible by Def. 8.*

Corollary 2 (Causal Effects Recovery by Adjustment). *Let G be a causal diagram augmented with a variable S representing the selection mechanism. Let \mathbf{V} be the set of variables measured under selection bias, and $\mathbf{T} \subset \mathbf{V}$ the set of variables measured externally in the overall population. Consider disjoint sets of variables $\mathbf{X}, \mathbf{Y} \subseteq \mathbf{V}$, then the causal effect $P(y | do(\mathbf{x}))$ is recoverable from $\{P(\mathbf{v} | S=1), P(\mathbf{t})\}$ by the adjustment expression (6) while*

$\mathbf{Z}^T \subseteq \mathbf{T}$, in every model inducing G if and only if $(\mathbf{Z}, \mathbf{Z}^T)$ is an admissible pair relative to \mathbf{X}, \mathbf{Y} in G according to Def. 8.

Corollary 2 answers to the proposed task of obtaining causal effects by adjustment from the data assumed as input. This means that the causal effect of X on Y can be estimated if an admissible pair $(\mathbf{Z}, \mathbf{Z}^T \subseteq \mathbf{T})$ exists.

As noted, Eq. (6) reduces to expression (2) and (3) when considering pairs of the form (\mathbf{Z}, \emptyset) and (\mathbf{Z}, \mathbf{Z}) , respectively. By the same token, GACT1 and GACT2 are special cases of GACT3. That is, Def. 8 will be equivalent to Def. 5 when $\mathbf{Z}^T = \emptyset$ and equivalent to Def. 6 when $\mathbf{Z}^T = \mathbf{Z}$, as stated in the following propositions.

Proposition 1. *If $\mathbf{Z}^T = \emptyset$, then $GACT1 \iff GACT3$.*

Proposition 2. *If $\mathbf{Z}^T = \mathbf{Z}$, then $GACT2 \iff GACT3$.*

Fig. 3 summarizes the inputs and adjustment expressions associated with each criterion.

Finding Admissible Sets

Once the admissibility of adjustment pairs has been characterized, it's natural to ask how to find them systematically and efficiently. This task is especially relevant since factors such as feasibility, cost, and statistical power may be relevant when choosing one of such sets.

To illustrate the complexity of this task, suppose we want to list all possible adjustment sets for the causal diagram given in Fig. 4. It contains ℓ non-causal paths from X to Y . For any pair $(\mathbf{Z}, \mathbf{Z}^T)$ admissible in this model, \mathbf{Z} and \mathbf{Z}^T must contain at least one variable in every one of the ℓ paths. For path i , either V_{i1} , V_{i2} , or both should be in those sets. In total, there are 3^ℓ different \mathbf{Z} , and for each one of them there are 3^k sets \mathbf{Z}^T , where k is the number of paths that contain both variables in \mathbf{Z} . The possible admissible pairs are in the Cartesian product of those sets, which amounts to $O(3^{2\ell})$ possibilities. It is clear that any algorithm that aims to output all admissible sets will take exponential time. Hence, no efficient algorithm exists for this task. In order to ameliorate this problem, we consider a special complexity class called polynomial delay (Takata 2010). Algorithms belonging to this class have the special property that the time required to output the first solution (or indicate failure), and the time between the outputs of consecutive solutions, is polynomial in the size of the input.

We show next an alternative, equivalent version of the criterion given in Def. 8 that will prove useful to operate within the polynomial delay class.

Definition 9 (Generalized Adjustment Criterion Type 3 (Alternative) GACT3A). *Given a causal diagram G augmented with selection variable S , disjoint sets of variables $\mathbf{X}, \mathbf{Y}, \mathbf{Z}$ and a set $\mathbf{Z}^T \subseteq \mathbf{Z}$; $(\mathbf{Z}, \mathbf{Z}^T)$ is an admissible pair relative to \mathbf{X}, \mathbf{Y} in G if:*

- (a) $\mathbf{Z} \cap Dpcp(\mathbf{X}, \mathbf{Y}) = \emptyset$
- (b) $(\mathbf{Y} \perp\!\!\!\perp \mathbf{X} \mid \mathbf{Z}, S)_{G_{\mathbf{X}\mathbf{Y}}^{pbd}}$
- (c) $(\mathbf{Y} \perp\!\!\!\perp S \mid \mathbf{Z}^T)_{G_{\mathbf{X}\mathbf{Y}}^{pbd}}$

where $Dpcp(\mathbf{X}, \mathbf{Y}) = De((De(\mathbf{X})_{G_{\overline{\mathbf{X}}}} \setminus \mathbf{X}) \cap An(\mathbf{Y})_{G_{\overline{\mathbf{X}}}})$.

