

# Counterfactual Prediction Under Selective Confounding

Sohaib Kiani<sup>a,b,\*</sup>, Jared Barton<sup>c</sup>, Jon Sushinsky<sup>c</sup>, Lynda Heimbach<sup>c</sup> and Bo Luo<sup>b</sup>

<sup>a</sup>Department of Mathematics and Computer Science, Beloit College, USA

<sup>b</sup>Department of Electrical Engineering and Computer Science, University of Kansas, USA

<sup>c</sup>School of Social Welfare, University of Kansas, USA

**Abstract.** This research addresses the challenge of conducting interpretable causal inference between a binary treatment and its resulting outcome when not all confounders are known. Confounders are factors that have an influence on both the treatment and the outcome. We relax the requirement of knowing all confounders under desired treatment, which we refer to as *Selective Confounding*, to enable causal inference in diverse real-world scenarios. Our proposed scheme is designed to work in situations where multiple decision-makers with different policies are involved and where there is a re-evaluation mechanism after the initial decision to ensure consistency. These assumptions are more practical to fulfill compared to the availability of all confounders under all treatments. To tackle the issue of *Selective Confounding*, we propose the use of *dual-treatment* samples. These samples allow us to employ two-step procedures, such as Regression Adjustment or Doubly-Robust, to learn counterfactual predictors. We provide both theoretical error bounds and empirical evidence of the effectiveness of our proposed scheme using synthetic and real-world child placement data. Furthermore, we introduce three evaluation methods specifically tailored to assess the performance in child placement scenarios. By emphasizing transparency and interpretability, our approach aims to provide decision-makers with a valuable tool. The source code repository of this work is located at <https://github.com/sohaib730/CausalML>.

## 1 Introduction

Accurately predicting treatment outcomes is crucial for decision-makers in various domains, including health care [3], child welfare [11], criminal justice [5], and marketing [4]. Predictive algorithms that can extract valuable insights from rich case information have the potential to enable better decision-making and improve future outcomes. Counterfactual prediction algorithms have been widely adopted in healthcare [39] and are now increasingly used in other domains, such as child welfare [10] and education [47, 48]. For example, in child welfare, practitioners often face challenges regarding whether to remove children from their homes and place them in foster care or keep them with their families [2]. Errors in judgment can occur on either end of this decision spectrum. Children who remain at home may experience additional maltreatment, while children erroneously removed from their homes may eventually return after a short stay in out-of-home placement, causing unnecessary trauma for the child and family [36]. A counterfactual prediction algorithm can address both issues by answering the question: What would be the risk of maltreatment if a child is kept in their home?

The problem of counterfactual prediction differs from the standard supervised learning problem [30, 42]. Confounders are the factors that influence both treatment assignment and outcome [45]. Standard predictive algorithms may either underestimate or overestimate the risk when trained over whole populations [11] or exhibit selective bias when trained only on the population that received the desired treatment. This scenario can be likened to a missing label problem, where outcomes for all treatments are not observed for each unit. Therefore, for applications involving interventional decisions, counterfactual learning is essential to avoid biases that can otherwise lead to erroneous training, performance, and fairness evaluations [11].

While randomized control trials (RCTs) are viewed as the ideal approach for establishing causal relationships, they are often impractical for numerous applications due to their cost, ethical considerations, or the unavailability of sufficient data [32]. Conversely, inferring causality from observed data necessitates the observation of all confounding variables [12]. This stringent requirement considerably restricts the practicality of these models in real-world scenarios where not all confounding factors are reported or obtainable [19, 10].

In various decision-making scenarios, such as child placement, ICU admission, or police arrests, limited information is available under the baseline treatment with no intervention. For example, in the field of child welfare, decision-making factors include demographic data, case records, and sensitive information such as maltreatment reports, mental health, and disabilities. It is common for the information reported for children who remain with their families to be less comprehensive compared to those who are placed out of the home [8, 9]. This poses a challenge of hidden confounders under the desired treatment, which we refer to as *Selective Confounding*.

Learning causal relationships from observed data is already a complex task, and it becomes even more challenging when confounding factors are hidden [17]. Existing techniques in the literature address this issue in specific settings, such as when hidden confounders are proxies of observed confounders [26], when there is limited randomized control trial (RCT) data available [19], or when confounders are hidden only during prediction [10]. However, to the best of our knowledge, no method has been proposed that provides an interpretable counterfactual model under the scenario of *Selective Confounding*.

In this work, we address the challenge of *Selective Confounding* using interpretable techniques, relying on two key assumptions: 1) there exists a variation in decision policies during the initial decision-making process. 2) the recidivism of the initial decision in appeal or reevaluation is attributed to errors or biases in the initial judgment, rather than being a consequence of covariate shift. These conditions are applicable in real-world scenarios involving child placement, ICU admission, and police arrests, where multiple decision-

\* Corresponding Author. Email: [kianis@beloit.edu](mailto:kianis@beloit.edu)

makers are involved, and opportunities for appeal or reevaluation exist. For instance, in child placement, we often observe varying decision policies, with initial screenings conducted by police officers or trained child welfare professionals [37, 28]. Then there exist short-stayer children who are reunited with their families soon after being removed due to erroneous out-of-home placements [13]. Similarly, in cases of short stays in the ICU, where patients are released after two days of monitoring, different hospitals or attending physicians correspond to varying decision policies, and the release without any intervention indicates that the stay could have been avoided. Likewise, in instances where individuals are arrested but subsequently released because no charges are filed against them by the attorney, different police officers and locations reflect varying decision policies, and the release without charges indicates the arrest could have been avoided.

This paper proposes a specialized approach that leverages the structure of *selective confounding* by using dual-treatment samples, such as *short stayers* in child placement applications. These samples are unconfounded under *selective confounding* and can be adjusted to obtain the desired treatment outcome, allowing us to identify the target estimator without any confounding bias [17, 10]. We adapt well-known two-stage counterfactual learning algorithms, Regression Adjustment (RA) and Doubly Robust (DR) [45], to train the target model. In this regard, we have opted for Linear regression as our target model, owing to its intrinsic interpretability, making it highly suitable for sensitive social applications.

We demonstrate the effectiveness of our approach through empirical evaluations on both synthetic and real-world data. Since evaluation over real-world data is not straightforward due to missing ground truth labels over subpopulations, we propose novel evaluation methods specifically tailored to child placement applications. Our work represents a significant contribution to the field of causal inference and has the potential to enhance decision-making in sensitive domains, including child welfare and ICU admissions, among others. The complete code for training and reproducing the results over synthetic data is accessible at <https://github.com/sohaib730/CausalML>.

