Robust Machine Learning for Treatment Effects in Multilevel Observational Studies Under Cluster-level Unmeasured Confounding

Youmi Suk ^{*1} and Hyunseung Kang ^{†2}

¹School of Data Science, University of Virginia

²Department of Statistics, University of Wisconsin-Madison

September 18, 2021

Abstract

Recently, machine learning (ML) methods have been used in causal inference to estimate treatment effects in order to reduce concerns for model mis-specification. However, many ML methods require that all confounders are measured to consistently estimate treatment effects. In this paper, we propose a family of ML methods that estimate treatment effects in the presence of cluster-level unmeasured confounders, a type of unmeasured confounders that are shared within each cluster and are common in multilevel observational studies. We show through simulation studies that our proposed methods are robust from biases from unmeasured cluster-level confounders in a variety of multilevel observational studies. We also examine the effect of taking an algebra course on math achievement scores from the Early Childhood Longitudinal Study, a multilevel observational educational study, using our methods. The proposed methods are available in the `CURobustML` R package.

Keywords: Causal inference, Machine learning methods, Unmeasured variables, Omitted variable bias, Fixed effects models

1 Motivation: Machine Learning for Treatment Effects in Multilevel Observational Studies with Cluster-Level Unmeasured Confounding

There has been a growing interest in causal inference to estimate treatment effects using machine learning (ML) methods (Athey & Imbens, 2016; Carvalho, Feller, Murray, Woody, & Yeager, 2019; Chernozhukov et al., 2018; Dorie, Hill, Shalit, Scott, & Cervone, 2019; Hill, 2011; Künzel,

^{*}eub6uw@virginia.edu

[†]hyunseung@stat.wisc.edu

This article has been accepted for publication in *Psychometrika*, published by Springer.

Sekhon, Bickel, & Yu, 2019; McCaffrey, Ridgeway, & Morral, 2004; Su, Tsai, Wang, Nickerson, & Li, 2009; Suk, Kang, & Kim, 2020; van der Laan, Polley, & Hubbard, 2007; van der Laan & Rose, 2011; Wager & Athey, 2018; Westreich, Lessler, & Funk, 2010). An attractive feature of ML-based methods is that they flexibly estimate the propensity score model or the outcome regression model, reducing the risk for model mis-specification. However, many of these methods rely on the assumption of no unmeasured confounding where the treatment assignment is conditionally ignorable, i.e., as-if random conditional on observed pre-treatment confounders. If unmeasured confounders are present and treatment is no longer conditionally ignorable, these methods will likely lead to inconsistent estimates of the treatment effect. Although it is generally impossible to eliminate bias from any unmeasured confounder, the overarching goal of this paper is to propose ML methods that are robust to a specific type of unmeasured confounders called cluster-level unmeasured confounders while retaining the ML methods’ robustness against model mis-specification.

Broadly speaking, cluster-level confounders are confounders that are shared by study units within a cluster. They differ from individual-level confounders which are confounders that vary across individuals. Cluster-level confounders are often encountered in studies where units are naturally clustered together, most notably in multilevel observational studies. For example, consider the Early Childhood Longitudinal Study-Kindergarten cohort (ECLS-K) study, a national longitudinal study about school performance and student experience where students, the study units, were clustered by school. One of the research questions from the ECLS-K study was to estimate the effect of students taking an algebra or higher-level math course on their math achievement scores. Potential cluster-level confounders include school funding and school’s emphasis on mathematics; these factors are not only shared across every student in the same school, but also affect students’ propensity to take an algebra course and their achievement scores. In contrast, potential individual-level confounders include students’ prior achievement scores in mathematics and parents’ educational levels. In particular, students with higher prior achievement scores in math likely find math courses enjoyable, leading them to take more math courses as well as do well on future math tests. Similarly, parents’ educational background may affect courses that students take as well as students’ overall academic performance.

If an investigator were able to measure all the individual-level and cluster-level confounders in a study, they can use one of the aforementioned ML-based estimators to estimate the treatment effect and be guarded against biases from model mis-specification. But when some confounders are not measured, these estimators may be biased. Specifically, going back to the algebra class example, unfortunately, ECLS-K was not able to precisely measure cluster-level confounders such as school funding for math instruction and school’s emphasis on science, technology, engineering, and mathematics (STEM). With these cluster-level variables unmeasured, using popular ML methods in causal inference, such as Bayesian additive regression trees (BART) and targeted maximum likelihood estimators (TMLE) with ensemble learning, may lead to biased estimates of the desired treatment effect even if they can flexibly estimate the propensity score or the outcome model.

2 Prior Work and Our Contribution

Most prior works on estimating treatment effects under cluster-level unmeasured confounding are based on estimating parametric models for the outcome or the treatment with fixed or random effects (Arpino & Mealli, 2011; Firebaugh, Warner, & Massoglia, 2013; Kim & Frees, 2006; Li, Zaslavsky, & Landrum, 2013; Li, Lee, Port, & Robinson, 2020). However, due to their parametric nature, these methods run the risk of introducing bias from mis-specifying either the outcome model or the propensity score model. For example, methods based on random effects assume that the net effect from cluster-level unmeasured confounders on the outcome is Normally distributed

with mean zero and constant variance; they also assume that measured confounders are not correlated with unmeasured confounders (Raudenbush & Bryk, 2002; Wooldridge, 2010). Methods based on fixed effects do not typically assume Normality, but require that the aforementioned net effects are linear; in fact, the linearity assumption is critical to allow a procedure called “demeaning” where every observed variable is centered to their within-cluster means (Wooldridge, 2010, 2012). Recent works by Yang (2018), He (2018), and Lee, Nguyen, and Stuart (2019) use methods based on propensity scores, but they require parametric modeling assumptions and, similar to the outcome-based methods, may lead to biased estimates if the propensity score model is mis-specified. Arkhangelsky and Imbens (2019) also use a propensity-score-based approach and assume that the conditional distributions of the covariates and the treatment are in the exponential family. Finally, Zetterqvist, Vansteelandt, Pawitan, and Sjölander (2016) propose a doubly robust estimator that is consistent for the treatment effect so long as either the outcome model or the propensity score is correctly specified, but not necessarily both. However, their method, called the doubly robust conditional generalized estimating equation (DRCGEE), still requires a parametric model for the propensity score.