The set $Dpcp(\mathbf{X}, \mathbf{Y})$ was originally introduced in (van der Zander, Liskiewicz, and Textor 2014) to account for the set of descendants of variables that lie in a proper causal path from \mathbf{X} to \mathbf{Y} .

Proposition 3. *Def. 9 is equivalent to Def. 8.*

In fact, Def. 9 is appealing to our task since each of the conditions can be easily verified, algorithmically, in a graph.

The following definition will be used to describe a collection of sets that separate variables in a causal model, subject to subset and superset constraints:

Definition 10 (Family of Separators). *Let $\mathbf{X}, \mathbf{Y}, \mathbf{R}$ be disjoint sets of variables in a causal diagram G , and let $\mathbf{I} \subseteq \mathbf{R}$ be another set. Define*

$$\mathcal{Z}_G(\mathbf{X}, \mathbf{Y}) \langle \mathbf{I}, \mathbf{R} \rangle := \{ \mathbf{Z} \mid (\mathbf{X} \perp\!\!\!\perp \mathbf{Y} \mid \mathbf{Z})_G \text{ and } \mathbf{I} \subseteq \mathbf{Z} \subseteq \mathbf{R} \} \quad (7)$$

to be the family of all sets \mathbf{Z} that d -separate \mathbf{X} and \mathbf{Y} in G and contain all elements in \mathbf{I} but no element outside \mathbf{R} .

For convenience, let the set of viable candidates for adjustment be denoted and defined as:

$$\mathcal{C} = \mathbf{V} \setminus (\mathbf{X} \cup \mathbf{Y} \cup Dpcp(\mathbf{X}, \mathbf{Y})) \quad (8)$$

Using this notation, the families that satisfy the conditions of our criterion can be specified. For conditions (a) and (b):

$$\mathcal{Z}_{a,b} = \left\{ \mathbf{Z} \mid \mathbf{Z} \cup \{S\} \in \mathcal{Z}_{G_{\mathbf{X}\mathbf{Y}}^{pbd}}(\mathbf{X}, \mathbf{Y}) \langle \{S\}, \mathcal{C} \cup \{S\} \rangle \right\} \quad (9)$$

We would like our algorithm to take into account the availability of external data over a set of covariates \mathbf{T} . In order to obtain admissible pairs for which the adjustment is estimable using the input as in Corollary 2, the set \mathbf{T} is incorporated in the definition of the family for condition (c):

$$\mathcal{Z}_c = \mathcal{Z}_{G_{\mathbf{X}\mathbf{Y}}^{pbd}}(\{S\}, \mathbf{Y}) \langle \emptyset, \mathbf{T} \rangle. \quad (10)$$

Our task can be summarized as finding pairs in the set:

$$\mathcal{Z}_{a,b,c} = \left\{ (\mathbf{Z}, \mathbf{Z}^T) \in \mathcal{Z}_{a,b} \times \mathcal{Z}_c \mid \mathbf{Z}^T \subseteq \mathbf{Z} \right\} \quad (11)$$

Algorithm 1 presents the procedure LISTADJPAIRS that solves this problem, as well as auxiliary routines used by it. Specifically it may be used to:

1. Given external data $P(\mathbf{t})$, list all admissible pairs such that $\mathbf{Z}^T \subseteq \mathbf{T}$.
2. List all admissible pairs (by setting $\mathbf{T} = \mathbf{V} \setminus (\mathbf{X} \cup \mathbf{Y})$) such that scientists know what external data to measure.

