

Bounding Causal Effects on Continuous Outcome

Junzhe Zhang, Elias Bareinboim

Causal Artificial Intelligence Laboratory
Columbia University
{junzhez,eb}@cs.columbia.edu

Abstract

We investigate the problem of bounding causal effects from experimental studies in which treatment assignment is randomized but the subject compliance is imperfect. It is well known that under such conditions, the actual causal effects are not point-identifiable due to uncontrollable unobserved confounding. In their seminal work, Balke and Pearl (1994) derived the tightest bounds over the causal effects in this settings by employing an algebra program to derive analytic expressions. However, Pearl’s approach assumes the primary outcome to be discrete and finite. Solving such a program could be intractable when high-dimensional context variables are present. In this paper, we present novel non-parametric methods to bound causal effects on the continuous outcome from studies with imperfect compliance. These bounds could be generalized to settings with the high-dimensional context.

Introduction

One of the most common methods for policy learning used throughout the empirical sciences is the use of randomization of the treatment assignment. This method is considered the gold standard within many disciplines and can be traced back, at least, to Fisher (Fisher 1935) and Neyman (Neyman 1923). Whenever human subjects are at the center of the experiment, unfortunately, issues of *non-compliance* arise, namely, subjects do not necessarily follow the experimental protocol and end up doing what they want. It is well-understood that under such conditions, *confounding bias* will emerge. For instance, subjects who did not comply with the treatment assignment may be precisely those who would have responded adversely to the treatment. Therefore, the actual *causal effects* of the treatment, when it is applied uniformly to the population, might be substantially less effective than of the data reveals.

To cope with this bias, analysts may resort to exploit theoretical assumptions underlying the interactions between compliance and response (Wright 1928; Angrist, Imbens, and Rubin 1996). The problem of identifying causal effects from observed data provided with causal assumptions about the data-generating mechanisms, represented in the form of a *directed acyclic causal diagram* (Pearl 2000, Ch. 1.2), has been extensively studied in the causal inference literature.

Several criteria and algorithms have been developed (Pearl 2000; Spirtes, Glymour, and Scheines 2000; Bareinboim and Pearl 2016). For example, a criterion called *back-door* (Pearl 2000, Ch. 3.2.2) permits one to determine whether causal effects can be obtained by covariate adjustment and subsequent inverse probability weighting. This condition is also known as *conditional ignorability* and *unconfoundedness* (Rosenbaum and Rubin 1983). Efficient estimators were developed based on the propensity score (Rosenbaum and Rubin 1983; Bang and Robins 2005) and off-policy learning (Dudík, Langford, and Li 2011; Li, Munos, and Szepesvari 2015; Munos et al. 2016; Thomas and Brunskill 2016).

By and large, the combination of causal assumptions and observational data does not always allow one to point-identify the causal effect, called the *non-identifiable*. That is, there exists more than one parametrization of the target effect that are compatible with the same observational data and qualitative assumptions (Pearl 2000, Def. 3.2.2). A causal effect is partially identifiable if it is not identifiable, but the set of its possible values is smaller than the original parameter space. Inferring about the treatment effect in the partially identifiable settings has been a target of growing interest in the domains of causal inference (Balke and Pearl 1995; Chickering and Pearl 1996; Richardson et al. 2014; Cinelli et al. 2019), and more recently, in machine learning (Kallus and Zhou 2018; Kallus, Puli, and Shalit 2018). Among these works, two approaches are often employed: (1) bounds are derived for the target effect under minimal assumptions; or (2) additional untestable assumptions are invoked under which the causal effect is identifiable, and then sensitivity analysis is conducted to assess how the target causal effect varies as the untestable assumptions are changed. This paper focuses on the bounding approach.

(Robins 1989; Manski 1990) derived the first informative bounds over the causal effects from studies with imperfect compliance, under a set of non-parametric assumptions called *instrumental variables*. In their seminal work (Balke and Pearl 1994a, 1997), Balke and Pearl improved earlier results by employing a computer algebraic program to derive analytic expressions of the causal bounds, which are provably optimal. Despite the optimality guarantees provided in their treatment, there are still significant challenges in performing the partial identification of the causal effects with the presence of instrumental variables. First, Pearl’s ap-

proaches assume the outcome is discrete and finite, which is often not the case in many practical applications. Second, in settings with the high-dimensional context, solving the formulated program is often intractable due to computational and sample complexity issues.

The goal of this paper is to overcome these challenges. We investigate the partial identification of the causal effect on the continuous outcome, with the presence of instrumental variables and the high-dimensional context. More specifically, our contributions are as follows. (1) We identify a set of novel non-parametric assumptions that explicate the inherent independence relationships among the latent counterfactual variables (also called the potential outcomes) when instrumental variables are present. (2) Using the proposed model, we formulate the linear programs that bound the target causal effect on the continuous outcome from studies with imperfect compliance, which is provable optimal. (3) We provide efficient estimation procedures for the derived bounds from finite observational sample under high-dimensional context. Finally, we apply the derived causal bounds to various bandit learning algorithms (Gittins 1979), showing that they could consistently improve the convergence for identifying the optimal treatment. Our results are validated on the International Stroke Trial data (Carolei et al. 1997). Given the space constraints, all proofs are provided in the full technical report (Zhang and Bareinboim 2020).

Preliminaries

In this section, we introduce the basic notations and definitions used throughout the paper. We use capital letters to denote variables (X) and small letters for their values (x). Let $\mathcal{X}$ stand for the domain of X and $|\mathcal{X}|$ for its dimension. We use $P(x)$ to represent probabilities $P(X = x)$.

The basic semantical framework of our analysis rests on *structural causal models* (SCM) (Pearl 2000, Ch. 7). A SCM M is a tuple $\langle \mathbf{U}, \mathbf{V}, \mathcal{F}, P(\mathbf{u}) \rangle$, where $\mathbf{U}$ is a set of exogenous (unobserved) variables and $\mathbf{V}$ is a set of endogenous (observed) variables. $\mathcal{F}$ is a set of structural functions where $f_{V_i} \in \mathcal{F}$ decides the values of $V_i \in \mathbf{V}$ taking as argument a combination of other endogenous and exogenous variables (i.e., $V_i \leftarrow f_{V_i}(Pa_{V_i}, U_{V_i})$, $Pa_{V_i} \subseteq \mathbf{V}$, $U_{V_i} \subseteq \mathbf{U}$). The values of $\mathbf{U}$ are drawn from the distribution $P(\mathbf{u})$, inducing an observational distribution $P(\mathbf{v})$ over $\mathbf{V}$. Each SCM is associated with a causal diagram in the form of a directed acyclic graph G , where nodes represent variables and arrows stand for functional relationships (e.g., see Fig. 1). By convention, whenever clear from the context, the exogenous $\mathbf{U}$ are left implicit. The bi-directed arrows between V_i and V_j indicate the existence of an unobserved confounder (UC) U_k affecting both V_i and V_j , i.e., $V_i \leftarrow U_k \rightarrow V_j$.