## 2 Setup

**Selective Confounding:** Our goal is to predict the outcome under a desired treatment  $a_1$  for a binary treatment application where possible treatments are  $\{a_1, a_2\}$ . Here  $a_1$  is the desired treatment that involves no intervention like a child kept in-home or no ICU admission. Let  $X, Z$  be the confounding factors. We have considered the problem setting of *selective confounding* where  $Z$  represents hidden confounders that are not available under desired treatment  $T = a_1$  and  $X$  are observable under all treatments. Our aim is to relax the condition of observable confounding features to accommodate more applications. According to the potential outcome framework [17], the target estimator is  $\nu_{a_1}(x) := E(Y_{a_1}|X = x)$ , where  $Y_{a_1} \in R$  is the potential outcome one would observe under treatment  $T = a_1$ .

**Dual-Treatment Samples:** In an observational study of a binary treatment application, few samples experience treatment reversal, e.g., a child reunited with their family after removal or a patient released from ICU without any intervention. These samples were previously ignored or considered only under one treatment to ensure training samples are independent. In our work, we still maintain that samples are independent by either including them under one treatment or identifying them with another treatment  $T = a_3$  - namely, dual-treatment. We argue that a change in treatment usually occurs due to two reasons: 1) Covariate Shift, when things im-

prove or deteriorate over time, or 2) Error in the Initial judgment. The dual-treatment samples are those where decisions have been reverted from  $a_2$  to the desired treatment  $a_1$  due to a mistake in the initial judgment. For example, in the child placement application, *Short-stayers* are considered dual-treatment samples. Short-stayers are the children who are reunited with their families within 30-days after out-of-home placement [13]. The 30-day separation window is quite small for a change in covariates, and the main reason behind reunification is errors in earlier judgment [36]. Thus, to separately identify dual-treatment samples, the observed treatment for each sample is given as  $T \in \{a_1, a_2, a_3\}$ .

**Data:** The  $n$  i.i.d. samples in a training data are described as  $D_{i=\{1..n\}} = (X_i, Z_i, T_i, Y_i)$ , where  $T_i \in \{a_1, a_2, a_3\}$  is the observed treatment,  $X_i \in \mathbb{R}^p$  are the observable confounding features,  $Z_i \in \{z : z \in \mathbb{R}^q \text{ and } T_i \neq a_1\}$  are the hidden confounding features not reported under desired treatment  $T = a_1$ , and  $Y_i \in \mathbb{R}$  is the observed outcome. The impact of hidden confounders is evident through the following propositions:

*P2.1 - Training Ignorability:* Decisions are unconfounded given  $X$  and  $Z$ :  $Y_{a_2}, Y_{a_3} \perp T|X, Z$ .

*P2.2 - Selective Confounding:* Decisions are confounded given only  $X$ :  $Y_{a_1} \not\perp T|X$ .

This means that when all confounders are observed, the potential outcomes  $Y_{a_2}$  or  $Y_{a_3}$  are directly identifiable from observed data, i.e.,  $E(Y_{a_2}|X, Z) = E(Y|X, Z, T = a_2)$ . Conversely, due to the hidden confounders, the target quantity is not identifiable, i.e.,  $E(Y_{a_1}|X) \neq E(Y|X, T = a_1)$  [17].

Next, the marginal and conditional propensity to receive any treatment  $t$  are defined as  $\pi_{T=t}(X, Z) := P(T = t|X = x, Z = z)$  and  $\pi_{T=t|T \neq t'}(X, Z) := P(T = t|X = x, Z = z, T \neq t')$  respectively. The outcome regression for treatment  $T = t$  is defined as  $\mu_t(X, Z) := E(Y_t|X = x, Z = z)$ .

**Assumptions:** To identify the target quantity  $\nu_{a_1}(x)$ , we adopt the following standard assumptions commonly used in causal inference research [30], which are easier to satisfy:

*Consistency:* A case that receives treatment  $t$  has outcome  $Y = Y_t$ . This gives  $Y = \mathbb{I}(T = a_1)Y_{a_1} + \mathbb{I}(T = a_2)Y_{a_2} + \mathbb{I}(T = a_3)Y_{a_3}$ .

*Positivity:* There is a non-zero probability of assigning any treatment to a unit, i.e.,  $\mathbb{P}(\pi_{T=t}(X, Z) \geq \epsilon > 0) = 1 \forall t \in \{a_1, a_2, a_3\}$ .

**Miscellaneous notation:** In the paper, we use both  $\hat{f}$  and  $\tilde{f}$  to denote estimate of  $f$ , and  $\mathbb{I}$  to represent the indicator function. We will use  $\mathbb{I}_t$  instead of  $\mathbb{I}_{T=t}$ . For simplicity, we denote the propensity score with short subscripts, such as  $\pi_t(\cdot)$  and  $\pi_{t|t \neq t'}(\cdot)$  for  $\pi_{T=t}(\cdot)$  and  $\pi_{T=t|T \neq t'}(\cdot)$ , respectively. Finally, we represent the target estimator as  $\hat{\nu}(x)$  instead of  $\hat{\nu}_{a_1}(x)$ .

## 3 Related Work

There has been a significant amount of research on causal learning with no hidden confounders. One of the classical methods for counterfactual estimation is matching [35, 1, 22]. Matching methods estimate the counterfactual outcomes by the nearest neighbor of each individual in terms of covariates. Because the curse of dimensionality makes finding appropriate nearest neighbors of each individual more difficult, propensity score matching, was developed [34]. Tree-based methods, such as Random forest and Bayesian additive regression trees (BART), have also been applied [7, 16]. Some literature

has focused on learning treatment-invariant representations [18, 46]. These methods are not applicable for counterfactual prediction with hidden confounders and will result in more loss [23].

Methods for addressing the problem of hidden confounders in counterfactual estimation can generally be divided into generic and specialized approaches. Generic approaches tend to focus on robustness and one such approach is sensitivity analysis [20]. It aims to minimize the worst-case estimated regret of a candidate decision policy compared to a baseline decision policy over an uncertainty set for propensity weights that control the extent of unobserved confounding. However, this approach is unsuitable for our problem settings as it requires observed outcomes from all treatments to estimate the uncertainty set for propensity weights. Another generic approach for addressing hidden confounders is learning latent representations using Variational Autoencoder (VAE) [26, 27, 15]. These methods assume that observed confounders are proxies of hidden confounders and that hidden confounders can be estimated to some extent through a generative model. The deep generative models such as VAE do not leverage the structure in *selective confounding* to outperform proposed methods. Also, they are not interpretable to be applicable in the social science domain.

In contrast to generic ones, specialized methods exploit the structure in problem setting to address hidden confounders. Similar to our approach, these methods are more specific to the applications at hand, and provide better performance gains than generic approaches.