Functions LISTSEPAB and LISTSEPC are modifications of the enumeration algorithm LISTSEP in (van der Zander, Liskiewicz, and Textor 2014). The function FINDSEP is also described in that paper, and works as follows: given a graph G , sets of variables $\mathbf{X}, \mathbf{Y}, \mathbf{I}, \mathbf{R}$, where $\mathbf{X}, \mathbf{Y}, \mathbf{R}$ are disjoint and $\mathbf{I} \subseteq \mathbf{R}$; FINDSEP is guaranteed to output a $\tilde{\mathbf{Z}} \in \mathcal{Z}_G(\mathbf{X}, \mathbf{Y}) \langle \mathbf{I}, \mathbf{R} \rangle$ whenever there exists a separator \mathbf{C} such that $\mathbf{I} \subseteq \mathbf{C} \subseteq \mathbf{R}$; otherwise it returns \perp denoting failure.

Proposition 4 (Correctness of LISTSEPC). *Given a graph G , a variable S , sets of variables $\mathbf{Y}, \mathbf{I}, \mathbf{R}, \mathbf{Z}$, where $\{S\}, \mathbf{Y}, \mathbf{Z}$ are disjoint and $\mathbf{I} \subseteq \mathbf{R} \subseteq \mathbf{Z}$; LISTSEPC outputs all pairs $(\mathbf{Z}, \mathbf{Z}^T)$, where $\mathbf{Z}^T \in \mathcal{Z}_G(\{S\}, \mathbf{Y}) \langle \mathbf{I}, \mathbf{R} \rangle$.*

Criterion	Input	Adjustment Expression			
			Biased Data		Unbiased Data
GACT1	$\{P(\mathbf{v} \mid S=1)\}$	$\sum_{\mathbf{z}}$	$P(\mathbf{y} \mid \mathbf{x}, \mathbf{z}, S=1)$	$P(\mathbf{z} \mid S=1)$	
GACT2	$\{P(\mathbf{v} \mid S=1), P(\mathbf{t})\}$	$\sum_{\mathbf{z}}$	$P(\mathbf{y} \mid \mathbf{x}, \mathbf{z}, S=1)$		$P(\mathbf{z})$
GACT3	$\{P(\mathbf{v} \mid S=1), P(\mathbf{t})\}$	$\sum_{\mathbf{z}}$	$P(\mathbf{y} \mid \mathbf{x}, \mathbf{z}, S=1)$	$P(\mathbf{z} \setminus \mathbf{z}^{\mathbf{T}} \mid \mathbf{z}^{\mathbf{T}}, S=1)$	$P(\mathbf{z}^{\mathbf{T}})$

Figure 3: Comparison of the Adjustment Types




Figure 4: Simple diagram where the number of different separators is exponential in the size of the graph

Proposition 5 (Correctness of LISTSEPAB). *Given a graph G , a variable S , sets of variables $\mathbf{X}, \mathbf{Y}, \mathbf{I}, \mathbf{R}, \mathbf{T}$, where $\mathbf{X}, \mathbf{Y}, \{S\}, \mathbf{R}$ are disjoint, $\mathbf{I} \subseteq \mathbf{R}$ and $\mathbf{T} \subseteq \mathbf{C}$; LISTSEPAB outputs all pairs in $\{(\mathbf{Z}, \mathbf{Z}^{\mathbf{T}}) \in \mathcal{Z}_G(\mathbf{X}, \mathbf{Y}) \langle \mathbf{I}, \mathbf{R} \rangle \times \mathcal{Z}_G(\{S\}, \mathbf{Y}) \langle \emptyset, \mathbf{T} \rangle \mid \mathbf{Z}^{\mathbf{T}} \subseteq \mathbf{Z}\}$.*

The following theorem states that the algorithm LISTADJPAIRS can solve the task proposed in this section.