Also, many prior works above have focused on two-level studies where there is one level of grouping; the ECLS-K study is an example of a two-level study where the students (i.e., the study units) are grouped at the school level. But most multilevel observational studies have complex grouping or hierarchical structures, and unmeasured cluster-level confounders may be present at different levels of the hierarchy. For example, in some educational studies, students, the study units, are clustered/nested at the classroom level, and classrooms are clustered/nested at the school level, leading to “three-level” studies. There may be different cluster-level confounders at each level, such as characteristics specific to classrooms (second level) or to schools (third level). In cross-classified studies, students may belong to two types of non-nested clusters simultaneously, say school and neighborhood, and cluster-level confounders could be related to either school or neighborhood (Raudenbush & Bryk, 2002). To the best of our knowledge, there has been little work on dealing with unmeasured cluster-level confounders in three-level, cross-classified, or other complex, potentially non-nested multilevel observational studies.

Our contribution is to provide ML-based methods for estimating treatment effects in multilevel observational studies with potential unmeasured cluster-level confounders. Compared to prior works that relied on parametric models to deal with cluster-level unmeasured confounding, our ML-based methods do not require such models; instead, we use flexible ML methods to estimate the propensity score or the outcome model and consequently, there is less potential for bias from model mis-specification. Also, our methods can be applied to two-level, three-level, k -level, cross-classified, or a large family of multilevel observational studies where cluster-level unmeasured confounders can be present at each level. We evaluate our proposed methods in a simulation study with different multilevel structures—two-level and cross-classified studies—and show that our methods estimate the treatment effects well in the presence of unmeasured cluster-level confounders. In particular, when the outcome or the propensity score model is complex, our methods are better than the DRCGEE estimator, which is robust to cluster-level confounding, but requires a parametric model. In addition, our methods outperform BART and TMLE, which can flexibly model the propensity score or the outcome model, but cannot handle cluster-level unmeasured confounders. We conclude by estimating the effect of taking an eighth-grade algebra course on math achievement scores in the ECLS-K study using our methods. We also conduct a sensitivity analysis to examine whether our conclusions about the estimated effect from the ECLS-K study would be changed by individual-level unmeasured confounders.

3 Setup

3.1 Notation, Definitions, and Potential Outcomes

Suppose there are n individuals in the study. Let $j = 1, \dots, J$ be an index for J groups/clusters of units and let $i = 1, \dots, n_j$ be individual units within group j . In two-level studies, j would denote clusters and i would denote individuals within a cluster. For three-level studies, j would index all the clusters and i would denote individuals within a cluster; specifically, we can expand the subscript j by using double subscripts, say jk where j indexes clusters at the second level and k indexes clusters at the third level. For cross-classified studies with non-nested cluster-level groups, j would denote clusters determined jointly by the two non-nested groups, say school and neighborhood, and i would denote individuals within a cluster; we can also expand the subscript j to jk where j denotes one group (e.g., schools) and k denotes another group (e.g., neighborhoods). To simplify notation, we suppress these double or possibly l -level subscripts.

Let $\mathbf{X}_{ij}$ denote p_x observed pre-treatment confounders for individual i in cluster j . Pre-treatment confounders can include both individual-level confounders, such as student's prior achievement scores, and cluster-level confounders, such as whether the school was public or private. To distinguish the two types of confounders, we let $\mathbf{X}_{ij}^{(2)}$ denote cluster-level measured confounders where they satisfy $\mathbf{X}_{ij}^{(2)} = \mathbf{X}_{i'j}^{(2)} = \mathbf{X}_j^{(2)}$ for every $i \neq i'$; we drop the subscript i to emphasize that $\mathbf{X}_j^{(2)}$ are cluster-level confounders. We then denote $\mathbf{X}_{ij}^{(1)}$ to be any measured confounders that exclude cluster-level confounders, or more concisely, individual-level confounders. This allows $\mathbf{X}_{ij}$ to be partitioned into cluster-level and individual-level confounders, i.e., $\mathbf{X}_{ij} = (\mathbf{X}_{ij}^{(1)}, \mathbf{X}_j^{(2)})$. Similarly, let $\mathbf{U}_{ij}$ denote unmeasured confounders for individual i in cluster j where we partition the unmeasured confounders $\mathbf{U}_{ij}$ into $\mathbf{U}_{ij} = (\mathbf{U}_{ij}^{(1)}, \mathbf{U}_j^{(2)})$; here, $\mathbf{U}_{ij}^{(1)}$ denote individual-level unmeasured confounders and $\mathbf{U}_j^{(2)}$ denote cluster-level unmeasured confounders. Let $Z_{ij} \in \{0, 1\}$ denote the treatment assignment where $Z_{ij} = 1$ indicates that individual i in cluster j was treated and $Z_{ij} = 0$ indicates that individual i in cluster j was untreated. Let Y_{ij} denote the observed outcome for individual i in cluster j .

We use the potential outcomes notation (Neyman, 1923; Rubin, 1974) and its extension to multilevel/clustered data by Hong and Raudenbush (2006) to define causal effects. Let $Y_{ij}(1)$ be the potential treatment outcome if individual i in cluster j were treated ($Z_{ij} = 1$) and let $Y_{ij}(0)$ be the potential control outcome if they were untreated ($Z_{ij} = 0$). We remark that the notation implicitly assumes the Stable Unit Treatment Value Assumption (SUTVA; Rubin, 1986) where $Y_{ij} = Z_{ij}Y_{ij}(1) + (1 - Z_{ij})Y_{ij}(0)$; see Hong and Raudenbush (2006, 2013) for discussion of SUTVA in multilevel studies.