An intervention on a set of endogenous variables $\mathbf{X}$, denoted by $do(\mathbf{x})$, is an operation where values of $\mathbf{X}$ are set to constants $\mathbf{x}$, regardless of how they were ordinarily determined (through the functions $\{f_X : \forall X \in \mathbf{X}\}$). For a SCM M , let M_x be a modified sub-model of M under intervention $do(\mathbf{x})$. The potential outcome of Y to intervention $do(\mathbf{x})$, denoted by $Y_x(\mathbf{u})$, is the solution for Y with $\mathbf{U} = \mathbf{u}$ in the sub-model M_x ; it can be read as the counterfactual sentence “the value that Y would have obtained in




Figure 1: Causal diagram of the instrumental variable (IV) model: Z represents the (randomized) treatment assigned, X the treatment actually received, and Y the outcome.

situation $\mathbf{U} = \mathbf{u}$, had $\mathbf{X}$ been $\mathbf{x}$.” Statistically, averaging $\mathbf{u}$ over the distribution $P(\mathbf{u})$ leads to the interventional distribution $P(y_x)$. For a detailed survey on the structural causal models, we refer readers to (Pearl 2000, Ch. 7).

One fundamental problem in causal inference is to estimate $P(y_x)$ from the combination of the observational distribution $P(\mathbf{v})$ and causal diagram G . An interventional distribution $P(y_x)$ is identifiable from G if for any pair of SCMs M_1 and M_2 compatible with G , $P_{M_1}(y_x) = P_{M_2}(y_x)$ whenever $P_{M_1}(\mathbf{v}) = P_{M_2}(\mathbf{v})$ (Pearl 2000, pp. 77). In other words, $P(y_x)$ are non-identifiable if there exists a pair of SCMs that give arise to the same $P(\mathbf{v})$ and G but induce different distributions $P(y_x)$.

New Bounds on Causal Effects

We will focus on the a special type of SCM called the *instrumental variable* (IV) models which represent experimental studies with imperfect compliance (Pearl 2000, Ch. 8.2). Fig. 1a shows the causal diagram of the IV model where Z represents the (randomized) treatment assigned, X the treatment actually received, and Y the observed outcome; exogenous variables $\mathbf{U}$ summarize the unknown factors about an individual subject that affect both X and Y . The values of Y are continuous, decided by a function $y \leftarrow f_Y(x, \mathbf{u})$ bounded in $[0, 1]$. We assume Z, X are both finite. For each $Z = z$, values of X are decided by an unknown mechanism $x \leftarrow f_X(z, \mathbf{u})$. The data collected from the studies are summarized as the observational distribution $P(x, y|z)$.

Given $P(x, y|z)$, we are interested in inferring the expected outcome on Y by of performing a treatment $do(x)$, i.e., the causal effect $E[Y_x]$. Fig. 1 graphically describes this learning settings. Unfortunately, the non-identifiability of the treatment effect $E[Y_x]$ from the surrogate Z and UCs between X and Y was shown in (Bareinboim and Pearl 2012). To overcome this challenge, we will consider the problem of partial identification in IV models (Manski 2003). Instead of pin-pointing the target quantity, the goal of partial identification is to derive bounds on the parameter space of the causal effect $E[Y_x]$ from the observational data $P(x, y|z)$, called the *causal bound*.

Restricted Instrumental Variable Models

Let $\mathcal{M}_{IV}[P(x, y|z)]$ denote a set of IV models described in Fig. 1a which are compatible with distribution $P(x, y|z)$; therefore, for any $M \in \mathcal{M}_{IV}[P(x, y|z)]$, $P_M(x, y|z) = P(x, y|z)$. We could derive causal bounds $E[Y_x] \in [l_x, h_x]$

by solving the optimization problems as follows:

$$\begin{aligned} l_x &= \min_{M \in \mathcal{M}_{\text{OBS}}} E_M[Y_x] \\ h_x &= \max_{M \in \mathcal{M}_{\text{OBS}}} E_M[Y_x] \end{aligned} \quad \left| \quad \mathcal{M}_{\text{OBS}} = \mathcal{M}_{\text{IV}}[P(x, y|z)] \right. \quad (1)$$

The challenge of solving Eq. (1) is that the parametric forms of the exogenous variables $\mathbf{U}$ and structural functions $\mathcal{F}$ are not explicitly specified. $\mathcal{M}_{\text{IV}}[P(x, y|z)]$ could be infinitely large, making it hard to derive the bounds $[l_x, h_x]$.

We will provide efficient methods to overcome this challenge. In particular, we propose a new non-parametric representation for the pair $\mathbf{U}, \mathcal{F}$. Let $Y_{\mathcal{X}}$ denote a vector of potential responses $(Y_{x_0}, \dots, Y_{x_{|\mathcal{X}|-1}})$ where each $x_i \in \mathcal{X}$ and let $y_{\mathcal{X}}$ be its realizations; $X_{\mathcal{Z}} = x_{\mathcal{Z}}$ is similarly defined. We first describe a family of IV models where $\mathbf{U}, \mathcal{F}$ are well-specified from counterfactuals $X_{\mathcal{Z}}, Y_{\mathcal{X}}$.

Definition 1 (Restricted IV (RIV) Model). A restricted instrumental variable model is a SCM $\langle \mathbf{U}, \mathbf{V}, \mathcal{F}, P(\mathbf{u}) \rangle$ in Fig. 1a where $\mathbf{V} = \{X, Y, Z\}$, $\mathbf{U} = \{X_{\mathcal{Z}}, Y_{\mathcal{X}}\}$. Given a vector $x_{\mathcal{Z}}$, x_z is an element in $x_{\mathcal{Z}}$ at the index $z_i = z$; similarly, y_x is an element in $y_{\mathcal{X}}$ at $x_i = x$. Values of X, Y are decided by functions $f_X, f_Y \in \mathcal{F}$ defined as:

$$x \leftarrow f_X(z, x_{\mathcal{Z}}) = x_z, \quad y \leftarrow f_Y(x, y_{\mathcal{X}}) = y_x, \quad (2)$$

Y_{x_i} are mutually independent given $X_{\mathcal{Z}}$, i.e., for any x_i ,

$$Y_{x_i} \perp\!\!\!\perp \{Y_{x_j} : \forall x_j \neq x_i\} | X_{\mathcal{Z}} \quad (3)$$

At first glance, the conditional independence among $Y_{\mathcal{X}}$ in Def. 1 may be surprising since it seems to impose additional constraints about the exogenous $\mathbf{U}$ in the original IV model. We will show in sequel that this restriction indeed captures the natural properties of the optimization problem in Eq. (1). Fig. 2a shows the graphical representation of a RIV model. The square labeled with $\mathcal{X}$ indicates that there are $|\mathcal{X}|$ nodes Y_{x_i} of this kind. The counterfactuals $X_{\mathcal{Z}}, Y_{\mathcal{X}}$ are the exogenous variables $\mathbf{U}$ affecting the treatment X and outcome Y , respectively. $X_{\mathcal{Z}}, Y_{\mathcal{X}}$ are correlated; each potential reward node Y_{x_i} is d-separated from other nodes Y_{x_j} where $x_j \neq x_i$ given $X_{\mathcal{Z}}$ (Pearl 2000, Def. 1.2.3). The counterfactual distribution $P(x_{\mathcal{Z}}, y_{\mathcal{X}})$ can be written as

$$P(x_{\mathcal{Z}}, y_{\mathcal{X}}) = P(x_{\mathcal{Z}}) \prod_{x_i \in \mathcal{X}} P(y_{x_i} | x_{\mathcal{Z}}).$$

Let $\mathcal{M}_{\text{RIV}}[P(x, y|z)]$ denote a set of RIV models compatible with $P(x, y|z)$. We will show that solving Eq. (1) is equivalent to optimizing $E[Y_x]$ over the feasible region $\mathcal{M}_{\text{OBS}} = \mathcal{M}_{\text{RIV}}[P(x, y|z)]$.