Theorem 3 (Correctness of LISTADJPAIRS). *Given a graph G , disjoint sets $\mathbf{X}, \mathbf{Y}, \mathbf{T}$, and a selection variable S , LISTADJPAIRS outputs all admissible pairs $(\mathbf{Z}, \mathbf{Z}^{\mathbf{T}})$ relative to \mathbf{X}, \mathbf{Y} in G such that $\mathbf{Z}^{\mathbf{T}} \subseteq \mathbf{Z}$.*

It is worth noting that a straightforward adaptation of the algorithm LISTSEP (van der Zander, Liskiewicz, and Textor 2014) may be used to find sets in $\mathbf{Z} \in \mathcal{Z}_{a,b}$ and $\mathbf{Z}^{\mathbf{T}} \in \mathcal{Z}_c$. However, the condition $\mathbf{Z}^{\mathbf{T}} \subseteq \mathbf{Z}$ has to be verified so as to produce admissible pairs. One strategy could be to search for sets in $\mathcal{Z}_{a,b}$ first, and then, while a second run outputs each set in \mathcal{Z}_c , validate if it is a subset of any output from the first batch of sets. In the worst case, exponential time is required to output the first admissible pair. A better idea would be to search for sets in $\{\mathbf{Z}^{\mathbf{T}} \in \mathcal{Z}_c \mid \mathbf{Z}^{\mathbf{T}} \subseteq \mathbf{Z}\}$ as soon as some $\mathbf{Z} \in \mathcal{Z}_{a,b}$ is found, and then output pairs made of \mathbf{Z} and the outputs of the secondary search. While improving over the original strategy, it may be the case that, for some sets in $\mathcal{Z}_{a,b}$, there is no set in \mathcal{Z}_c , which would lead to an exponential waiting time to get the first output.

Prop. 6 and Thm. 4 show that LISTADJPAIRS is, in fact, able to achieve $O(n(n+m))$ delay by carefully combining the search for the components of the pairs, where n, m are the number of variables and edges in G , respectively.

Proposition 6 (Complexity of LISTSEPAB). LISTSEPAB works with $O(n(n+m))$ delay.

Theorem 4 (Complexity of LISTADJPAIRS). LISTADJPAIRS outputs all admissible pairs such that $\mathbf{Z}^{\mathbf{T}} \subseteq \mathbf{T}$ with $O(n(n+m))$ polynomial delay.

Algorithm 1 Routines used to list admissible pairs

```

1: function LISTADJPAIRS( $G, \mathbf{X}, \mathbf{Y}, S, \mathbf{V}, \mathbf{T}$ )
2:    $G_{\mathbf{X}\mathbf{Y}}^{pbd} \leftarrow$  Compute proper backdoor graph from  $G$ 
3:    $\mathbf{R} \leftarrow (\mathbf{V} \cup \{S\}) \setminus (\mathbf{X} \cup \mathbf{Y} \cup D_{pcp}(\mathbf{X}, \mathbf{Y}))$ 
4:   LISTSEPAB( $G_{\mathbf{X}\mathbf{Y}}^{pbd}, \mathbf{X}, \mathbf{Y}, S, \{S\}, \mathbf{R}, \mathbf{T}$ )
5: end function
6: function LISTSEPAB( $G, \mathbf{X}, \mathbf{Y}, S, \mathbf{I}, \mathbf{R}, \mathbf{T}$ )
7:   if FINDSEP( $G, \mathbf{X}, \mathbf{Y}, \mathbf{I}, \mathbf{R}$ )  $\neq \perp$  ^
8:     FINDSEP( $G, \{S\}, \mathbf{Y}, \emptyset, \mathbf{R} \cap \mathbf{T}$ )  $\neq \perp$  then
9:       if  $\mathbf{I} = \mathbf{R}$  then
10:        LISTSEPC( $G, S, \mathbf{Y}, \emptyset, \mathbf{I} \cap \mathbf{T}, \mathbf{I} \setminus \{S\}$ )
11:       else
12:         $V \leftarrow$  arbitrary variable from  $\mathbf{R} \setminus \mathbf{I}$ 
13:        LISTSEPAB( $G, \mathbf{X}, \mathbf{Y}, \mathbf{I} \cup \{V\}, \mathbf{R}, \mathbf{T}$ )
14:        LISTSEPAB( $G, \mathbf{X}, \mathbf{Y}, \mathbf{I}, \mathbf{R} \setminus \{V\}, \mathbf{T}$ )
15:       end if
16:     end if
17:   end function
18: function LISTSEPC( $G, S, \mathbf{Y}, \mathbf{I}, \mathbf{R}, \mathbf{Z}$ )
19:   if FINDSEP( $G, \{S\}, \mathbf{Y}, \mathbf{I}, \mathbf{R}$ )  $\neq \perp$  then
20:     if  $\mathbf{I} = \mathbf{R}$  then
21:       output  $(\mathbf{Z}, \mathbf{I})$ 
22:     else
23:        $V \leftarrow$  arbitrary variable from  $\mathbf{R} \setminus \mathbf{I}$ 
24:       LISTSEPC( $G, \mathbf{X}, \mathbf{Y}, \mathbf{I} \cup \{V\}, \mathbf{R}, \mathbf{Z}$ )
25:       LISTSEPC( $G, \mathbf{X}, \mathbf{Y}, \mathbf{I}, \mathbf{R} \setminus \{V\}, \mathbf{Z}$ )
26:     end if
27:   end if
28: end function