3.2 Causal Estimands

The first causal estimand we define is the average treatment effect (ATE), which is the average linear contrast of potential outcomes between the treated and untreated groups and is denoted as τ :

$$\tau = E[Y_{ij}(1) - Y_{ij}(0)]. \quad (1)$$

In the ECLS-K study, τ would be the overall effect of taking an algebra course in eighth grade on math achievement scores. The second causal estimand is a type of conditional average treatment effect (CATE) known as the best linear projection of the CATE (Ding, Feller, & Miratrix, 2019; Semenova & Chernozhukov, 2020). To motivate this estimand, consider the usual CATE which measures the ATE among a subgroup of individuals defined by p_v pre-treatment covariates $\mathbf{V}_{ij} \in \mathbb{R}^{p_v}$, i.e.,

$$\tau(\mathbf{v}) = E[Y_{ij}(1) - Y_{ij}(0) \mid \mathbf{V}_{ij} = \mathbf{v}]. \quad (2)$$

In the ECLS-K study, if V_{ij} is defined by sex so that $V_{ij} = 1$ indicates females and $V_{ij} = 0$ indicates males, $\tau(1)$ is the ATE of taking an algebra course on math achievement scores among females and $\tau(0)$ is the ATE among males. Or if $\mathbf{V}_{ij}$ is defined by baseline achievement scores, i.e., $V_{ij} = X_{ij}$, $\tau(80)$ is the ATE among students who received an 80 on their baseline scores and $\tau(85)$ is the ATE among students who received an 85 on their baseline scores. But unlike the case where V_{ij} is defined by sex, a bounded discrete variable, the case where V_{ij} is defined by achievement scores requires estimating a continuous function. Specifically, if we want to estimate the ATE among students with a baseline score of 83, i.e., $\tau(83)$, ideally, our data should contain students whose baseline scores are exactly 83. Otherwise, we would have to interpolate (or extrapolate) from students with baseline scores close to 83, say 80 or 85 and the quality of the interpolation depends on the assumptions about the functional form of $\tau(v)$.

To rely less on such functional assumptions, Semenova and Chernozhukov (2020) proposed to estimate the best linear projection of the CATE, or formally:

$$\beta = \operatorname{argmin}_{\beta'} E \left[(\tau(\mathbf{X}_{ij}, \mathbf{U}_{ij}) - \mathbf{V}_{ij}^T \beta')^2 \right]. \quad (3)$$

Briefly, the parameter $\mathbf{V}_{ij}^T \beta$ can be interpreted as the best linear approximation of a potentially complex function $\tau(\mathbf{X}_{ij}, \mathbf{U}_{ij})$; see Figure 1 for an illustration. In particular, even if the true function $\tau(\mathbf{X}_{ij}, \mathbf{U}_{ij})$ is non-linear, involving quadratic or cubic terms, β can still be interpreted as the linear effect of changing $\mathbf{V}_{ij}$ on increasing or decreasing the treatment effect. Finally, β collapses to the “usual” CATE in (2) if $\mathbf{V}_{ij}$ only contains discrete variables. For example, in the above example where $\tau(v)$ was defined by sex, β_0 would equal to the ATE among males, i.e., $\beta_0 = \tau(0)$, and β_1 would equal to the ATE among females, i.e., $\beta_1 = \tau(1)$. Also, if V_{ij} only contains a constant of 1s, β collapses to the ATE τ . Here, we slightly abuse the definition of $\mathbf{V}_{ij}$ by allowing a constant of 1 to be part of $\mathbf{V}_{ij}$. While this constant can be thought of as a “null confounder” where everyone has the same confounding value 1, its inclusion is primarily meant to help re-parametrize $\tau(\mathbf{v})$ and interpret the coefficient β based on usual linear regression models involving an intercept term. In subsequent sections, we will focus on estimating both the ATE and β .

Figure 1: Illustration of $\beta = (\beta_0, \beta_1)$ defined in (3) with one continuous confounder X_{ij} , no unmeasured confounder, and $V_{ij} = X_{ij}$. The gray lines indicate the true functional form of $\tau(X_{ij})$ and the red line represents the best linear projection defined in (3).

3.3 Assumptions

If the unmeasured confounders $\mathbf{U}_{ij}$ were observed, the usual set of the working assumptions to identify the ATE and the CATE is:

- (A1) Conditional Ignorability: $Y_{ij}(1), Y_{ij}(0) \perp Z_{ij} \mid \mathbf{X}_{ij}, \mathbf{U}_{ij}$

(A2) Overlap: $0 < e(\mathbf{X}_{ij}, \mathbf{U}_{ij}) = P(Z_{ij} = 1 \mid \mathbf{X}_{ij}, \mathbf{U}_{ij}) < 1$

Briefly, (A1) states that within every value of the confounders $\mathbf{X}_{ij}, \mathbf{U}_{ij}$, the treatment assignment Z_{ij} is randomly assigned to individuals and thus, is independent of $Y_{ij}(1)$ and $Y_{ij}(0)$. (A2) states that for every value of the confounders, the probability of getting treated or untreated, called the propensity score and denoted as $e(\mathbf{X}_{ij}, \mathbf{U}_{ij})$, is between zero and one. In particular, under the true propensity score, (A2) says that there are no clusters with only treated units or only control units. But (A2) may not hold in finite samples, especially if the sample size is generally small. (A1) and (A2) are jointly referred to as *strong ignorability* (Rosenbaum & Rubin, 1983), and under strong ignorability, the ATE or the CATE can be identified; see Hernan and Robins (2020), Schafer and Kang (2008), Austin (2011), and Steiner and Cook (2013) for details.