For applications where limited RCT data is available, Kallus et al. proposed an approach to address hidden confounders [19]. Since scope of experimental data does not cover observational data, they have proposed to learn calibration parameter from experimental data to adjust counterfactual outcome of observational data. Another approach utilized the network information to unravel patterns of hidden confounders [14].

Structural nested models (SNMs) have been applied in various contexts to address the issue of hidden confounders in data that has multiple levels, such as cross-sectional or longitudinal settings [33, 44]. In our case, the focus is on safety when a child stays at home, and not how state-level predictors or programs play any role in the outcome [31]. Additionally, the longitudinal aspect of the data is irrelevant to this research question as the aim is not to study how the likelihood of abuse changes over time [31].

Another widely used approach for addressing hidden confounding in causal inference is the use of Instrument Variables (IV) [24, 6, 40]. It is important to note that when using IV, certain assumptions must be met such as: 1) the IV being strongly correlated with the treatment variable, 2) independent of the outcome variable and any unobserved confounding factors, and 3) not affecting the outcome variable through any other path than through the treatment variable. In our study, it's not possible that decision policy will always be a valid IV. For instance, in child placement example, a lenient decision policy results in other programs offered by state that ensures In-Home treatment is safe.

Negative control outcomes are used to address the problem of hidden confounding by providing a way to control for the effects of hidden confounding factors [43, 29]. The idea is to find an outcome variable that is not affected by the treatment or exposure of interest, but is affected by the same unobserved confounding factors. By comparing the treatment effect on the negative control outcome to the treatment effect on the outcome of interest, it is possible to control for the effects of unobserved confounding factors. These approaches are only suitable for applications where negative control exposures or outcomes exist. This technique is not suitable for societal applica-

tions, as it's not possible to record additional outcomes due to privacy concerns.

The work closest to ours deals with run-time confounding only [10]. For them, training nuisance estimator without confounding bias was straightforward due to no hidden confounders during training. Our work address *selective confounding*, where confounding bias is present during both training and prediction.

## 4 Methodology

The target model  $\hat{\nu}(x)$  is usually termed as a counterfactual predictor [10]. Since in an observational study, all samples did not get the desired-treatment to have the factual outcome  $Y_{a_1}$ . In this section, we start by introducing the standard method to learn the target model. Then will present the proposed two-step approaches.

### 4.1 Standard Predictor (SP)

One straightforward method to estimate the target quantity  $\hat{\nu}(x)$  is to ignore all samples that did not receive the proposed treatment, i.e., those with  $T \neq a_1$ . This approach involves learning  $\hat{\mathbb{E}}[Y|T = a_1, X = x]$  to estimate the target quantity. However, this method may suffer from selective bias, which is relatively easier to handle if we have access to all the confounders, meaning  $T \perp Y_{a_1}|X$ .

The challenge arises under *Selective Confounding*, where hidden confounders  $Z$  are present, and the quantities are no longer conditionally independent. In such cases, this method does not target the correct counterfactual quantity and can produce misleading predictions due to confounding bias [41].

### 4.2 Counterfactual Predictors

Standard method involves training the model on samples with available ground truth labels  $Y_{a_1}$ , leading to an estimator  $\hat{\omega}_{a_1}(x) = \hat{\mathbb{E}}(Y_{a_1}|X) = \hat{\mathbb{E}}(Y|X, T = a_1)$  [38, 11]. However, this approach will have no confounding bias under certain conditions, such as when all confounders are observed as given in proposition P.2.1, or in randomized controlled trials (RCT) [19].

To overcome confounding bias, we propose leveraging dual-treatment samples, which are defined in Section 2. The intuition behind our solution is that dual-treatment samples are unconfounded and it is possible to estimate potential outcome  $\hat{Y}_{a_1}$  for them. In that case, the outcome estimator  $\hat{\mu}_{a_1}(x, z) = \hat{\mathbb{E}}(\hat{Y}_{a_1}|X, Z) = \hat{\mathbb{E}}(\hat{Y}_{a_1}|X, Z, T = a_3)$  and the conditional propensity score estimator  $\hat{\pi}_{a_3|\neq a_1} = \hat{P}(T = a_3|X = x, Z = z, T \neq a_1)$  can be identified without any confounding bias. These two nuisance estimators with no confounding bias can then be used to train the target estimator  $\hat{\nu}(x)$  using two-step methods like Regression Adjustment (RA) and Doubly Robust (DR).

#### 4.2.1 Data Preprocessing - Treatment Adjustment for Dual-Treatment Samples

In order to learn target estimator using a two-step procedure, we require potential outcome label  $Y_{a_1}$  for dual-treatment samples but the observed label is  $Y_{a_3}$  for them. The difference between the outcomes of two treatments is usually given as the Conditional Average Treatment Effect (CATE) given in eq.1. By design, samples from desired-treatment  $T = a_1$  are closely related to dual-treatment samples. It is possible for certain applications that  $\text{CATE}(a_1, a_3)$  is negligible. In that case, the observed label  $Y_{a_3}$  for dual-treatment samples

is equivalent to potential outcome  $Y_{a_1}$  and the outcome regressor is identifiable by  $\mu_{a_1}(x, z) = \mathbb{E}(Y_{a_1}|X, Z) = \mathbb{E}(Y_{a_3}|X, Z) = \mathbb{E}(Y|X, Z, T = a_3)$ .

$$\text{CATE}(a_1, a_3) = \mathbb{E}(Y_{a_1}|X, Z) - \mathbb{E}(Y_{a_3}|X, Z) \quad (1)$$

In case there is a possibility that CATE might be non-zero between two treatments  $a_1, a_3$ , then it's necessary to estimate potential outcome  $\hat{Y}_{a_1}$  for dual-treatment samples.

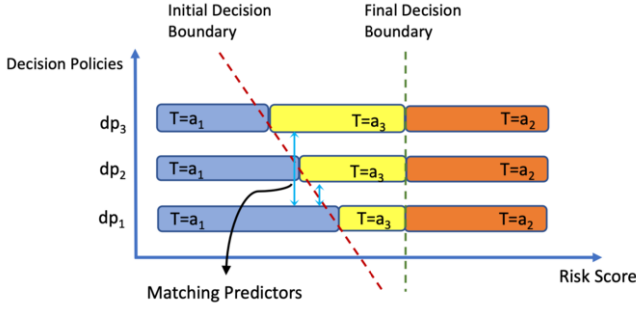


Figure 1: Decision policies across different locations

The proposed solution works irrespective of the CATE value. It assumes there exist multiple locations with slightly different initial decision policies (dp) as shown in Figure 1. It represents a real-world scenario when some judges are strict and some are lenient [25]. The same is the case in child welfare, where some initial placement screenings are performed by police officers or by trained child welfare professionals [37] or state policies are different [28].