```

Using covariates from $An(\mathbf{X}\cup\mathbf{Y})$ is sufficient to block any biasing path when controlling for confounding bias, which does not hold when selection bias comes into play. The proposition below constitutes a natural extension of this result when searching for adjusting pairs, in particular, considering the set $An(\mathbf{X}\cup\mathbf{Y}\cup\{S\})$.

Proposition 7. *Suppose a pair $(\mathbf{Z}, \mathbf{Z}^{\mathbf{T}})$ is admissible relative to \mathbf{X}, \mathbf{Y} in G . Then, the pair $(\mathbf{Z}_A, \mathbf{Z}_A^{\mathbf{T}})$, where $\mathbf{Z}_A^{\mathbf{T}} = \mathbf{Z}^{\mathbf{T}} \cap An(\mathbf{X}\cup\mathbf{Y}\cup\{S\})$ and $\mathbf{Z}_A = \mathbf{Z} \cap An(\mathbf{X}\cup\mathbf{Y}\cup\{S\})$, is also admissible.*

If the data scientist is not interested in deciding among different adjustment pairs to use, it is possible to explicitly construct a pair if one exists, namely:

Theorem 5 (Explicit admissible set construction). *There exists an admissible pair in a causal diagram G relative to disjoint sets of variables \mathbf{X}, \mathbf{Y} if and only if the pair $(\mathbf{Z}, \mathbf{Z}^{\mathbf{T}})$*

is admissible, where

$$\mathbf{Z} = An(\mathbf{X} \cup \mathbf{Y} \cup \{S\})_{G_{\mathbf{X}\mathbf{Y}}^{pbd}} \cap \mathcal{C} \quad (12)$$

$$\mathbf{Z}^T = (An(\{S\} \cup \mathbf{Y}))_{G_{\mathbf{X}\mathbf{Y}}^{pbd}} \cap \mathbf{T} \cap \mathcal{C} \quad (13)$$

Corollary 6 (Admissible pair can be constructed in linear time). *One can determine the existence of an admissible pair and construct one in $O(n + m)$ time.*

Inverse Probability Weighting Estimation

Covariate adjustment is currently the most widely used method for causal effect estimation in practice, even when more powerful identification methods have been developed in recent years (Pearl 2000). Adjusting for covariates \mathbf{Z} involves finding the conditional probability of \mathbf{Y} given \mathbf{X} for each stratum defined by the possible values of the covariates, which may present computational and sample complexity challenges. The number of different strata may grow rapidly with the cardinality of the set \mathbf{Z} , and the number of samples falling under each stratum may be too small to provide a reliable estimate of the conditional distribution.

There exist robust statistical estimation procedures for the adjustment expression (1) that circumvent this issue. To do so, one can rewrite the adjustment expression (1) as $P(\mathbf{y} \mid do(\mathbf{x})) = \sum_{\mathbf{z}} P(\mathbf{y}, \mathbf{x}, \mathbf{z})/P(\mathbf{x} \mid \mathbf{z})$. If a reliable estimate of the conditional distribution $P(\mathbf{x} \mid \mathbf{z})$ could be obtained, which is known as the ‘‘propensity score’’ (Pearl, Glymour, and Jewell 2016), then the causal effect could be estimated by ‘‘weighting’’ every observed sample by the factor $1/P(\mathbf{x} \mid \mathbf{z})$, leading to the widely used ‘‘inverse probability weighed (IPW) estimator’’ (Lunceford and Davidian 2004). Assume we are interested in the mean causal effect $\mu = E[\mathbf{Y} \mid do(\mathbf{x})]$. Given observed i.i.d. data $\{(\mathbf{X}_i, \mathbf{Y}_i, \mathbf{Z}_i)\}_{i=1}^n$, the IPW estimator for μ is given by