Theorem 1. *Given $P(x, y|z)$, for any IV model $M_1 \in \mathcal{M}_{\text{IV}}[P(x, y|z)]$, there exists an IV model $M_2 \in \mathcal{M}_{\text{RIV}}[P(x, y|z)]$ such that (s.t.) $E_{M_1}[Y_x] = E_{M_2}[Y_x]$, and vice versa.*

Thm. 1 says that for any IV model M of Fig. 1a inducing the observational data $P(x, y|z)$, we could always reduce it into a RIV model in $\mathcal{M}_{\text{RIV}}[P(x, y|z)]$ while preserving its treatment effects $E[Y_x]$. Optimizing Eq. (1) within the feasible region $\mathcal{M}_{\text{OBS}} = \mathcal{M}_{\text{RIV}}[P(x, y|z)]$ thus induces causal bounds $E[Y_x] \in [l_x, h_x]$. The sharpness of $[l_x, h_x]$ follows immediately from Def. 1, i.e., there exist SCMs $M_1, M_2 \in \mathcal{M}_{\text{IV}}[P(x, y|z)]$ such that $E_{M_1}[Y_x] = l_x, E_{M_2}[Y_x] = h_x$.




Figure 2: Causal diagrams of (a) a RIV model where Y_{x_i} are mutually independent given $X_{\mathcal{Z}}$; (b) an unrestricted RIV model. The square labelled with $\mathcal{X}$ represents $|\mathcal{X}|$ nodes of which only a single example Y_{x_i} is shown explicitly.

A Linear Program Formulation

We now turn our attention to solving the optimization problem in Eq. (1) within the feasible region $\mathcal{M}_{\text{OBS}} = \mathcal{M}_{\text{RIV}}[P(x, y|z)]$. We will use probabilities $P(x_{\mathcal{Z}})$ and $E[Y_{x_i} | x_{\mathcal{Z}}]P(x_{\mathcal{Z}})$ as unknown parameters. Basic probabilistic properties and $y \leftarrow f_Y(x, \mathbf{u}) \in [0, 1]$ imply:

$$\sum_{x_{\mathcal{Z}}} P(x_{\mathcal{Z}}) = 1, \quad 0 \leq E[Y_{x_i} | x_{\mathcal{Z}}]P(x_{\mathcal{Z}}) \leq P(x_{\mathcal{Z}}) \leq 1 \quad (4)$$

By Def. 1, $P(x, y|z)$ could be written as linear combinations of $P(x_{\mathcal{Z}})$ and $E[Y_{x_i} | x_{\mathcal{Z}}]P(x_{\mathcal{Z}})$ as

$$P(x|z) = \sum_{x_{\mathcal{Z}}} I_{x_z=x} P(x_{\mathcal{Z}}), \quad (5)$$

$$E[Y|x, z]P(x|z) = \sum_{x_{\mathcal{Z}}} I_{x_z=x} E[Y_x | x_{\mathcal{Z}}]P(x_{\mathcal{Z}}). \quad (6)$$

where $I_{\{\cdot\}}$ is an indicator function. Similarly, the causal effects $E[Y_x]$ can be written as a linear function:

$$E[Y_x] = \sum_{x_{\mathcal{Z}}} E[Y_x | x_{\mathcal{Z}}]P(x_{\mathcal{Z}}).$$

Eq. (1) is reducible to linear programs (LP) optimizing the objective function $E[Y_x]$ subject to probabilistic constraints Eq. (4) and observational constraints Eq. (6), i.e.,

$$\begin{aligned} l_x &= \min E[Y_x] \\ h_x &= \max E[Y_x] \end{aligned} \quad \left| \quad \text{subject to Eqs. (4) to (6)} \right. \quad (7)$$

Solving such a linear program leads to a valid causal bound over the expected reward $E[Y_x]$.

Comparison with Existing Methods The idea of modeling the exogenous variables $\mathbf{U}$ and functional relationships in $\mathcal{F}$ using its projection to the latent counterfactuals $X_{\mathcal{Z}}, Y_{\mathcal{X}}$ has been explored in the literature, including the canonical partition (Balke and Pearl 1994b) and principal stratification (Frangakis and Rubin 2002). These methods can be seen as a RIV model without the independence restriction among $Y_{\mathcal{X}}$ (see Fig. 2b). For discrete X, Y, Z , representing $P(x_{\mathcal{Z}}, y_{\mathcal{X}})$ of this unrestricted model requires a probability table of size $\mathcal{O}(|\mathcal{X}|^{|\mathcal{Z}|} |\mathcal{Y}|^{|\mathcal{X}|})$; while a RIV model (Def. 1) requires a table of only $\mathcal{O}(|\mathcal{X}|^{|\mathcal{Z}|+1} |\mathcal{Y}|)$, removing the exponential dependence on $|\mathcal{X}|$. In addition, our representation does not require a particular parametric form of outcome (e.g., Y is discrete). This allows one to derive causal bounds on the continuous outcome by employing standard LP methods, which are provably optimal.

Contextual Settings

We now study the contextual IV (CIV) model shown in Fig. 3a where an additional context C is now observed. We denote by $do(\pi(x|c))$ a stochastic intervention where values of treatment are decided following a conditional distribution $\pi(x|c)$ mapping from C to X . The expected reward $E[Y_{\pi(x|c)}]$ induced by $do(\pi(x|c))$ is given by

$$E[Y_{\pi(x|c)}] = \sum_{x,c} E[Y_x|c] \pi(x|c) P(c). \quad (8)$$

We are interested in inferring the causal effect $E[Y_{\pi(x|c)}]$ from the observational distribution $P(x, y, c|z)$. The non-identifiability of this learning settings, shown in Fig. 3, has been acknowledged in the causal inference literature (Tian 2008; Correa and Bareinboim 2019).

We thus consider the partial identification problem that bounds $E[Y_{\pi(x|c)}]$ from $P(x, y, c|z)$ in CIV models. Among quantities in Eq. (8), $\pi(x|c)$, $P(c)$ are provided. It is thus sufficient to bound the conditional causal effect $E[Y_x|c]$. Let $\mathcal{M}_{CIV}[P(x, y, c|z)]$ denote contextual IV models that are compatible with the observational distribution $P(x, y, c|z)$. We show that the formulation of Def. 1 are also applicable in the contextual settings.