As mentioned earlier, the regressor  $\tilde{\omega}_{a_1}(x) = \hat{\mathbb{E}}(Y|X, T = a_1)$  will be erroneous due to confounding bias. However, in the case of randomized controlled trials (RCT), where  $Y_{a_1} \perp T$ , there will be no confounding bias for  $\tilde{\omega}_{a_1}(x)$  even with hidden confounders [19]. Although observed samples from the whole population do not result from RCT, here we only need to estimate  $\hat{Y}_{a_1}$  for dual-treatment samples. In cases where we have varying decision boundaries between locations, the scope of samples with treatment  $T = a_1$  from one location covers samples with treatment  $T = a_3$  from another location, resembling an RCT between two treatments  $a_1$  and  $a_3$ .

For instance, let's consider there are three different decision policies (DPs) resulting in different decision boundaries for the initial decision as shown in Figure 1. It shows  $dp_1$  is the lenient, and  $dp_3$  is the strictest. Also, it is expected that the final decision boundary is uniform among locations since the review process is to ensure removing any personal biases in the judgment process. As a result, the scope of desired-treatment sample  $T = a_1$  in  $dp_1$  will cover the scope of dual-treatment samples  $T = a_3$  in  $dp_2$  and  $dp_3$ . Eventually when samples from all locations are considered as a whole, it will resemble a scenario of RCT between treatments  $a_1, a_3$  with  $P(T = a_1|X, Z, T \neq a_2) = P(T = a_3|X, Z, T \neq a_2) \approx 0.5$ .

The two subpopulations with treatments  $a_1$  and  $a_3$  satisfy  $Y_{a_1} \perp T|T \neq a_2$ . Thus, it is possible to correctly estimate potential outcome  $\hat{Y}_{a_1}$  for dual-treatment samples from  $\tilde{\omega}_{a_1}(x)$  without any confounding bias. The detailed procedure is provided in Algorithm 1. First, it learns  $\tilde{\omega}_{a_1}(x) = \hat{\mathbb{E}}(Y|T = a_1, x)$ , and then estimates potential outcome  $\hat{Y}_{a_1}$  for dual-treatment samples. In Section 5.1, we have presented the effectiveness of this approach with varying CATE over synthetic data.

---

**Algorithm 1** Label Estimation for dual-treatment samples

---

**Input:** Data  $D_{d \in \{dp_1 \dots dp_L\}} = \{(X_i, Z_i, T_i, Y_i)\}_{i=1}^{n_d}$

1. Train regressor  $\tilde{\omega}_{a_1}(x)$  over  $Y \sim X|T = a_1$
2. Estimate  $\hat{Y}_{a_1}$  from  $\tilde{\omega}_{a_1}(x)$  for samples with  $T = a_3$

**Return:** Desired Labels  $\hat{Y}_{a_1}$  for dual-treatment samples

---

4.2.2 Regression-Adjustment (RA) Predictor

In the Regression Adjustment (RA) method, target estimator  $\hat{\nu}(x)$  is learned in two stages. The procedure is presented in Algorithm 2. First a nuisance estimator  $\hat{\mu}_{a_1}(X, Z) = \hat{\mathbb{E}}(\hat{Y}_{a_1}|X, Z, T = a_3)$  is learned with no confounding bias using dual-treatment samples. Then similar to the semi-supervised learning approach the nuisance estimator will be used to provide a pseudo-outcome for the whole population. Then target model  $\hat{\nu}(x)$  is estimated over the whole population using ground truth label for samples received treatment  $T = a_1$  and pseudo outcome  $\hat{\mu}_{a_1}(X, Z)$  for other samples with  $T = \{a_2, a_3\}$ .

For theoretical error bounds, we have adopted a procedure similar to [21] which bounds the error for a two-stage regression on the full set of confounding variables.

**Theorem 1** Under sample-splitting to learn  $\hat{\mu}_{a_1}(x, z), \hat{\nu}(x)$ , the RA method has pointwise regression error that is bounded by:

$$\begin{aligned} \mathbb{E} \left[ (\hat{\nu}(x) - \nu(x))^2 \right] &\lesssim \mathbb{E} \left[ (\tilde{\nu}(X) - \nu(X))^2 \right] \\ &+ \mathbb{E} \left[ (\hat{\mu}_{a_1}(x, Z) - \mu_{a_1}(x, Z))^2 | X = x \right] \\ &+ \mathbb{E} \left[ (\Delta(x, Z))^2 | X = x \right] \quad (2) \end{aligned}$$

where  $\hat{\nu}(x)$  is the final regressor. The right-hand side splits the error bounds according to the different stages of the algorithm, where  $\tilde{\nu}(x) = \hat{\mathbb{E}}(\mathbb{I}_{a_1} Y + \mathbb{I}_{\neq a_1} \hat{\mu}_{a_1}(X, Z) | X = x)$  denotes second-stage regression with oracle access to the first-stage output.  $\hat{\mu}_{a_1}(X, Z) = \hat{\mathbb{E}}(\hat{Y}_{a_1} | X, Z)$  is the first stage regressor with oracle access to  $Y_{a_1}$  and  $\Delta(X, Z) = \mathbb{E}(\hat{Y}_{a_1} - Y_{a_1} | X, Z, T = a_3)$  is the label estimation error for dual-treatment samples. Notably, the bound depend linearly on the Mean square error of  $\hat{\mu}_{a_1}(x, Z)$ . The next approach provides a more robust mechanism by introducing another first-stage estimator.

---

**Algorithm 2** Regression-Adjustment (RA) Algorithm

---

**Input:** Data  $D = \{(X_i, Z_i, T_i, Y_i)\}_{i=1}^n$

Randomly divide  $D$  into two partitions  $W^1, W^2$

**for**  $(i, j) \in \{(1, 2), (2, 1)\}$  **do**

**Stage 1:** On  $W^i$ , learn regression function  $\hat{\mu}_{a_1}(X, Z)$  over  $\hat{Y}_{a_1} \sim X = x, Z = z | T = a_3$

**Stage 2:** On  $W^j$ , learn regressor  $\hat{\nu}^j(x)$  over  $\mathbb{I}_{a_1} Y + \mathbb{I}_{\neq a_1} \hat{\mu}_{a_1}(X, Z) \sim X$

**end for**

**Prediction:**  $\hat{\nu}(x) = \sum_{q=1}^2 \hat{\nu}^q(x)$

---

### 4.2.3 Doubly-Robust (DR) Predictor

This method introduces two nuisance estimators, i.e.  $\hat{\mu}_{a_1}(X, Z)$  and  $\hat{\pi}_{a_3|a_1}(X, Z)$ . It guarantees robustness since it requires only one of the nuisance estimators to be closer to the true estimator. The DR method is presented in Algorithm 3. Again it's a two-step procedure. At the first, it trains regressor  $\hat{\mu}_{a_1}(X, Z)$  and conditional propensity score model  $\hat{\pi}_{a_3|a_1}(X, Z)$ . The propensity score model allows us to give higher weightage to those dual-treatment samples that resemble samples from treatment  $T = a_2$ . The second step of the DR algorithm involves training the target model  $\hat{\nu}(x)$  over the whole population.