$$\hat{\mu}_{IPW} = \frac{1}{n} \sum_{i=1}^n w_i I_{\mathbf{X}_i=\mathbf{x}} \mathbf{Y}_i \quad (14)$$

where $I_{\mathbf{X}_i=\mathbf{x}}$ is the indicator function, $w_i = 1/\hat{P}(\mathbf{X}_i \mid \mathbf{Z}_i)$, and $\hat{P}(\mathbf{X}_i \mid \mathbf{Z}_i)$ is the estimator of the propensity score. In practice, $\hat{P}(\mathbf{X}_i \mid \mathbf{Z}_i)$ is estimated from data by assuming some parametric model (typically a logistic regression model). $\hat{\mu}_{IPW}$ is a consistent estimator for μ if the model for $P(\mathbf{x} \mid \mathbf{z})$ is correctly specified.

Next, we show that IPW style estimator could be constructed for causal effect estimation in the presence of selection bias using the generalized adjustment given in this paper. We rewrite the adjustment expression (6) as follows:

$$\begin{aligned} P(\mathbf{y} \mid do(\mathbf{x})) &= \sum_{\mathbf{z}} P(\mathbf{y} \mid \mathbf{x}, \mathbf{z}, S=1) P(\mathbf{z} \setminus \mathbf{z}^T \mid \mathbf{z}^T, S=1) P(\mathbf{z}^T) \\ &= \sum_{\mathbf{z}} \frac{P(\mathbf{y}, \mathbf{x}, \mathbf{z} \mid S=1)}{P(\mathbf{x} \mid \mathbf{z}, S=1)} \frac{P(\mathbf{z}^T)}{P(\mathbf{z}^T \mid S=1)} \end{aligned} \quad (15)$$

$$= \sum_{\mathbf{z}} \frac{P(\mathbf{y}, \mathbf{x}, \mathbf{z} \mid S=1)}{P(\mathbf{x} \mid \mathbf{z}, S=1)} \frac{P(S=1)}{P(S=1 \mid \mathbf{z}^T)} \quad (16)$$

The quotient $P(\mathbf{z}^T)/P(\mathbf{z}^T \mid S=1)$ (in Eq. (15)), which is directly computable from the combination of the external and biased datasets, can be equivalently expressed

as $P(S=1)/P(S=1 \mid \mathbf{z}^T)$ (in Eq. (16)). The later is usually known as the ‘‘inverse probability-of-selection weight (IPSW)’’ (Cole and Stuart 2010), and, in practice, is estimated by assuming some parametric model such as logistic regression. Given observed data $\{(\mathbf{X}_i, \mathbf{Y}_i, \mathbf{Z}_i)\}_{i=1}^n$ under selection bias (from $P(\mathbf{v} \mid S=1)$), assume we could obtain reliable estimate of the propensity score $P(\mathbf{x} \mid \mathbf{z}, S=1)$ and the inverse probability-of-selection $P(S=1)/P(S=1 \mid \mathbf{z}^T)$ from selection biased data and additional unbiased external data. Let $w'_i = 1/\hat{P}(\mathbf{X}_i \mid \mathbf{Z}_i, S=1)$ and $w_i^S = \hat{P}(S=1)/\hat{P}(S=1 \mid \mathbf{Z}_i^T)$. The causal effect can be estimated by first weighting every observed sample under selection bias by the weight $w'_i \cdot w_i^S$. We propose the following estimator for μ

$$\hat{\mu}_{IPWS} = \frac{1}{n} \sum_{i=1}^n w'_i w_i^S I_{\mathbf{X}_i=\mathbf{x}} \mathbf{Y}_i. \quad (17)$$

Theorem 7. $\hat{\mu}_{IPWS}$ is a consistent estimator for $\mu = E[\mathbf{Y} \mid do(\mathbf{x})]$ if the models for $P(\mathbf{x} \mid \mathbf{z}, S=1)$ and $P(S=1)/P(S=1 \mid \mathbf{z}^T)$ are correctly specified.