Unfortunately, the above identification strategy requires observing unmeasured confounders $\mathbf{U}_{ij}$ and without any restrictions on $\mathbf{U}_{ij}$, estimation of the ATE or the CATE is generally impossible. As such, a common identifying restriction on $\mathbf{U}_{ij}$ in multilevel observational studies is that all unmeasured confounders are cluster-level confounders or formally

(A3) Cluster-Level Unmeasured Confounding: $\mathbf{U}_{ij} = \mathbf{U}_j^{(2)}$

Under (A1)-(A3), identification of the ATE is achieved by comparing the differences in the outcomes between the treated and untreated units in the same cluster and aggregating these differences across clusters; a similar strategy also works for the CATE. Intuitively, this identification strategy works because within each cluster j , every study unit's unobserved confounders are identical under assumption (A3). In other words, the treated and untreated units in cluster j do not differ with respect to their unmeasured confounders $\mathbf{U}_j$. Then, by assumptions (A1) and (A2) where the treatment is as-if randomly assigned among units with identical confounder variables, the difference in the outcomes between treated and untreated units within cluster j will be unbiased for the treatment effect. While this identification strategy is well-known (see Imai and Kim (2019) for a graphical argument and Arkhangelsky and Imbens (2019) for an argument based on potential outcomes), few works have explored nonparametric estimators of treatment effects, especially with flexible, ML-based based methods. Subsequent sections will focus on estimation, specifically how to utilize assumption (A1)-(A3) and ML algorithms to obtain robust estimators of the ATE and the CATE.

4 Estimation Under Cluster-Level Unmeasured Confounding

4.1 Approach 1: Proxy Regression Estimator

The first approach is motivated by the doubly robust estimator for the ATE where we replace the propensity score and the outcome models that depend on $\mathbf{U}_j^{(2)}$ with observable proxies. Formally, suppose (A1)-(A3) hold and let $m_1(\mathbf{X}_{ij}, \mathbf{U}_j^{(2)}) = E[Y_{ij} \mid Z_{ij} = 1, \mathbf{X}_{ij}, \mathbf{U}_j^{(2)}]$ and $m_0(\mathbf{X}_{ij}, \mathbf{U}_j^{(2)}) = E[Y_{ij} \mid Z_{ij} = 0, \mathbf{X}_{ij}, \mathbf{U}_j^{(2)}]$ be the mean outcomes among treated and control units, respectively. If $\mathbf{U}_j^{(2)}$ were observed, a doubly robust estimator of the ATE (Bang & Robins, 2005; Glynn & Quinn, 2010) is:

$$\hat{\tau} = \frac{1}{n} \sum_{ij} \hat{\tau}_{ij},$$

$$\hat{\tau}_{ij} = \frac{Y_{ij}(Z_{ij} - m_1(\mathbf{X}_{ij}, \mathbf{U}_j^{(2)}))}{e(\mathbf{X}_{ij}, \mathbf{U}_j^{(2)})} + m_1(\mathbf{X}_{ij}, \mathbf{U}_j^{(2)}) - \frac{Y_{ij}(1 - Z_{ij} - m_0(\mathbf{X}_{ij}, \mathbf{U}_j^{(2)}))}{1 - e(\mathbf{X}_{ij}, \mathbf{U}_j^{(2)})} - m_0(\mathbf{X}_{ij}, \mathbf{U}_j^{(2)}). \quad (4)$$

In particular, the estimator $\hat{\tau}$ is consistent for the ATE as long as the propensity score model or the outcome model is correctly specified, but not necessarily both. Additionally, following Semenova and Chernozhukov (2020), we can use the doubly robust estimator to estimate the best linear approximation coefficients β by running a linear regression of $\hat{\tau}_{ij}$ on $\mathbf{V}_{ij}$:

$$\hat{\beta} = \operatorname{argmin}_{\beta'} \sum_{ij} \left(\hat{\tau}_{ij} - \mathbf{V}_{ij}^\top \beta' \right)^2. \quad (5)$$

Unfortunately, as written, the estimators (4) and (5) are not useful in our setting since they depend on unmeasured cluster-level confounders $\mathbf{U}_j^{(2)}$. To resolve this, we replace the regression functions m_1 , m_0 , and e that depend on unmeasured confounders $\mathbf{U}_j^{(2)}$ with *proxy regressions* (PR) $\tilde{m}_1$, $\tilde{m}_0$ and $\tilde{e}$ that do not depend on unmeasured confounders and the resulting estimators of the treatment effects are referred to as PR estimators. Specifically, for each individual i in cluster j , let $\mathbf{D}_{ij} \in \{0, 1\}^J$ be a J dimensional binary vector variable where the j th element is 1 and the rest are zero; $\mathbf{D}_{ij}$ can also be interpreted as the i th row of an n by J dummy matrix indicating individual i 's cluster membership in one of the J clusters. Under (A3), cluster-level unmeasured confounders $\mathbf{U}_j^{(2)}$ can be written as a function of $\mathbf{D}_{ij}$ as follows:

$$\mathbf{U}_j^{(2)} = f_{\mathbf{\Gamma}}(\mathbf{D}_{ij}) = \mathbf{D}_{ij} \mathbf{\Gamma}, \quad \mathbf{D}_{ij} \in \mathbb{R}^{1 \times J}, \quad \mathbf{\Gamma} \in \mathbb{R}^{J \times p_u}. \quad (6)$$