DR method provides better theoretical guarantees as compared to RA due to two nuisance estimators.

**Theorem 2** *Under sample-splitting to learn  $\hat{\mu}_{a_1}(x, z)$ ,  $\hat{\nu}(x)$  and  $\hat{\pi}_{a_3|a_1}(x, z)$  the pointwise error bound for proposed DR method is given as:*

$$\begin{aligned} \mathbb{E} \left[ (\hat{\nu}(x) - \nu(x))^2 \right] &\lesssim \mathbb{E} \left[ (\tilde{\nu}(x) - \nu(x))^2 \right] \\ &+ \mathbb{E} \left[ (\tilde{\mu}_{a_1}(x, Z) - \mu_{a_1}(x, Z))^2 | X=x \right] \\ &\times \mathbb{E} \left[ (\hat{\pi}_{a_3|a_1}(x, Z) - \pi_{a_3|a_1}(x, Z))^2 | X=x \right] \\ &+ \mathbb{E} \left[ (\Delta(x, Z))^2 | X=x \right] \quad (3) \end{aligned}$$

where  $\hat{\nu}(x)$  is the final regressor and the error bound is split between different stages at the right-hand side.  $\tilde{\nu}(x)$  denotes second-stage regressor with oracle access to true first-stage estimators  $\mu_{a_1}(x, Z)$ ,  $\pi_{a_3|a_1}(x, Z)$ .  $\tilde{\mu}_{a_1}(x, Z)$  is the first stage regressor with oracle access to  $Y_{a_1}$  and  $\hat{\pi}_{a_3|a_1}(x, Z)$  is the propensity score estimator. Finally,  $\Delta(x, Z) = \mathbb{E}(Y_{a_1} - Y_{a_3} | X, Z, T = a_3)$  is the desired label estimation error for dual-treatment samples. Here, in the error bound there is a product between errors of nuisance estimators. It guarantees robustness because only one of the nuisance estimators needs to be converged. The error bound still linearly depends on the label estimation error of dual-treatment samples. However, that estimation error over subpopulation is much simpler to handle as compared to the target problem of the whole population. Also, in some cases label estimation for dual-treatment samples might not be needed at all where  $\text{CATE}(a_1, a_3)$  is negligible, as mentioned in Section 4.2.1.

---

#### Algorithm 3 Doubly-Robust (DR) Algorithm

---

**Input:**  $D = \{(X_i, Z_i, T_i, Y_i)\}_{i=1}^n$

Randomly divide  $D$  into three partitions  $W^1, W^2, W^3$

**for**  $(i, j, k) \in \{(1, 2, 3), (3, 1, 2), (2, 3, 1)\}$  **do**

**Stage 1:** On  $W^i$ , train regression function  $\hat{\mu}_{a_1}(X, Z)$  over  $\hat{Y}_{a_1} \sim X = x, Z = z | T = a_3$

On  $W^j$ , train conditional propensity estimator  $\hat{\pi}_{a_3|a_1}(X, Z)$  over  $\mathbb{I}_{a_3} \sim X = x, Z = z | T \neq a_1$

**Stage 2:** On  $W^k$ , train  $\hat{\nu}^k(x)$  by regressing  $\mathbb{I}_{a_1} Y + \mathbb{I}_{\neq a_1} \left[ \frac{\mathbb{I}_{a_3}}{\hat{\pi}_{a_3|a_1}(X, Z)} (\hat{Y}_{a_1} - \hat{\mu}_{a_1}(X, Z)) + \hat{\mu}_{a_1}(X, Z) \right] \sim X$

**end for**

**Prediction:**  $\hat{\nu}(x) = \sum_{q=1}^3 \hat{\nu}^q(x)$

---

## 5 Experiments

The evaluation of counterfactual prediction algorithms is challenging due to the absence of ground truth labels for subpopulations with different treatments. To overcome this challenge, counterfactual evaluation is typically performed using synthetic, semi-synthetic, or RCT data [45, 18]. In this study, we evaluated the proposed counterfactual prediction algorithms using both synthetic and real-world child placement data. For comparison we have either used SP as baseline approach or human decision maker in child-placement data. We are not aware that there exist any other techniques that are both interpretable and able to predict under *selective confounding*, as discussed in Section 3.

### 5.1 Synthetic-Data Evaluation

**Data:** The synthetic data is generated for  $L = 20$  locations to simulate different decision policies. For each location, the acceptance rate between  $ar = [0.3, 0.5]$  is randomly picked which will determine how strict or lenient the initial decision policy is. Once determined, the acceptance rate will stay fixed for each location. There is  $n = 1000$  number of data points in each location.

The confounders  $X, Z$  are sampled from Normal Distribution and the correlation among them is controlled using parameter  $\rho$ .

$$X_i \sim \mathbb{N}(0, 1), \quad Z_i \sim \mathbb{N}(\rho X_i, 1 - \rho^2) \quad ; 0 \leq i \leq d \quad (4)$$

where  $2 * d$  is the total number of confounders present in a single sample. The potential outcome  $\mu_{a_1}, \mu_{a_3}$  under treatments  $a_1, a_3$  is determined by:

$$\mu_{a_1}(X, Z) = \frac{k_x}{k_x + \rho k_z} \left( \sum_{i=0}^{k_x} X_i + \sum_{i=0}^{k_z} Z_i \right) + \epsilon \quad (5)$$

$$\begin{aligned} \mu_{a_3}(X, Z) = \frac{k_x}{k_x + \rho k_z} &\left[ \sum_{i=0}^d X_i + \sum_{i=0}^d Z_i \right. \\ &\left. - \tau \left( \sum_{i=0}^d X_i + \sum_{i=0}^d Z_i^2 \right) \right] + \epsilon \quad (6) \end{aligned}$$

where,

$$\epsilon \sim \mathbb{N}\left(0, \frac{1}{2n} \|\mu_{a_1}(X, Z)\|_2^2\right)$$

and  $k_x, k_z$  determines the number of confounders impacting the outcome. CATE between two treatments  $a_1, a_3$  is controlled with parameter  $\tau$  in eq. 6. To compare learned estimator  $\hat{\nu}(x)$ , the ground truth outcome based on confounders  $X$  is given as:


$$\nu(X) = \frac{k_x}{k_x + \rho k_z} \left[ \sum_{i=0}^{k_x} X_i + \rho \sum_{i=0}^{k_z} X_i \right] \quad (7)$$