Theorem 2. *Given $P(x, y, c|z)$, fix a context $C = c$, for any CIV model $M_1 \in \mathcal{M}_{CIV}[P(x, y, c|z)]$, there exists an IV model $M_2 \in \mathcal{M}_{RIV}[P(x, y, c, z)]$ such that $E_{M_1}[Y_x|c] = E_{M_2}[Y_x]$, and vice versa.*

Thm. 2 implies that bounding $E[Y_x|c]$ from $P(x, y, c|z)$ in CIV models is equivalent to bounding $E[Y_x]$ from the observational data $P(x, y, c, z)$ in RIV models. For any $M \in \mathcal{M}_{CIV}[P(x, y, c|z)]$, one could always translate it into a solution of Eq. (1) within the feasible region $\mathcal{M}_{RIV}[P(x, y, c, z)]$. Let $E[Y_x|c] \in [l_x(c), h_x(c)]$ denote the solutions of Eq. (1) with $\mathcal{M}_{OBS} = \mathcal{M}_{RIV}[P(x, y, c, z)]$. Causal bounds $E[Y_{\pi(x|c)}] \in [l_\pi, h_\pi]$ are computable from $[l_x(c), h_x(c)]$ following Eq. (8).

Theorem 3. *Given $P(x, y, c|z)$, there exist CIV models $M_1, M_2 \in \mathcal{M}_{CIV}[P(x, y, c|z)]$ such that $E_{M_1}[Y_{\pi(x|c)}] = l_\pi$ and $E_{M_2}[Y_{\pi(x|c)}] = h_\pi$.*

Thm. 3 guarantees that $[l_\pi, h_\pi]$ are optimal in CIV models. Suppose there exists a bound $E[Y_{\pi(x|c)}] \in [l'_\pi, h'_\pi]$ strictly contained in $[l_\pi, h_\pi]$. We can always find CIV models M_1, M_2 compatible with $P(x, y, c|z)$ while their $E[Y_{\pi(x|c)}]$ lie outside $[l'_\pi, h'_\pi]$, which is a contradiction.

The conditional bound $E[Y_x|c] \in [l_x(c), h_x(c)]$ is obtainable by solving Eq. (7) where observational constraints $P(x|z), E[Y|x, z]P(x|z)$ in Eq. (6) are replaced with conditional quantities $P(x|c, z), E[Y|x, z, c]P(x|z, c)$. Approximating solution to Eq. (7) generalizes the natural bounds in (Manski 1990) to the contextual settings.

Theorem 4. *For an CIV model M , given $P(x, y, c|z)$, $\max_z l_\pi(z) \leq E[Y_{\pi(x|c)}] \leq \min_z h_\pi(z)$ where*

$$l_\pi(z) = \sum_{x,c} E[Y|x, c, z] \pi(x|c) P(x, c|z),$$

$$h_\pi(z) = l_\pi(z) + \sum_{x,c} \pi(X \neq x|c) P(x, c|z).$$




Figure 3: Causal diagram of the contextual instrumental variable (CIV) model: Z represents the (randomized) treatment assigned, X the treatment actually received, and Y the observed outcome; C is an additional observed context.

Estimation in High-Dimensional Context

Causal bounds developed so far are functions of the observational distribution, which are identifiable from the sampling process, and so generally can be consistently estimated. However, computational and sample complexity challenges could arise when context C is high-dimensional.

In this section, we will introduce robust estimation procedures to circumvent this issue. We first assume that the models $\hat{P}(x|c, z)$ and $\hat{E}[YI_{\{X=x\}}|c, z]$ of observational quantities $P(x|c, z)$ and $E[YI_{\{X=x\}}|c, z]$ are provided. Let $\{X_i, Y_i, C_i, Z_i\}_{i=1}^n$ denote finite samples drawn from an observational distribution $P(x, y, c, z)$. For a sampled instance with context C_i , we could compute functions $l_{X_i}(C_i), h_{X_i}(C_i)$ by solving LPs in Eq. (7) with observational constraints $P(x|z), E[Y|x, z]P(x|z)$ in Eq. (6) replaced with $\hat{P}(x|C_i, z)$ and $\hat{E}[YI_{\{X=x\}}|C_i, z]$. Since we consider only a fixed $C = C_i$, Eq. (7) could be solved efficiently despite of dimensionalities of context C . Thm. 2 ensures that the conditional causal effect $E[Y_x|C_i] \in [l_x(C_i), h_x(C_i)]$. We could obtain causal bounds $E[Y_{\pi(x|c)}] \in [l_\pi, h_\pi]$ by summing over finite samples. Formally, the empirical estimates $\hat{l}_\pi, \hat{h}_\pi$ of causal bounds l_π, h_π are given by:

$$\hat{l}_\pi = \frac{1}{n} \sum_{i=1}^n \sum_x \pi(x|C_i) l_x(C_i), \quad \hat{h}_\pi = \frac{1}{n} \sum_{i=1}^n \sum_x \pi(x|C_i) h_x(C_i).$$

Lemma 1. *If $\hat{P}(x|c, z)$ and $\hat{E}[YI_{\{X=x\}}|c, z]$ are models of $P(x|c, z)$ and $E[YI_{\{X=x\}}|c, z]$, $\hat{l}_\pi, \hat{h}_\pi$ are consistent estimators of l_π, h_π .*

When the models of $P(x|c, z)$ and $E[YI_{\{X=x\}}|c, z]$ are not available, one could obtain an efficient approximation of the causal bounds using the empirical estimates of the natural bounds in Thm. 4. For any $Z = z$, let $n(z) = \sum_{i=1}^n I_{Z_i=z}$. The estimators $\hat{l}_\pi(z), \hat{h}_\pi(z)$ are defined as follows:

$$\hat{l}_\pi(z) = \frac{1}{n(z)} \sum_{i=1}^n Y_i I_{Z_i=z} \pi(X_i|C_i),$$

$$\hat{h}_\pi(z) = \hat{l}_\pi(z) + \frac{1}{n(z)} \sum_{i=1}^n I_{Z_i=z} \pi(X \neq X_i|C_i).$$

Lemma 2. *$\hat{l}_\pi(z), \hat{h}_\pi(z)$ are consistent estimators of functions $l_\pi(z), h_\pi(z)$ defined in Thm. 4.*

The causal effect $E[Y_{\pi(x|c)}]$ could then be bounded from the finite samples $\{X_i, Y_i, C_i, Z_i\}_{i=1}^n$ by inequalities $\max_z \hat{l}_\pi(z) \leq E[Y_{\pi(x|c)}] \leq \min_z \hat{h}_\pi(z)$.

Bandit Algorithms with Causal Bounds

The causal bounds derived so far may seem to be uninformative since they do not immediately identify the optimal treatment in IV models. We will show that this is not the case. More specifically, we will introduce a systematic procedure to incorporate causal bounds in online bandit algorithms (Auer, Cesa-Bianchi, and Fischer 2002; Sen, Shanmugam, and Shakkottai 2018; Audibert and Bubeck 2010) for identifying the optimal treatment. Our analysis reveals that causal bounds consistently improve the performance of bandit algorithms in various learning settings.