Here, p_u is the dimension of the cluster-level unmeasured confounders and f is a linear function parametrized by a matrix of unknown coefficients $\mathbf{\Gamma}$. Using equation (6), we can rewrite m_1 , m_0 , and e in terms of $\mathbf{D}_{ij}$ and $\mathbf{\Gamma}$:

$$m_1(\mathbf{X}_{ij}, \mathbf{U}_j^{(2)}) = E[Y_{ij} \mid Z_{ij} = 1, \mathbf{X}_{ij}, f_{\mathbf{\Gamma}}(\mathbf{D}_{ij})] = \tilde{m}_1(\mathbf{X}_{ij}, \mathbf{D}_{ij}, \mathbf{\Gamma}), \quad (7a)$$

$$m_0(\mathbf{X}_{ij}, \mathbf{U}_j^{(2)}) = E[Y_{ij} \mid Z_{ij} = 0, \mathbf{X}_{ij}, f_{\mathbf{\Gamma}}(\mathbf{D}_{ij})] = \tilde{m}_0(\mathbf{X}_{ij}, \mathbf{D}_{ij}, \mathbf{\Gamma}), \quad (7b)$$

$$e(\mathbf{X}_{ij}, \mathbf{U}_j^{(2)}) = P(Z_{ij} = 1 \mid \mathbf{X}_{ij}, f_{\mathbf{\Gamma}}(\mathbf{D}_{ij})) = \tilde{e}(\mathbf{X}_{ij}, \mathbf{D}_{ij}, \mathbf{\Gamma}). \quad (7c)$$

The functions $\tilde{m}_1$, $\tilde{m}_0$, and $\tilde{e}$ on the right-hand side of (7) are functions of the observed variables, i.e., $\mathbf{X}_{ij}$ and $\mathbf{D}_{ij}$, and $\mathbf{\Gamma}$ can be thought of as an unknown parameter that parameterizes these functions. In particular, so long as we use a flexible ML-based method to estimate these proxy regressions (see Section 3.4 for details), we can approximate the term $\mathbf{D}_{ij} \mathbf{\Gamma}$ in equation (6) and consequently, approximate the original regressions m_1 , m_0 , and e .

To illustrate this idea more concretely, consider a simple example where there are two clusters $J = 2$, say two schools, and there is a single unmeasured cluster-level covariate $U_j^{(2)}$, say a summary score measuring school funding that ranges from 1 to 10. Let the first school's funding score be 10, and let the second school's funding score be 5. Then, the function f that links the unmeasured cluster-level confounder $\mathbf{U}_j^{(2)}$ to the observed cluster membership indicator $\mathbf{D}_{ij}$ is $U_j^{(2)} = \mathbf{D}_{ij} \mathbf{\Gamma}$ where $\mathbf{\Gamma} = (10, 5)^\top$. Now, suppose the true outcome model $m_1(U_j^{(2)})$ is linear, i.e., $m_1(U_j^{(2)}) = \alpha_0 + \alpha_1 U_j^{(2)}$. From (6), we can rewrite $m_1(U_j^{(2)})$ as $m_1(U_j^{(2)}) = \alpha_0 + \mathbf{D}_{ij} \mathbf{\Gamma} \alpha_1$. Thus, as long as we use a linear regression model for the proxy regression $\tilde{m}_1(\mathbf{D}_{ij}, \mathbf{\Gamma})$, we can estimate α_0 and $\mathbf{\Gamma} \alpha_1$, which leads to an approximation of $m_1(U_j^{(2)})$ with $\tilde{m}_1(\mathbf{D}_{ij}, \mathbf{\Gamma})$.

Using the proxy regression functions, we can rewrite the estimator in equation (4) as:

$$\begin{aligned} \tilde{\tau} &= \frac{1}{n} \sum_{ij} \tilde{\tau}_{ij}, \\ \tilde{\tau}_{ij} &= \frac{Y_{ij}(Z_{ij} - \tilde{m}_1(\mathbf{X}_{ij}, \mathbf{D}_{ij}, \mathbf{\Gamma}))}{\tilde{e}(\mathbf{X}_{ij}, \mathbf{D}_{ij}, \mathbf{\Gamma})} + \tilde{m}_1(\mathbf{X}_{ij}, \mathbf{D}_{ij}, \mathbf{\Gamma}) - \frac{Y_{ij}(1 - Z_{ij} - \tilde{m}_0(\mathbf{X}_{ij}, \mathbf{D}_{ij}, \mathbf{\Gamma}))}{1 - \tilde{e}(\mathbf{X}_{ij}, \mathbf{D}_{ij}, \mathbf{\Gamma})} - \tilde{m}_0(\mathbf{X}_{ij}, \mathbf{D}_{ij}, \mathbf{\Gamma}). \end{aligned} \quad (8)$$

Similarly, to estimate β , we replace $\hat{\tau}_{ij}$ in (5) with $\tilde{\tau}_{ij}$:

$$\tilde{\beta} = \operatorname{argmin}_{\beta'} \sum_{ij} \left(\tilde{\tau}_{ij} - \mathbf{V}_{ij}^\top \beta \right)^2. \quad (9)$$

For standard errors, following Semenova and Chernozhukov (2020), we use the standard errors from ordinary least squares (OLS) regression where we treat $\tilde{\tau}_{ij}$ as the outcome variable and $\mathbf{V}_{ij}$ as the predictors and extract the OLS standard errors for the regression coefficients β associated with $\mathbf{V}_{ij}$.