The propensity score, probability of treatment assignment, is given as  $\pi(X, Z)$  in eq. 8. To replicate real-world scenarios where decision reversal is possible through appeals and gives dual-treatment samples, we simulate two decision stages  $d_1, d_2$ , given in eq. 9 and 10 respectively, corresponding to the initial and final decision.


$$\pi(X, Z) = \sigma \left( \frac{1}{\sqrt{k_x + k_z}} \left( \sum_{i=0}^{k_x} X_i + \sum_{i=0}^{k_z} Z_i \right) \right) \quad (8)$$

$$d_1 = \text{Bernoulli}(\min(\frac{\pi(X, Z)}{\text{ar}} * 0.5, 0.99)) \quad (9)$$


$$d_2 = \text{Bernoulli}(\pi(X, Z)) \quad (10)$$



**Figure 2:** Performance under Varying Treatment Effects



**Figure 3:** Performance Comparison over Confounders Correlation



**Figure 4:** Performance comparison over Hidden-Confounders

where  $\sigma(x) = \frac{1}{1+e^{-x}}$  and  $ar = [0.3, 0.5]$  is the acceptance rate randomly picked for each location in the dataset. Lower  $ar$  values mean a strict decision policy as it brings down the decision boundary from 0.5. Based on  $d_1, d_2$ , the assigned treatment is given as:

$$T = \begin{cases} a_1 & d_1 = 0 \\ a_2 & d_1 = 1 \ \& \ d_2 = 1 \\ a_3 & d_1 = 1 \ \& \ d_2 = 0 \end{cases} \quad (11)$$

**Evaluation Metric:** The dataset contain samples from all  $L$  locations. For evaluation, the dataset is partitioned into train and test sets. The metric used to evaluate the performance of target estimator  $\hat{\nu}(x)$  is Mean Square Error (MSE) given as:

$$MSE = \frac{1}{n} \sum_{i=1}^n [(\nu(x_i) - \hat{\nu}(x_i))^2] \quad (12)$$

where  $n$  is the total number of points in the test set.

**Results:** Our solution comprises of two main components: 1) Estimation of potential outcome  $\hat{Y}_{a_1}$  for dual-treatment samples using Algorithm 1 and 2) Training a counterfactual predictor model  $\hat{\nu}(x)$  using RA (Algorithm 2) and DR (Algorithm 3) schemes.

Figure 2 demonstrates Algorithm 1’s effectiveness by comparing its performance over varying CATE( $a_1, a_3$ ) with three techniques for handling labels in dual-treatment samples. The estimator  $\hat{\nu}(x)$  used only the DR scheme to ensure a fair comparison, and setting **A** in Table 1 was used to generate data for this experiment. The error lower bound is obtained with oracle access to potential outcome  $Y_{a_1}$  for dual-treatment samples. In the case of label estimation  $\hat{Y}_{a_1}$  is not performed for dual-treatment samples and their observed label  $Y_{a_3}$  is used as it is, performance deteriorates with increasing  $\tau$ . However, when Algorithm 1 is used to estimate  $\hat{Y}_{a_1}$  for dual-treatment samples, performance does not degrade with  $\tau$ . Moreover, the proposed algorithm 1 performs similarly to using observed label  $Y_{a_3}$  as it is when the treatment effect is negligible (i.e., when  $\tau \approx 0$ ). Thus, the proposed algorithm 1 can be used in any situation irrespective of the value of  $\tau$ .

**Table 1:** Experimental settings over synthetic-data

Setting	$d$	$k_x$	$k_z$	$\rho$	$\tau$
<b>A</b>	250	25	25	0.25	*
<b>B</b>	250	25	25	*	0.5
<b>C</b>	250	25	*	0.25	0.5

In Figure 3, we compare the final predictor approaches of AR and DR with the baseline SP. To generate data, we have used Setting **B** in Table 1. We plot the performances of three algorithms against

the correlation  $\rho$  between hidden and observed confounders. As  $\rho$  increases, the observed features  $X$  serve as proxies for  $Z$ . In such cases, SP performs comparably due to negligible confounding bias. However, when the correlation is low meaning confounding bias will be higher in case of SP, there is a clear advantage of using proposed counterfactual estimation approaches AR and DR presented in Algorithm 2, 3 respectively.

Figure 4 also demonstrate prediction algorithms’ performance over confounding bias. Here, we control confounding bias using  $k_z$ . The parameter setting for data generation is Setting **C** in Table 1. As expected, the performance of all algorithms decreases as the value of  $k_z$  increases. However, the proposed counterfactual approaches, AR and DR, outperform the baseline SP when the confounding bias is non-zero.

It is important to note that in our experiments, there was no clear advantage of using the DR scheme over RA. Since, DR guarantees robustness, in case the outcome regressor  $\hat{\mu}_{a_1}(x, z)$  does not converge properly, instead of gain in performance.

### 5.2 Real-World Data Evaluation

**Data:** The real-world data used in this study consists of placement decisions from eight states in the USA, with In-Home placement ( $T = a_1$ ) as the desired treatment. While the eight locations cover a wide area with overlapping attributes, each state may handle cases differently. The observed features  $X$  include demographics, case records, and complaints, while hidden features  $Z$  include substantiated maltreatment, mental health, and disability type. The relevant outcome is whether or not further maltreatment occurred within six months of In-Home placement.

Figure 5 displays  $N = 5000$  samples randomly drawn from each location. To compare the characteristics of sub-populations or the effectiveness of initial decision criteria among each state, two new metrics are introduced: Acceptance Rate (AR) and Failure Rate (FR), given as follows:

$$AR = \frac{\text{Number of In-Home samples}}{N}$$

$$FR = \frac{\text{Number of maltreated In-Home samples}}{N}$$

While a high AR usually results in a high FR and vice versa, not all locations follow the same trend, as shown in Figure 5. This suggests that In-home sample distribution may vary among locations, and certain locations may have room for improvement in terms of FR.

**Evaluation:** To circumvent the issue of missing ground truth labels in evaluation, we have devised three new methods that are applicable






Figure 5: Real-World data Statistics

in applications similar to child placement. The first method enables us to quantify the benefit of using a prediction model as compared to human decision-makers [25]. Meanwhile, the second and third methods allow us to compare different prediction models to select the best.

**Method 1 - Comparison with Humans:** In the child placement application, initial decisions may be wrong in two ways 1) higher FR among the In-home population, and 2) significant proportions of dual-treatment samples. For comparison, we see if the proposed prediction model (DR) can switch high-risk in-home samples with low-risk dual-treatment samples. The risk-score for dual-treatment samples can be estimated using Algorithm 1.