For the IV model of Fig. 1(a), we denote by μ_x the expected reward $E[Y_x]$ of performing a treatment $do(x)$. Let $\mu^* = \max_x \mu_x$ and let x^* denote the optimal treatment (so, $\mu_{x^*} = \mu^*$). A bandit algorithm learns the optimal treatment through repeated episodes of experiments $t = 1, \dots, T$. At each episode t , the algorithm performs an intervention $do(X_t)$ and observes an outcome Y_t . The cumulative regret $E[R_T]$ after T episodes is defined by $E[R_T] = T\mu^* - \sum_{t=1}^T E[Y_t]$, i.e., the loss due to the fact that the algorithm does not always play the optimal arm. A desirable asymptotic property of an algorithm is to have $\lim_{T \rightarrow \infty} E[R_T]/T = 0$, meaning that the procedure converges and finds the optimal treatment x^* .

We also consider the contextual IV model of Fig. 3(a). Let Π be a finite set of candidate policies $\{\pi_1, \dots, \pi_N\}$. Let $\mu_\pi = E[Y_{\pi(x|c)}]$ and let π^* denote the optimal policy $\pi^* = \arg \max_{\pi \in \Pi} \mu_\pi$. At each episode $t = 1, \dots, T$, a bandit algorithm has access to a context C_t , picks a policy π_t , assigns a treatment X_t following $\pi_t(x|C_t)$ and observes a reward Y_t . Observational data $P(x, y, c|z)$ are provided prior to the experiments. Similarly, the cumulative regret $E[R_T^\Pi]$ after T episodes is defined as $E[R_T^\Pi] = T\mu_{\pi^*} - \sum_{t=1}^T E[Y_t]$. We will assess and compare the performance of bandit algorithms in terms of the cumulative regrets.

Causal UCB

Our methods follow the well celebrated principle of *optimism in the face of uncertainty* (OFU). This principle leads to efficient bandit strategies, in the form of the upper confidence bound (UCB) algorithms (Auer, Cesa-Bianchi, and Fischer 2002). We will next describe UCB strategy in bandit models with a set of candidate policies Π (an arm choice x can be seen as a policy in $\mathcal{X}$). At each round t , a UCB agent constructs a set of “plausible” models $\mathcal{M}_t$ that are consistent with the data. The agent then identifies the most “favorable” model from $\mathcal{M}_t$ and prescribes the optimal policy for the identified model. In practice, this decision is determined by the concentration bounds over possible models. An upper (lower) confidence bound $U_\pi(t)$ ($L_\pi(t)$) for a policy π at time t is the maximum (minimum) expected reward of π of models in $\mathcal{M}_t$. The agent then prescribes a policy based on $U_x(t)$, $L_x(t)$ and observes a reward.

We now incorporate causal bounds using the OFU principle, called the Causal-UCB (for short, UCB^c). Let μ_{M_π} denote the expected reward of a policy $\pi \in \Pi$ in model M . At each round t , we obtain a subset $\mathcal{M}_t^c$ from $\mathcal{M}_t$ by removing models inconsistent with the causal bounds, i.e.,

Algorithm 1: Causal-UCB (UCB^c)

- 1: **Input:** Causal bounds $\{[l_\pi, h_\pi]\}_{\pi \in \Pi}$.
- 2: **for all** t **do**
- 3: For each $\pi \in \Pi$, compute $\bar{U}_\pi(t)$, $\bar{L}_\pi(t)$ as:

$$\begin{aligned}\bar{U}_\pi(t) &= \max \{ \min \{ U_\pi(t), h_\pi \}, l_\pi \}, \\ \bar{L}_\pi(t) &= \max \{ \min \{ L_\pi(t), h_\pi \}, l_\pi \},\end{aligned}\tag{9}$$

where $U_\pi(t)$, $L_\pi(t)$ are, respectively, the upper and lower confidence bounds for policy π .

- 4: Play an arm $do(X_t \sim \pi_t)$ and observe Y_t .
 - 5: **end for**
-

$\mathcal{M}_t^c = \{M \in \mathcal{M}_t : \mu_{M_\pi} \in [l_\pi, h_\pi]\}$. The agent then prescribes a decision that is optimal to the most favorable model in the subset $\mathcal{M}_t^c$. When the causal bounds are beneficial and significantly reduce the complexities of the search space $\mathcal{M}_t$, it is expected that a UCB^c agent outperforms the standard (non-causal) method. Alg. 1 describes an implementation of UCB^c in bandit settings. At trial t , we clip confidence bounds $U_\pi(t)$, $L_\pi(t)$ using the causal bound $[l_\pi, h_\pi]$. The agent then proceeds ordinarily with the clipped bounds $\bar{U}_\pi(t)$, $\bar{L}_\pi(t)$. Likewise, UCB^c in IV models of Fig. 1a follows Alg. 1 with π replaced with arm $x \in \mathcal{X}$.

For the remainder of this section, we will apply the UCB^c strategy (Alg. 1) to the state-of-the-art bandit algorithms, showing its consistent improvements for various tasks. We refer the interested readers to the technical report (Zhang and Bareinboim 2020) for detailed implementations.

Multi-Armed Bandits We start with IV models of Fig. 1a. (Cappé et al. 2013) proposed the kl-UCB procedure that is applicable for bandit settings with bounded reward. It computes confidence bound $U_x(t)$ for each arm x using Cramer’s theorem. The agent then plays an arm X_t with the largest $U_x(t)$. Let $\Delta_x = \mu^* - \mu_x$ and let $\mathcal{X}^-$ be the subset $\{x \in \mathcal{X} : \mu_x < \mu^*\}$. kl-UCB guarantees a bound on the cumulative regret of:

$$E[R_T] \leq \sum_{x \in \mathcal{X}^-} \left(\frac{\Delta_x}{kl(\mu_x, \mu^*)} \right) \log(T) + o(\log(T)), \tag{10}$$

where $kl(\mu_x, \mu^*) = \mu_x \log(\mu_x/\mu^*) + (1 - \mu_x) \log((1 - \mu_x)/(1 - \mu^*))$, i.e., the Kullback-Leibler divergence between Bernoulli distributions with mean μ_x and μ^* .

We applied Causal UCB strategy (Alg. 1) to kl-UCB by replacing policy π with arm $x \in \mathcal{X}$. We denote the resultant algorithm kl-UCB^c. At trial t , a kl-UCB^c agent pulls an arm X_t with the largest clipped confidence bound $\bar{U}_x(t)$ where $\bar{U}_x(t)$ is obtained following Eq. (9). Let $\mathcal{X}_{h_x \geq c}^-$ denote a set $\{x \in \mathcal{X}^- : h_x \geq c\}$. kl-UCB^c guarantees the asymptotic regret bound as follows:

Theorem 5. *Given $E[Y_x] \in [l_x, h_x]$, the regret $E[R_T]$ of kl-UCB^c after $T \geq 3$ is bounded by*

$$E[R_T] \leq \sum_{x \in \mathcal{X}_{h_x \geq c}^-} \left(\frac{\Delta_x}{kl(\mu_x, \mu^*)} \right) \log(T) + o(\log(T)).$$