4.2 Approach 2: Double Demeaning Estimator

The second approach is based on modifying the DRCGEE estimator (Zetterqvist et al., 2016) and before we present our modification, we briefly review the DRCGEE estimator. Consider the following models for the outcome and the treatment.

$$E[Y_{ij} \mid Z_{ij}, \mathbf{X}_{ij}, \mathbf{U}_j^{(2)}] = g_x(\mathbf{X}_{ij}) + g_u(\mathbf{U}_j^{(2)}) + Z_{ij} \mathbf{V}_{ij}^\top \beta, \quad (10)$$

$$E[Z_{ij} \mid \mathbf{X}_{ij}, \mathbf{U}_j^{(2)}] = h_x(\mathbf{X}_{ij}) + h_u(\mathbf{U}_j^{(2)}). \quad (11)$$

Here, g_x, h_x and g_u, h_u are some functions of the measured and unmeasured confounders, respectively. Zetterqvist et al. (2016) assumed that g_x and h_x have parametric specifications; specifically, if g_x and h_x are parametric, (10) and (11) are identical to the identity link models in equations (4.1), (4.2), and (4.4) of Zetterqvist et al. (2016). Once g_x and h_x are known, β can be estimated using an estimating equation of the form

$$\mathbf{0} = \sum_{ij} \mathbf{V}_{ij} (Z_{ij}^* - h_x^*(\mathbf{X}_{ij})) \cdot \left(Y_{ij}^* - g_x^*(\mathbf{X}_{ij}) - (Z_{ij} \mathbf{V}_{ij})^*{}^\top \hat{\beta} \right) \quad (12)$$

where $Y_{ij}^* = Y_{ij} - n_j^{-1} \sum_{i=1}^{n_j} Y_{ij}$, $Z_{ij}^* = Z_{ij} - n_j^{-1} \sum_{i=1}^{n_j} Z_{ij}$, $(Z_{ij} \mathbf{V}_{ij})^* = Z_{ij} \mathbf{V}_{ij} - n_j^{-1} \sum_{i=1}^{n_j} Z_{ij} \mathbf{V}_{ij}$, $h_x^*(\mathbf{X}_{ij}) = h_x(\mathbf{X}_{ij}) - n_j^{-1} \sum_{i=1}^{n_j} h_x(\mathbf{X}_{ij})$, and $g_x^*(\mathbf{X}_{ij}) = g_x(\mathbf{X}_{ij}) - n_j^{-1} \sum_{i=1}^{n_j} g_x(\mathbf{X}_{ij})$. Zetterqvist et al. (2016) showed that the expectation of this estimating equation is zero so long as either h_x^* or g_x^* is correctly specified; see Appendix S3 of Zetterqvist et al. (2016). Also, some matrix algebra reveals that a closed-form solution to equation (12) is:

$$\begin{aligned} \hat{\beta} &= (\mathbf{Q}^\top \mathbf{Q})^{-1} \mathbf{Q}^\top \mathbf{R}, \\ \mathbf{Q} &= \sum_{ij} \mathbf{V}_{ij} (Z_{ij}^* - h_x^*(\mathbf{X}_{ij})) (Z_{ij} \mathbf{V}_{ij})^*{}^\top, \quad \mathbf{Q} \in \mathbb{R}^{p_v \times p_v}, \quad \mathbf{V}_{ij} \in \mathbb{R}^{p_v \times 1}, \\ \mathbf{R} &= \sum_{ij} \mathbf{V}_{ij} (Z_{ij}^* - h_x^*(\mathbf{X}_{ij})) (Y_{ij}^* - g_x^*(\mathbf{X}_{ij})), \quad \mathbf{R} \in \mathbb{R}^{p_v \times 1}. \end{aligned}$$

While not noted in Zetterqvist et al. (2016), we remark that another way to obtain $\hat{\beta}$ is by running an instrumental variables (IV) regression based on generalized method of moments (GMM) (Hansen, 1982) where the instrumental variables are $\mathbf{V}_{ij} (Z_{ij}^* - h_x^*(\mathbf{X}_{ij}))$, the endogeneous variables are $(Z_{ij} \mathbf{V}_{ij})^*$, the outcome variable is $Y_{ij}^* - g_x^*(\mathbf{X}_{ij})$, and the weighting matrix is a $p_v \times p_v$ identity matrix. This connection to IV and GMM allows us to obtain standard errors using existing standard error formulas in IV and GMM; see Chapters 5.2 and 14 in Wooldridge (2010) and Kim and Frees (2007) for additional details.

Our main modification to the DRCGEE estimator is proposing to estimate g_x and h_x non-parametrically. Specifically, we borrow ideas from the works by Henderson, Carroll, and Li

(2008), Sun, Carroll, and Li (2009) and Lin, Li, and Sun (2014) on kernel regression in panel data and propose to solve the following optimization problem:

$$\begin{aligned}\hat{g}_x &= \operatorname{argmin}_{g_x, \beta'} \sum_{ij} \left(Y_{ij}^* - g_x(\mathbf{X}_{ij}) - n_j^{-1} \sum_{j=1}^{n_j} g_x(\mathbf{X}_{ij}) - (Z_{ij} \mathbf{V}_{ij})^* \beta' \right)^2, \\ \hat{h}_x &= \operatorname{argmin}_{h_x} \sum_{ij} \left(Z_{ij}^* - h_x(\mathbf{X}_{ij}) - n_j^{-1} \sum_{j=1}^{n_j} h_x(\mathbf{X}_{ij}) \right)^2.\end{aligned}\tag{13}$$

Like Sun et al. (2009) and Lin et al. (2014), there is an implicit assumption on g_u and h_u , usually that they are mean-zero so that g_x and h_x are identified; this is the nonparametric equivalent to the constraints on fixed effects in linear models where the individual fixed effects must sum to zero when the overall effect term is present (Henderson et al., 2008). Also, while Sun et al. (2009) and Lin et al. (2014) directly estimated g_x and h_x through kernel regressions, our proposal in (13) takes a different approach by proposing a loss function for ML algorithms so that they are not only robust to cluster-level unmeasured confounding, but the investigator can choose their own ML learners for g_x and h_x beyond kernel regression, say those based on trees, support vector machines, or neural networks. Once we have nonparametric estimates of $\hat{g}_x$ and $\hat{h}_x$, we plug it into the estimating equation in (12) to compute $\hat{\beta}$.