First using DR method, we estimate the risk score for the In-Home and dual-treatment samples. Then without modifying the AR of each location, assign In-Home treatment based on predicted risk scores of In-Home and dual-treatment samples. It can be seen in Figure 6, that prediction models bring down the FR for all locations. Also, it means that the proportion of dual-treatment samples will decrease as some of them are now assigned in-home treatment by the prediction model. After further analyses, we have discovered that the most common predictor responsible for *short-stays* is reported as child neglect [36], and our prediction algorithm rightfully assigns negative weightage to this predictor while determining the risk score.

**Method 2 - Evaluation over Outlier State:** In Figure 5, AR and FR for different states are presented. Two locations, A and G, stand out with relatively higher FR and lower AR as compared to other locations. It means their *In-Home* samples contain a significant number of high-risk samples that are usually *Out-of-Home* samples in other locations. It allows us to train predictors over other locations B to F and evaluate them over In-Home samples of Location A and G. The results are reported in Table 2, and it can be seen that MSE is different between Standard Predictors (SP) and counterfactual predictors (RA, DR). First, it validates the evaluation method and also shows a gain in performance by circa 10% due to counterfactual predictors.

Table 2: MSE Comparison from Proposed Evaluation Methods

Prediction Algorithm	Method 2	Method 3
Standard Predictor (SP)	0.8	0.83
Regression Adjustment (RA)	0.69	0.39
Doubly-Robust (DR)	0.68	0.26

**Method 3 - Doubly-Robust Evaluation:** The prediction error is identifiable when there are no hidden confounders under desired treatment as  $\mathbb{E}((Y_{a_1} - \hat{\nu}(x))^2 | X, Z) = \mathbb{E}((\hat{Y}_{a_1} - \hat{\nu}(x))^2 | X, Z, T = a_3)$ . We propose the Doubly-Robust Evaluation approach similar to




Figure 6: FR Comparison with Similar AR

[10], with the difference that it can handle *selective confounding*. Our evaluation procedure identifies error term using dual-treatment samples. Defining identifiable error regressor as  $\eta(X, Z) := \mathbb{E}((\hat{Y}_{a_1} - \hat{\nu}(x))^2 | X, Z, T = a_3)$ , the doubly-robust estimate of MSE is given as:

$$\frac{1}{n} \sum_{i=1}^n \left[ \mathbb{I}_{\neq a_1}^i \left[ \frac{\mathbb{I}_{a_3}^i}{\hat{\pi}_{a_3 | \neq a_1}(X_i, Z_i)} ((\hat{Y}_{a_1}^i - \hat{\nu}(x_i))^2 - \hat{\eta}(X_i, Z_i)) + \hat{\eta}(X_i, Z_i) \right] + \mathbb{I}_{a_1}^i (Y_i - \hat{\nu}(x_i))^2 \right] \quad (13)$$

The results are reported in Table 2. It can be seen that there is a significant gain in using counterfactual predictors RA and DR as compared to SP. It is interesting to note that DR performs slightly better than RA which is highlighted by this method.

## 6 Summary

Counterfactual predictions for desired treatments play a crucial role in aiding decision-makers in real-world applications. However, assuming all confounders or their proxies are observable might not be realistic. In this work, we address this limitation by considering the scenario of *selective confounding*, where confounders can be hidden under a desired treatment. Despite this relaxation, our proposed method guarantees the identifiability of the target quantity and offers an interpretable target model, subject to two key assumptions: 1) The existence of different decision makers and 2) the possibility of appeal. To train the target model, we employ a two-stage counterfactual learning technique with no confounding bias over nuisance estimators. This allows us to use Linear Regression as the target model, making it inherently interpretable and suitable for societal applications. Our empirical results demonstrate that the proposed framework surpasses both human decision-makers and conventional statistical learning models in terms of performance. While we provide results for child-placement examples, the proposed method can be extended to other applications such as ICU admission and police arrests that satisfy the two conditions.

## References

- [1] Alberto Abadie and Guido W Imbens, 'Large sample properties of matching estimators for average treatment effects', *econometrica*, (2006).
- [2] Charles Bhatti-Sinclair, Kish Sutcliffe, 'Challenges in identifying factors which determine the placement of children in care? an international review', *Child and Adolescent Social Work Journal*, (2013).
- [3] Ioana Bica, Ahmed M. Alaa, James Jordon, and Mihaela van der Schaar, 'Estimating counterfactual treatment outcomes over time through adversarially balanced representations', in *8th International Conference on Learning Representations ICLR*, (2020).