SCMs	IV models (Fig. 1)		Contextual IV (Fig. 3)
Tasks	Cumulative Regret $E[R_T]$	Best-Arm	Cumulative Regret $E[R_T^{\Pi}]$
Standard	$\mathcal{O}\left(\sum_{x \in \mathcal{X}^-} \left(\frac{\Delta_x}{kl(\mu_x, \mu^*)}\right) \log(T)\right)$	$\mathcal{O}\left(\sum_{x \in \mathcal{X}^-} \Delta_x^{-2} \log(\delta^{-1} K \log(\Delta_x^{-2}))\right)$	$\mathcal{O}(C\lambda(\Pi^-)M^2 \log(T))$
Causal	$\mathcal{O}\left(\sum_{x \in \mathcal{X}^-_{h_x \geq \mu^*}} \left(\frac{\Delta_x}{kl(\mu_x, \mu^*)}\right) \log(T)\right)$	$\mathcal{O}\left(\sum_{x \in \mathcal{X}^-_{h_x \geq \mu_{1,2}}} \Delta_x^{-2} \log(\delta^{-1} K \log(\Delta_x^{-2}))\right)$	$\mathcal{O}(C\lambda(\Pi^-_{h_\pi \geq \mu_{\pi^*}})M^2 \log(T))$

Table 1: Summary of bandit results presented in this paper. “SCMs” represents the causal diagrams of the corresponding learning settings. “Obj.” stands for the objective that the target agent aims to optimize. “Standard” stands for the asymptotics of the standard UCB-style algorithms, “Causal” for the proposed strategy leveraging the observational distribution.

The regret bound in Thm. 5 is guaranteed to be smaller than Eq. (10) if $\mathcal{X}^-_{h_x \geq \mu^*}$ is strictly contained in $\mathcal{X}^-$. This result coincides with the optimal regret of B-kl-UCB (Zhang and Bareinboim 2017) when bounds over the expected reward are provided. Since $\Delta_x/kl(\mu_x, \mu^*) \leq 1/(2\Delta_x)$, the improvement of kl-UCB^c is significant when the gap Δ_x of $x \in \mathcal{X}^-_{h_x < \mu^*}$ is small and close to zero.

Best Arm Identification We also consider the settings of pure exploration in IV models (Mannor and Tsitsiklis 2004). Rather than looking at the cumulative regret, we are concerned with PAC-style (“probably approximately correct”) bound on the sample complexity to identify the optimal treatment. In (Jamieson and Nowak 2014), a sampling strategy, called lil’LUCB, were provided, which finds the optimal arm with probability at least $1 - \frac{2+\epsilon}{\epsilon/2}(\log(1+\epsilon))^{-(1+\epsilon)}\delta$ in at most $\mathcal{O}(\sum_{x \in \mathcal{X}^-} \Delta_x^{-2} \log(\delta^{-1} |\mathcal{X}| \log(\Delta_x^{-2})))$ trials for any $\epsilon \in (0, 1)$ and $\delta \in (0, \log(1+\epsilon)/e)$.

We apply Alg. 1 to lil’LUCB and denote the resulting algorithm lil’LUCB^c. Assume (without loss of generality) that arms are ordered such that $\mu_1 > \mu_2 \geq \dots \geq \mu_N$ (so, $\mu^* = \mu_1$). Let $\mu_{1,2} = (\mu_1 + \mu_2)/2$. The following theorem provides the sample complexity analysis of lil’LUCB^c.

Theorem 6. *Given $E[Y_x] \in [l_x, h_x]$, with probability (w.p.) at least $1 - \frac{2+\epsilon}{\epsilon/2}(\log(1+\epsilon))^{-(1+\epsilon)}\delta$, lil’LUCB^c returns the optimal treatment x^* with $\mathcal{O}\left(\sum_{x \in \mathcal{X}^-_{h_x \geq \mu_{1,2}}} \Delta_x^{-2} \log(\delta^{-1} |\mathcal{X}| \log(\Delta_x^{-2}))\right)$ samples.*

Compared with lil’LUCB, the sample complexity bound in Thm. 6 is tighter if $\mathcal{X}^-_{\mu_{1,2}} \subset \mathcal{X}^-$.

Contextual Bandits Finally, we consider the regret minimization in the contextual IV models of Fig. 3a. (Sen, Shanmugam, and Shakkottai 2018) proposed a UCB-style algorithm for contextual bandits, called D-UCB, when a set of stochastic policies $\pi \in \Pi$ are provided, i.e., $\pi(x|z) > 0$. At each round t , D-UCB estimates $U_\pi(t)$ for each policy π using the concentration bounds of the clipped importance sampling estimator (Sen, Shanmugam, and Shakkottai 2018), and apply the policy with the highest $U_\pi(t)$ estimation. Let $M(\pi_i, \pi_j)$ denote the log divergence between two arbitrary policies π_i, π_j (Sen, Shanmugam, and Shakkottai 2018, Def. 2), and let $M = \max_{\pi_i, \pi_j \in \Pi} M(\pi_i, \pi_j)$. For any $\Pi' \subseteq \Pi$, let policies in Π' be ordered such that $\mu_{\pi_1} \geq \mu_{\pi_2} \geq \dots \geq \mu_{\pi_n}$ where $n = |\Pi'|$. For any $\pi \in \Pi$, let $\Delta_\pi = \mu_{\pi^*} - \mu_\pi$. We define function $\lambda(\Pi') = \Delta_{\pi_n} \gamma(\Delta_{\pi_n}) + \sum_{i=1}^{n-1} \Delta_{\pi_i} (\gamma(\Delta_{\pi_i}) -$

$\gamma(\Delta_{\pi_{i+1}}))$, where $\gamma(\Delta) = \log^2(6/\Delta)/\Delta^2$. Let Π^- denote the set of sub-optimal policies $\{\pi \in \Pi : \mu_\pi < \mu_{\pi^*}\}$. D-UCB obtains an asymptotic regret bound as follow:

$$E[R_T^{\Pi}] \leq C\lambda(\Pi^-)M^2 \log(T) + o(\log(T)), \quad (11)$$

where C is a constant. We apply the causal strategy (Alg. 1) to D-UCB and denote the new procedure D-UCB^c. Let $\Pi^-_{h_x \geq c}$ be a set of sub-optimal policies $\{\pi \in \Pi^- : h_\pi \geq c\}$. We now provides the regret bound of D-UCB^c.

Theorem 7. *Given $E[Y_{\pi(x|c)}] \in [l_\pi, h_\pi]$, the regret $E[R_T^{\Pi}]$ of D-UCB^c after $T \geq 2$ is bounded by*

$$E[R_T^{\Pi}] \leq C\lambda(\Pi^-_{h_x \geq \mu_{\pi^*}})M^2 \log(T) + o(\log(T)). \quad (12)$$

Compared with Eq. (11), the bound in Thm. 7 only differs in the hardness measure $\lambda(\Pi^-_{\mu_{\pi^*}})$. The following lemma guarantees that $\lambda(\Pi^-_{\mu_{\pi^*}})$ is never larger than $\lambda(\Pi^-)$.