Further, whenever $\mathbf{Z}^T = \emptyset$, the IPW estimator for the adjustment expression (2) is given by

$$\hat{\mu}_{IPW2} = \frac{1}{n} \sum_{i=1}^n w'_i I_{\mathbf{X}_i=\mathbf{x}} \mathbf{Y}_i \quad (18)$$

Note that in (18) the samples only need to be weighted by the propensity score but do not need to be weighted by the IPSW in order to adjust for selection bias. The difference between (18) and (14) is that in (18) the samples are observed under selection bias. One of the contributions of this paper is that we specify conditions over the adjustment set on how the biased data samples should be weighted in order to obtain unbiased estimates of causal effects.

Conclusions

This work generalizes the notion of adjustment set to that of *adjustment pairs* (Def. 7), that when admissible, recover causal effects via adjustment from a distribution under *selection bias* and auxiliary external data, while simultaneously controlling for *confounding bias*. We present a sufficient and necessary graphical condition (Def. 8) to determine if a pair is admissible, valid for any causal diagram G with latent variables in non-parametric settings. We develop the algorithm LISTAJDPAIRS that lists all admissible pairs for given \mathbf{X}, \mathbf{Y} in G (Theorem 3) with polynomial delay (Theorem 4). These results allow scientists to take into consideration the effort of measuring covariates, such as associated cost, availability, or feasibility. Finally, we describe how to use the inverse probability weighting technique to estimate adjustment under selection bias. Adjustment is not the only method to estimate causal effects, but it is still the most popular one in the empirical sciences. We hope the results presented in this paper will help the broad scientific community to account for selection bias in their studies.

Acknowledgments

This research was supported in parts by grants from NSF [IIS-1704352] and DARPA [ARO W911NF-16-1-0579].