- [4] Léon Bottou, Jonas Peters, Joaquin Quiñero-Candela, Denis X. Charles, D. Max Chickering, Elon Portugaly, Dipankar Ray, Patrice Simard, and Ed Snelson, 'Counterfactual reasoning and learning systems: The example of computational advertising', *Journal of Machine Learning Research (JMLR)*, (2013).
- [5] Tim Brennan, William Dieterich, and Beate Ehret, 'Evaluating the predictive validity of the compas risk and needs assessment system', *Criminal Justice and Behavior*, (2009).
- [6] Peter Bühlmann, Jonas Peters, and Jan Ernest, 'Cam: Causal additive models, high-dimensional order search and penalized regression', *The Annals of Statistics*, (2014).
- [7] Hugh A Chipman, Edward I George, and Robert E McCulloch, 'Bart: Bayesian additive regression trees', *The Annals of Applied Statistics*, (2010).
- [8] Ka Ho Brian Chor, Gary M. McClelland, Dana A. Weiner, Neil Jordan, and John S. Lyons, 'Patterns of out-of-home placement decision-making in child welfare', *Child Abuse & Neglect*, (2013).
- [9] Ka Ho Brian Chor, Gary M. McClelland, Dana A. Weiner, Neil Jordan, and John S. Lyons, 'Out-of-home placement decision-making and outcomes in child welfare: A longitudinal study', *Administration and Policy in Mental Health and Mental Health Services Research*, (2014).
- [10] Amanda Coston, Edward H. Kennedy, and Alexandra Chouldechova, 'Counterfactual predictions under runtime confounding', in *33: Annual Conference on Neural Information Processing Systems (NeurIPS)*, (2020).
- [11] Amanda Coston, Alan Mishler, Edward H. Kennedy, and Alexandra Chouldechova, 'Counterfactual risk assessments, evaluation, and fairness', in *Proceedings of the Conference on Fairness, Accountability, and Transparency (FAT)*, (2020).
- [12] Michele Jonsson Funk, Daniel Westreich, Chris Wiesen, Til Stürmer, M. Alan Brookhart, and Marie Davidian, 'Doubly robust estimation of causal effects', *American Journal of Epidemiology*, (2011).
- [13] Zoe Getz, Cassandra Simmel, Liwei Zhang, and Brett Greenfield, "'short-stayers" in child welfare: Characteristics and system experiences', *Children and Youth Services Review*, (2022).
- [14] Ruo Cheng Guo, Jundong Li, and Huan Liu, 'Learning individual causal effects from networked observational data', in *Proceedings of the 13th International Conference on Web Search and Data Mining, WSDM '20*. Association for Computing Machinery, (2020).
- [15] Shonosuke Harada and Hisashi Kashima, 'Infocevae: treatment effect estimation with hidden confounding variables matching', *Machine Learning*, 1–19, (2022).
- [16] Jennifer L Hill, 'Bayesian nonparametric modeling for causal inference', *Journal of Computational and Graphical Statistics*, (2011).
- [17] Guido W. Imbens and Donald B. Rubin, *Causal Inference for Statistics, Social, and Biomedical Sciences: An Introduction*, Cambridge University Press, USA, 2015.
- [18] Fredrik D. Johansson, Uri Shalit, and David Sontag, 'Learning representations for counterfactual inference', in *Proceedings of the 33rd International Conference on International Conference on Machine Learning (ICML)*, (2016).
- [19] Nathan Kallus, Aahlad Manas Puli, and Uri Shalit, 'Removing hidden confounding by experimental grounding', in *Advances in Neural Information Processing Systems (NeurIPS)*, eds., S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, (2018).
- [20] Nathan Kallus and Angela Zhou, 'Confounding-robust policy improvement', in *Advances in Neural Information Processing Systems (NeurIPS)*, eds., S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, (2018).
- [21] Edward H. Kennedy, 'Optimal doubly robust estimation of heterogeneous causal effects', *arXiv: Statistics Theory*, (2020).
- [22] Gary King and Richard Nielsen, 'Why propensity scores should not be used for matching', *Political Analysis*, 27(4), 435–454, (2019).
- [23] Toru Kitagawa and Aleksey Tetenov, 'Who should be treated? empirical welfare maximization methods for treatment choice', *Econometrica*, (2018).
- [24] Zhaobin Kuang, Frederic Sala, Nimit Sohoni, Sen Wu, Aldo Córdova-Palomera, Jared Dunnmon, James Priest, and Christopher Ré, 'Ivy: Instrumental variable synthesis for causal inference', in *International Conference on Artificial Intelligence and Statistics*. PMLR, (2020).
- [25] Himabindu Lakkaraju, Jon Kleinberg, Jure Leskovec, Jens Ludwig, and Sendhil Mullainathan, 'The selective labels problem: Evaluating algorithmic predictions in the presence of unobservables', in *Conference on Knowledge Discovery and Data Mining (KDD)*, (2017).
- [26] Christos Louizos, Uri Shalit, Joris Mooij, David Sontag, Richard Zemel, and Max Welling, 'Causal effect inference with deep latent-variable models', in *Proceedings of the 31st International Conference on Neural Information Processing Systems (NeurIPS)*, (2017).
- [27] David Madras, Elliot Creager, Toniann Pitassi, and Richard Zemel, 'Fairness through causal awareness: Learning causal latent-variable models for biased data', in *Proceedings of the Conference on Fairness, Accountability, and Transparency (FAT)*, (2019).
- [28] Ferol E. Mennen, Matthew Brensilver, and Penelope K. Trickett, 'Do maltreated children who remain at home function better than those who are placed?', *Children and youth services review*, (2010).
- [29] Wang Miao, Zhi Geng, and Eric J Tchetgen Tchetgen, 'Identifying causal effects with proxy variables of an unmeasured confounder', *Biometrika*, (2018).
- [30] Judea Pearl, 'Theoretical impediments to machine learning with seven sparks from the causal revolution', *Proceedings of the Eleventh International Conference on Web Search and Data Mining*, (2018).
- [31] James L. Peugh, 'A practical guide to multilevel modeling', *Journal of School Psychology*, (2010).
- [32] Yoshihiko Raita, Carlos Camargo, Liming Liang, and Kohei Hasegawa, 'Big data, data science, and causal inference: A primer for clinicians', *Frontiers in Medicine*, (2021).
- [33] James M. Robins, 'Correcting for non-compliance in randomized trials using structural nested mean models', *Communications in Statistics - Theory and Methods*, (1994).
- [34] Paul R. Rosenbaum and Donald B. Rubin, 'The central role of the propensity score in observational studies for causal effects', *Biometrika*, (1983).
- [35] Donald B Rubin, 'Matching to remove bias in observational studies', *Biometrics*, (1973).
- [36] Vivek S. Sankaran and Christopher E. Church, 'Easy come, easy go: The plight of children who spend less than thirty days in foster care', *University of Pennsylvania Journal of Law and Social Change*, (2016).
- [37] Vivek S. Sankaran, Christopher E. Church, and Monique T. Mitchell, 'A cure worse than the disease? the impact of removal on children and their families', *Marquette Law Review*, (2019).
- [38] Peter Schulam and Suchi Saria, 'Reliable decision support using counterfactual models', in *Advances in Neural Information Processing Systems (NeurIPS)*, (2017).
- [39] Uri Shalit, 'Can we learn individual-level treatment policies from clinical data?', *Biostatistics*, (2019).
- [40] Rahul Singh, Maneesh Sahani, and Arthur Gretton, 'Kernel instrumental variable regression', *Advances in Neural Information Processing Systems (NeurIPS)*, (2019).
- [41] Andrea C Skelly, Joseph R Dettori, and Erika D Brodt, 'Assessing bias: the importance of considering confounding', *Evidence-based spine-care journal*, (2012).
- [42] Peter Spirtes. A tutorial on causal inference, 2018.
- [43] Eric Tchetgen Tchetgen, 'The control outcome calibration approach for causal inference with unobserved confounding', *American journal of epidemiology*, (2014).
- [44] Stijn Vansteelandt and Marshall Joffe, 'Structural Nested Models and G-estimation: The Partially Realized Promise', *Statistical Science*, (2014).
- [45] Liuyi Yao, Zhixuan Chu, Sheng Li, Yaliang Li, Jing Gao, and Aidong Zhang, 'A survey on causal inference', *Conference on Knowledge Discovery and Data Mining (KDD)*, (2021).
- [46] Weijia Zhang, Lin Liu, and Jiuyong Li, 'Treatment effect estimation with disentangled latent factors', *Proceedings of the AAAI Conference on Artificial Intelligence*, (2021).
- [47] Siyuan Zhao and Neil T. Heffernan, 'Estimating individual treatment effect from educational studies with residual counterfactual networks', in *Educational Data Mining*, (2017).
- [48] Min Zheng, Jesseca K. Marsh, Jeffrey V. Nickerson, and Samantha Kleinberg, 'How causal information affects decisions', *Cognitive Research: Principles and Implications*, (2020).