Lemma 3. *For any $\Pi_1 \subset \Pi_2 \subseteq \Pi$, $\lambda(\Pi_1) \leq \lambda(\Pi_2)$. If $\Delta_{\pi_1} < \dots < \Delta_{\pi_{|\Pi|}}$, $\lambda(\Pi_1) < \lambda(\Pi_2)$.*

Thm. 7, together with Lem. 3, says that D-UCB^c improves over D-UCB if there exists some sub-optimal π with $h_\pi < \mu_{\pi^*}$ (given that μ_π of each $\pi \in \Pi^-$ are not all equal).

We summarize in Table. 1 the results discussed in this section. “Standard” stands for the asymptotics of the standard algorithms and “Causal” is our strategy leveraging the causal bounds obtained from the observation. The interesting aspect is that the causal approach is guaranteed to rival the standard algorithms, even when the observation is biased towards the wrong decision. If the causal bounds are beneficial (e.g., $h_x < \mu^*$), our approach could eliminate a sub-optimal arm x (or policy π) early during the trials, thus outperforming the standard methods. Such improvements could be significant in more difficult instances, when the gap Δ_x between x and x^* is small, e.g., close to zero.

Experiments: International Stroke Trials

We now use a real-world dataset to investigate the performance of proposed bandit strategies. Specifically, we study the International Stroke Trial (IST) (Carolei et al. 1997), focusing on the effect of the aspirin allocation X on a composite score Y , a continuous value in $[0, 1]$ predicting the likelihood of patients’ recovery. We also measure the pre-treatment attributes U , including age, gender, and conscious state. To emulate unobserved confounding, as discussed in the paper, we filter the IST data following a inclusion rule $f(x|z, u)$ and hide some columns of U . We repeat this procedure and generate observational samples using 4 different




Figure 4: Simulations comparing solvers that are causally enhanced (kl-UCB^c, lil'LUCB^c, D-UCB^c) and standard (kl-UCB, lil'LUCB and D-UCB) on the International Stroke Trials data. Graphs are rendered in high resolution and can be zoomed in.

inclusion rules $\{f(x|z_i, \mathbf{u})\}_{i=1, \dots, 4}$. To evaluate the performance of the different bandit strategies, we make bootstrap estimates of the patient's true response from the IST data. For details on the experimental setups, we refer readers to the full technical report (Zhang and Bareinboim 2020).

Multi-Armed Bandits The expected reward of not giving ($X = 0$) and giving aspirin ($X = 1$) are respectively $\mu_0 = 0.6201$ and $\mu_1 = 0.6948$, suggesting an increased chance of recovery from aspirin. The causal bounds $\mu_0 \in [0.5905, 0.6506]$, $\mu_1 \in [0.4839, 0.7527]$ estimated using all $n = 9650$ samples do not permit the identification of the optimal treatment, since one is contained in the other. We deploy a kl-UCB^c agent provided with these causal bounds. For comparison, we also include a standard kl-UCB agent, a kl-UCB^c agent with the causal bounds estimated using 300 samples (kl-UCB^c-), and kl-UCB warm-started with the empirical estimates of $E[Y|x, z]$ (kl-UCB⁻). The results (Fig. 4a) reveal a significant difference in the cumulative regret (CR) between kl-UCB (CR = 38.63) and kl-UCB^c (CR = 0.07); kl-UCB^c- agent (CR = 37.94) coincides with kl-UCB since the sample size $n = 300$ are not sufficient for obtaining any informative estimate of the causal bounds (Thm. 5); kl-UCB⁻ (CR = 373.42) performs worst among all strategies due to unobserved confounding.

Best Arm ID We also run the lil'LUCB^c agent with empirical bounds estimated from $n = 9711$ observations (lil'LUCB^c-all) to efficiently identify the optimal treatment ($X = 1$). For comparison, we include the standard lil'LUCB, lil'LUCB^c with empirical bounds obtained from $n = 100$ observations (lil'LUCB^c-100), and lil'LUCB warm-started with estimates of $E[Y|x]$ from observations (lil'LUCB^o). The stopping times T of each algorithm are compared in Fig. 4(b). We can immediately note a dramatic difference in the sample complexities experienced by lil'LUCB^c ($T = 4.4 \times 10^3$) compared to lil'LUCB ($T = 6.5 \times 10^3$) and lil'LUCB^c-100 ($T = 6.5 \times 10^3$). lil'LUCB^o ($T = 8.2 \times 10^3$) performs worst among all algorithms.

Contextual Bandits Suppose we now have access to a context $C = \{\text{sex}\}$. Our goal is to find the optimal treat-

ment among two policies $\pi_0(x|c)$ and $\pi_1(x|c)$. We also include experiments for more involved candidate policies in the technical report (Zhang and Bareinboim 2020). We estimate causal bounds over μ_π using $n = 9650$ observational samples and provide them to a D-UCB^c agent. For comparison, we include the standard D-UCB, D-UCB^c with causal bounds from $n = 300$ samples (D-UCB^c-), and D-UCB seeded with samples of confounded observations (D-UCB⁻). The cumulative regrets (CR) of each strategy are measured and compared in Fig. 4c. The analysis reveals a significant difference in CR experienced by D-UCB^c (CR = 0.07) compared to D-UCB (CR = 111.8) and D-UCB^c- (CR = 110.8). Unsurprisingly, D-UCB⁻ (CR = 153.9) performs worst among all strategies.

These results corroborate with our findings: useful information could be extracted from the confounded, passively-collected data to improve the performance of a learning agent. The causal approaches (e.g, kl-UCB^c) dominate the standard, non-causal methods (kl-UCB) given sufficient observational samples. When the number of observations is not statistically significant, our approaches could still rival the standard methods, i.e., no negative transfer occurs.

Conclusions

In this paper, we investigated the problem of bounding causal effects from experimental studies in which treatment assignment is randomized but the subject compliance is imperfect. Under such conditions, the actual causal effects are not identifiable due to uncontrollable confounding. In particular, we derived informative bounds over the causal effect and accounted for challenging issues due to high-dimensional context and the lack of discreteness. We incorporated these bounds into UCB-like algorithms and proved that the causal approach, leveraging observational data, consistently dominates non-causal, state-of-the-art procedures. We hope that our framework can be useful in practical settings since, even though imperfect, observational data contains causal information and is abundantly available.

Acknowledge

The authors were partially supported by grants from NSF IIS-1704352 and IIS-1750807 (CAREER).