Now, suppose that an investigator cannot modify the loss function inside an existing ML method; the investigator may only have access to existing ML algorithms that use squared loss function. In such cases, we propose the following heuristic to directly use these existing ML methods to estimate h_x and g_x . For estimating g_x , we use a supervised ML method where the outcome is Y_{ij}^* and the predictors are $\mathbf{X}_{ij}^*$, $\bar{\mathbf{X}}_j$, and Z_{ij}^* . For estimating h_x , we also use a supervised ML method where the outcome is Z_{ij}^* and the predictors are $\mathbf{X}_{ij}^*$ and $\bar{\mathbf{X}}_j$. Then, we compute h_x^* and g_x^* from h_x and g_x . We call this approach the *double demeaning (DD) approach* and the resulting estimators for the treatment effects as DD estimators; the naming comes from how we de-mean twice in the estimation procedure, first for the outcome variables of each of the supervised ML methods (i.e., Y_{ij}^* and Z_{ij}^*) and second for the predictions from these methods to estimate h_x^* and g_x^* . Also, while there are inherent limitations with heuristics, we found in our simulation study that it works well in practice, especially if an investigator is unable to modify loss functions in ML algorithms.

4.3 Approach 3: Double Demeaning Estimator with Proxy Regression

One major downside of the double demeaning approach is that it assumes models for the outcome and the treatment where the unmeasured cluster-level confounders $\mathbf{U}_j^{(2)}$ are separable from the individual-level confounders $\mathbf{X}_{ij}$. To remedy this, our third approach combines the two approaches as follows. Consider the following generalization of the models in equations (10)-(11):

$$E[Y_{ij} \mid Z_{ij}, \mathbf{X}_{ij}, \mathbf{U}_j^{(2)}] = g(\mathbf{X}_{ij}, \mathbf{U}_j^{(2)}) + Z_{ij} \mathbf{V}_{ij}^\top \beta, \tag{14}$$

$$E[Z_{ij} \mid \mathbf{X}_{ij}, \mathbf{U}_j^{(2)}] = h(\mathbf{X}_{ij}, \mathbf{U}_j^{(2)}). \tag{15}$$

The difference between models (10)-(11) and models (14)-(15) is that we no longer assume that the functions concerning the confounders $\mathbf{X}_{ij}, \mathbf{U}_j^{(2)}$ are separable. Now, if g and h are known, we can directly use (12) to estimate β where g_x and h_x are replaced by g and h , respectively. To estimate g and h , we use the proxy regressions in Section 4.1. Specifically, for g , we use the proxy regression $\tilde{m}_0(\mathbf{X}_{ij}, \mathbf{D}_{ij}, \Gamma)$ and for h , we use the proxy regression $\tilde{e}(\mathbf{X}_{ij}, \mathbf{D}_{ij}, \Gamma)$. We

remark that Evdokimov (2010) proposed another approach to estimating h and g from models that are similar to equations (14) and (15). Once we have the estimated g and h , we use the DD estimator’s formulation based on instrumental variables regression to obtain standard errors.

4.4 Ensemble Machine Learning

All three methods in Sections 4.1, 4.2, and 4.3 require estimating nuisance functions. Specifically, for the PR estimator, we need to estimate proxy regressions for the treatment $\tilde{e}$ and the outcome $\tilde{m}_1, \tilde{m}_0$. For the DD estimator, we need to estimate the functions g_x and h_x . For the DDPR estimator, we need to estimate the functions g and h . As mentioned in Section 4.2, Zetterqvist et al. (2016) used parametric models to estimate g_x and h_x , which could lead to biased estimates if the parametric models are mis-specified. Instead, in this section, we detail how to use ML methods to estimate these nuisance functions. We use $\tilde{m}_0$ as an example in this section, but the same principle applies for all the other functions; see Appendix A where we detail the estimation procedure behind every nuisance function in the PR estimator, the DD estimator, and the DDPR estimator.

Broadly speaking, our estimator of $\tilde{m}_0$ is based on an existing ensemble machine learning algorithm called SuperLearner, which combines predictions from different supervised learning methods (Polley & van der Laan, 2010; van der Laan et al., 2007; van der Laan & Rose, 2011). Briefly, a series of works by van der Laan et al. (2007) and Polley and van der Laan (2010) have shown that under some assumptions, an ensemble estimator of nuisance functions like $\tilde{m}_0$ will never be worse (in terms of asymptotic risk) than individual estimators of these functions, even if one of the individual estimator is incorrect/inconsistent. In other words, if the true outcome model $\tilde{m}_0$ is not based on a linear model with fixed effects, but instead is based on a more complex functional form that can be modeled by more advanced ML methods like a neural network, the ensemble estimator of the outcome model that takes a weighted combination of linear regression with fixed effects and a neural network will consistently estimate the true outcome model. Also, if the true outcome model is based on a linear regression with fixed effects, the ensemble estimator will still consistently estimate the true model. Finally, if the true outcome model cannot be modeled by either a linear model with fixed effects or a neural network, the ensemble estimator will have risk that is at least as small as the risk based on using either one alone.

In our setup, we use two prediction methods, one from a generalized linear model used in prior works and another from a neural network, to obtain our ensemble estimator of $\tilde{m}_0$; we remark that one can include additional ML learners like random forests and gradient boosting models to further improve prediction. The generalized linear model is implemented with fixed cluster effects (Arpino & Mealli, 2011; Li et al., 2013; Wooldridge, 2010). A neural network uses a rectifier activation function with 2 hidden layers and each layer has $J + p_x + 150$ neurons with L1 regularization of 10^{-8} and L2 regularization of 10^{-8} ; see the supplementary materials and the CURobustML package in the first author’s GitHub repository for more details. We aggregate the predictions from the generalized linear model and the neural network using non-negative least squares. Algorithm 1 summarizes these steps. We remark that one can also incorporate sample splitting and cross fitting (Chernozhukov et al., 2018) where a subset of the sample is used to estimate the nuisance functions and the rest is used to estimate the treatment effects.