References

- Angrist, J. D. 1997. Conditional independence in sample selection models. *Economics Letters* 54(2):103–112.
- Bareinboim, E., and Pearl, J. 2012a. Causal Inference by Surrogate Experiments: z-Identifiability. In Murphy, and Freitas., eds., *Proceedings of the 28th Conference on Uncertainty in Artificial Intelligence*, 113–120. AUAI Press.
- Bareinboim, E., and Pearl, J. 2012b. Controlling Selection Bias in Causal Inference. In Lawrence, and Girolami., eds., *Proc. of the 15th International Conference on Artificial Intelligence and Statistics (AISTATS)*. JMLR. 100–108.
- Bareinboim, E., and Pearl, J. 2016. Causal inference and the data-fusion problem. *Proceedings of the National Academy of Sciences* 113(27):7345–7352.
- Bareinboim, E.; Tian, J.; and Pearl, J. 2014. Recovering from Selection Bias in Causal and Statistical Inference. In Brodley, C. E., and Stone, P., eds., *Proceedings of the Twenty-eighth AAAI Conference on Artificial Intelligence*, 2410–2416. Palo Alto, CA: AAAI Press.
- Cole, S. R., and Stuart, E. A. 2010. Generalizing evidence from randomized clinical trials to target populations: The ACTG 320 trial. *American Journal of Epidemiology* 172(1):107–115.
- Cooper, G. 1995. Causal Discovery from Data in the Presence of Selection Bias. In *Fifth International Workshop on Artificial Intelligence and Statistics*. 5–7.
- Correa, J. D., and Bareinboim, E. 2017. Causal Effect Identification by Adjustment under Confounding and Selection Biases. *Proceedings of the 31th Conference on Artificial Intelligence (AAAI 2017)* 3740–3746.
- Correa, J. D.; Tian, J.; and Bareinboim, E. 2017. Adjustment under confounding and selection bias. Technical report, R-29, Purdue AI Lab, Dep. of Computer Science, Purdue University.
- Cortes, C.; Mohri, M.; Riley, M.; and Rostamizadeh, A. 2008. Sample Selection Bias Correction Theory. In *Proceedings of the 19th International Conference on Algorithmic Learning Theory*, 38–53.
- Elkan, C. 2001. The foundations of cost-sensitive learning. In *Proc. of the 17th IJCAI*, volume 17, 973–978.
- Evans, R. J., and Didelez, V. 2015. Recovering from selection bias using marginal structure in discrete models. In *Proc. of UAI-15 Workshop Advances in Causal Inference*, ACI'15, 46–55. Aachen, Germany: CEUR-WS.org.
- Glymour, M. M., and Greenland, S. 2008. Causal Diagrams. In Lash, R., and Greenland., eds., *Modern Epidemiology*. Lippincott Williams & Wilkins, 3rd edition. 183–209.
- Heckman, J. J. 1979. Sample Selection Bias as a Specification Error. *Econometrica* 47(1):153.
- Huang, Y., and Valtorta, M. 2006. Pearl's Calculus of Intervention Is Complete. In Richardson, and Dechter., eds., *Proceedings of the 22nd Conference on Uncertainty in Artificial Intelligence*, 217–224. Corvallis, OR: AUAI Press.
- Kuroki, M., and Cai, Z. 2006. On recovering a population covariance matrix in the presence of selection bias. *Biometrika* 93(3):601–611.
- Little, R., and Rubin, D. 1987. *Statistical analysis with missing data*. New York, NY, USA: Wiley & Sons, Inc.
- Lunceford, J. K., and Davidian, M. 2004. Stratification and weighting via the propensity score in estimation of causal treatment effects: a comparative study. *Stat Med* 23(19):2937–2960.
- Mefford, J., and Witte, J. S. 2012. The covariate's dilemma. *PLoS Genet* 8(11):e1003096.
- Pearl, J., and Paz, A. 2010. Confounding equivalence in causal equivalence. In *Proc. of the 26th Conference on Uncertainty in Artificial Intelligence*. AUAI Press. 433–441.
- Pearl, J.; Glymour, M.; and Jewell, N. P. 2016. *Causal inference in statistics: A Primer*. John Wiley & Sons.
- Pearl, J. 1993. Aspects of graphical models connected with causality. *Proceedings of the 49th Session of the International Statistical Institute* (August):399–401.
- Pearl, J. 1995. Causal diagrams for empirical research. *Biometrika* 82(4):669–688.
- Pearl, J. 2000. *Causality: models, reasoning, and inference*. NY, USA: Cambridge University Press, 2nd edition.
- Pirinen, M.; Donnelly, P.; and Spencer, C. C. A. 2012. Including known covariates can reduce power to detect genetic effects in case-control studies. *Nature Genetics* 44(8).
- Robins, J. M. 2001. Data, Design, and Background Knowledge in Etiologic Inference. *Epidemiology* 12(3):313–320.
- Robinson, L. D., and Jewell, N. P. 1991. Some Surprising Results about Covariate Adjustment in Logistic Regression Models. *International Statistical Review / Revue Internationale de Statistique* 59(2):227–240.
- 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*, 1219–1226.
- Shpitser, I.; VanderWeele, T. J.; and Robins, J. M. 2010. On the validity of covariate adjustment for estimating causal effects. *Proceedings of the Twenty Sixth Conference on Uncertainty in Artificial Intelligence (UAI-10)* 527–536.
- Spirites, P.; Glymour, C. N.; and Scheines, R. 2001. *Causation, Prediction, and Search*. MIT Press, 2nd edition.
- Takata, K. 2010. Space-optimal, backtracking algorithms to list the minimal vertex separators of a graph. *Discrete Applied Mathematics* 158(15):1660–1667.
- Tian, J., and Pearl, J. 2002. A General Identification Condition for Causal Effects. In *Proceedings of the Eighteenth National Conference on Artificial Intelligence (AAAI 2002)*, number August, 567–573. Menlo Park, CA: AAAI Press.
- van der Zander, B.; Liskiewicz, M.; and Textor, J. 2014. Constructing separators and adjustment sets in ancestral graphs. In *Proceedings of UAI 2014*, 907–916.
- Whittemore, A. 1978. Collapsibility of multidimensional contingency tables. *Journal of the Royal Statistical Society. Series B* 40(3):328–340.
- Zadrozny, B. 2004. Learning and evaluating classifiers under sample selection bias. In *Proc. of the 21th International Conference on Machine learning*, 114–121.