References

- Angrist, J.; Imbens, G.; and Rubin, D. 1996. Identification of causal effects using instrumental variables (with Comments). *Journal of the American Statistical Association* 91(434): 444–472.
- Audibert, J.-Y.; and Bubeck, S. 2010. Best arm identification in multi-armed bandits. In *COLT-23th Conference on Learning Theory-2010*, 13–p.
- Auer, P.; Cesa-Bianchi, N.; and Fischer, P. 2002. Finite-time analysis of the multiarmed bandit problem. *Machine learning* 47(2-3): 235–256.
- Balke, A.; and Pearl, J. 1994a. Counterfactual Probabilities: Computational Methods, Bounds, and Applications. In de Mantaras, R. L.; and Poole, D., eds., *Uncertainty in Artificial Intelligence 10*, 46–54. San Mateo, CA: Morgan Kaufmann.
- Balke, A.; and Pearl, J. 1994b. Counterfactual Probabilities: Computational Methods, Bounds and Applications. In *Proceedings of the Tenth International Conference on Uncertainty in Artificial Intelligence*, UAI’94, 46–54.
- Balke, A.; and Pearl, J. 1995. Counterfactuals and policy analysis in structural models. In *Proceedings of the Eleventh Conference on Uncertainty in Artificial Intelligence*, 11–18.
- Balke, A.; and Pearl, J. 1997. Bounds on treatment effects from studies with imperfect compliance. *Journal of the American Statistical Association* 92(439): 1172–1176.
- Bang, H.; and Robins, J. M. 2005. Doubly robust estimation in missing data and causal inference models. *Biometrics* 61(4): 962–973.
- Bareinboim, E.; and Pearl, J. 2012. Causal inference by surrogate experiments: z -identifiability. In de Freitas, N.; and Murphy, K., eds., *Proceedings of the Twenty-Eighth Conference on Uncertainty in Artificial Intelligence*, 113–120. Corvallis, OR: AUAI Press.
- Bareinboim, E.; and Pearl, J. 2016. Causal inference and the data-fusion problem. *Proceedings of the National Academy of Sciences* 113: 7345–7352.
- Cappé, O.; Garivier, A.; Maillard, O.-A.; Munos, R.; Stoltz, G.; et al. 2013. Kullback–leibler upper confidence bounds for optimal sequential allocation. *The Annals of Statistics* 41(3): 1516–1541.
- Carolei, A.; et al. 1997. The International Stroke Trial (IST): a randomized trial of aspirin, subcutaneous heparin, both, or neither among 19435 patients with acute ischaemic stroke. *The Lancet* 349: 1569–1581.
- Chickering, D.; and Pearl, J. 1996. A clinician’s apprentice for analyzing non-compliance. In *Proceedings of the Twelfth National Conference on Artificial Intelligence*, volume Volume II, 1269–1276. Menlo Park, CA: MIT Press.
- Cinelli, C.; Kumor, D.; Chen, B.; Pearl, J.; and Bareinboim, E. 2019. Sensitivity Analysis of Linear Structural Causal Models. In Chaudhuri, K.; and Salakhutdinov, R., eds., *Proceedings of the 36th International Conference on Machine Learning*, volume 97, 1252–1261. Long Beach, CA: PMLR.
- Correa, J.; and Bareinboim, E. 2019. From Statistical Transportability to Estimating the Effect of Stochastic Interventions. In Kraus, S., ed., *Proceedings of the 28th International Joint Conference on Artificial Intelligence*, 1661–1667. Macao, China: International Joint Conferences on Artificial Intelligence Organization.
- Dudík, M.; Langford, J.; and Li, L. 2011. Doubly robust policy evaluation and learning. In *Proceedings of the 28th International Conference on International Conference on Machine Learning*, 1097–1104. Omnipress.
- Fisher, R. 1935. *The Design of Experiments*. Edinburgh: Oliver and Boyd.
- Frangakis, C.; and Rubin, D. 2002. Principal Stratification in Causal Inference. *Biometrics* 1(58): 21–29.
- Gittins, J. C. 1979. Bandit processes and dynamic allocation indices. *Journal of the Royal Statistical Society. Series B (Methodological)* 148–177.
- Jamieson, K.; and Nowak, R. 2014. Best-arm identification algorithms for multi-armed bandits in the fixed confidence setting. In *Information Sciences and Systems (CISS)*, 1–6. IEEE.
- Kallus, N.; Puli, A. M.; and Shalit, U. 2018. Removing Hidden Confounding by Experimental Grounding. In Bengio, S.; Wallach, H.; Larochelle, H.; Grauman, K.; Cesa-Bianchi, N.; and Garnett, R., eds., *Advances in Neural Information Processing Systems 31*, 10911–10920. Curran Associates, Inc.
- Kallus, N.; and Zhou, A. 2018. Confounding-Robust Policy Improvement. In Bengio, S.; Wallach, H.; Larochelle, H.; Grauman, K.; Cesa-Bianchi, N.; and Garnett, R., eds., *Advances in Neural Information Processing Systems 31*, 9289–9299. Curran Associates, Inc.
- Li, L.; Munos, R.; and Szepesvari, C. 2015. Toward Minimax Off-policy Value Estimation. In *Proceedings of the 18th International Conference on Artificial Intelligence and Statistics*.
- Mannor, S.; and Tsitsiklis, J. N. 2004. The sample complexity of exploration in the multi-armed bandit problem. *Journal of Machine Learning Research* 5(Jun): 623–648.
- Manski, C. 1990. Nonparametric bounds on treatment effects. *American Economic Review, Papers and Proceedings* 80: 319–323.
- Manski, C. F. 2003. *Partial identification of probability distributions*. Springer Science & Business Media.
- Munos, R.; Stepleton, T.; Harutyunyan, A.; and Bellemare, M. 2016. Safe and efficient off-policy reinforcement learning. In *Advances in Neural Information Processing Systems*.
- Neyman, J. 1923. On the application of probability theory to agricultural experiments. Essay on principles. Section 9. *Statistical Science* 5(4): 465–480.
- Pearl, J. 2000. *Causality: Models, Reasoning, and Inference*. New York: Cambridge University Press. 2nd edition, 2009.
- Richardson, A.; Hudgens, M. G.; Gilbert, P. B.; and Fine, J. P. 2014. Nonparametric bounds and sensitivity analysis of

treatment effects. *Statistical science: a review journal of the Institute of Mathematical Statistics* 29(4): 596.

Robins, J. 1989. The analysis of randomized and non-randomized AIDS treatment trials using a new approach to causal inference in longitudinal studies. In Sechrest, L.; Freeman, H.; and Mulley, A., eds., *Health Service Research Methodology: A Focus on AIDS*, 113–159. Washington, D.C.: NCHSR, U.S. Public Health Service.

Rosenbaum, P.; and Rubin, D. 1983. The central role of propensity score in observational studies for causal effects. *Biometrika* 70: 41–55.

Sen, R.; Shanmugam, K.; and Shakkottai, S. 2018. Contextual Bandits with Stochastic Experts. In *International Conference on Artificial Intelligence and Statistics*, 852–861.

Spirtes, P.; Glymour, C. N.; and Scheines, R. 2000. *Causation, prediction, and search*. MIT press.

Thomas, P.; and Brunskill, E. 2016. Data-efficient off-policy policy evaluation for reinforcement learning. In *International Conference on Machine Learning*, 2139–2148.

Tian, J. 2008. Identifying dynamic sequential plans. In *Proceedings of the Twenty-Fourth Conference on Uncertainty in Artificial Intelligence*.

Wright, P. 1928. *The Tariff on Animal and Vegetable Oils*. New York, NY: The MacMillan Company.

Zhang, J.; and Bareinboim, E. 2017. Transfer learning in multi-armed bandits: a causal approach. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence*, 1340–1346. AAAI Press.

Zhang, J.; and Bareinboim, E. 2020. Bounding Causal Effects on Continuous Outcome. Technical report. URL <https://causalai.net/r61.pdf>.