For software, we use the `h2o` package (LeDell et al., 2020) and `lme4` package (Bates, Mächler, Bolker, & Walker, 2015) in R (R Core Team, 2020) to fit the neural network and the generalized linear regression model, respectively. We use the `nnls` package (Mullen & van Stokkum, 2012) for non-negative least squares. All of our estimators are available in the R package CURobustML.

Algorithm 1 Ensemble Machine Learning for $\tilde{m}_0$

Input: Outcome: Y_{ij}

Input: Predictors: $Z_{ij}, \mathbf{X}_{ij}, \mathbf{D}_{ij}$

- 1: Predict the outcome based on linear regression with Y_{ij} as the outcome and $Z_{ij}, \mathbf{X}_{ij}, \mathbf{D}_{ij}$ as the predictors. Denote the predictions under $Z_{ij} = 0$ as $\hat{Y}_{ij}^{\text{GLM}}$.
- 2: Predict the outcome based on a neural network. Denote the predictions under $Z_{ij} = 0$ as $\hat{Y}_{ij}^{\text{NN}}$.
- 3: Combine predictions from different ML methods using non-negative least squares.

$$\begin{aligned}\hat{\zeta}_1, \hat{\zeta}_2 &= \underset{\zeta_1, \zeta_2}{\operatorname{argmin}} \sum_{ij} \left(Y_{ij} - \zeta_1 \hat{Y}_{ij}^{\text{GLM}} - \zeta_2 \hat{Y}_{ij}^{\text{NN}} \right)^2 \\ 1 &= \zeta_1 + \zeta_2, \quad \zeta_k \geq 0 \\ \hat{Y}_{ij} &= \hat{\zeta}_1 \hat{Y}_{ij}^{\text{GLM}} + \hat{\zeta}_2 \hat{Y}_{ij}^{\text{NN}}\end{aligned}$$

Output: Predictions from $\tilde{m}_0$, i.e., $\hat{Y}_{ij}$.

5 Simulation Study

We conducted a large-scale simulation study to assess the performance of our three approaches. Our simulation study consists of four designs. Design 1 is based on a two-level structure with large cluster sizes and there is no cross-level interaction between a measured individual-level confounder and an unmeasured cluster-level confounder. Design 1 is similar to prior works by Arpino and Cannas (2016) and Zetterqvist et al. (2016) and is meant to measure the performance of the proposed estimators under existing simulation designs. Design 2 extends Design 1 by allowing cross-level interactions and is intended to demonstrate the flexibility of our methods to handle complex cluster-level confounders. Design 3 is a variation on Design 1 where we examine the performance of the proposed methods when the cluster size is small and provides a more closer examination of finite-sample properties of our estimators. Design 4 uses a cross-classified structure to demonstrate the robustness of our methods under different types of clustering structures.

Throughout the simulation, we estimate the ATE τ and the CATE β where $\mathbf{V}_{ij}$ is a two-dimensional vector with the first element being a constant term and the second element being a measured confounder. Also, the component of the parameter β associated with the constant term is referred to as the “main” treatment effect, and the component of the parameter β associated with the measured confounder is referred to as the “interaction” treatment effect. For convenience, we will denote the first and second components of β as β_1 (i.e., main effect) and β_2 (i.e., interaction effect).

We repeat our simulation 500 ($m=1, \dots, 500$) times and we measure the absolute bias, the standard deviation (SD), and the root mean squared errors (RMSE), i.e.,

$$\begin{aligned}|\text{Bias}|_{\theta} &= \frac{1}{500} \left| \sum_{m=1}^{500} (\hat{\theta}^{(m)} - \theta) \right| \\ \text{SD}_{\theta} &= \sqrt{\frac{1}{500-1} \sum_{m=1}^{500} (\hat{\theta}^{(m)} - \bar{\hat{\theta}})^2}\end{aligned}$$

$$\text{RMSE}_\theta = \sqrt{\frac{1}{500} \sum_{m=1}^{500} (\hat{\theta}^{(m)} - \theta)^2}$$

Here, θ represents either τ or one element of β . Also, $\hat{\theta}^{(m)}$ is the m -th estimate of the parameter θ , and $\bar{\hat{\theta}}$ is the sample mean of $\hat{\theta}^{(m)}$ across 500 replications. R codes for our simulation study are available in the supplementary materials and can also be found at the first author's GitHub repository.

5.1 Design 1: Two-level Data, No Cross-level Interaction, and Large Cluster Size

5.1.1 Model and Main Result

The data generating model for two-level data is based on those from Arpino and Cannas (2016) and Zetterqvist et al. (2016). For cluster structure, we have $j = 1, \dots, 50$ clusters, each with $n_j = 100$ individuals per cluster.

1. For each individual $i = 1, \dots, n_j$ in cluster j , generate three individual-level confounders, $\mathbf{X}_{ij}^{(1)} = (X_{1ij}^{(1)}, X_{2ij}^{(1)}, X_{3ij}^{(1)})$, a cluster-level measured confounder $X_j^{(2)}$, and a cluster-level unmeasured confounder $U_j^{(2)}$ as follows:

$$\begin{aligned} X_{1ij}^{(1)} &\sim \text{Unif}(-1, 1), & X_{2ij}^{(1)} &\sim N(0, 1), & X_{3ij}^{(1)} &\sim \text{Unif}(0, 1), \\ X_j^{(2)} &\sim N(0, 1), & U_j^{(2)} &\sim N(0, 1) \end{aligned}$$

2. Generate individual treatment status Z_{ij} from the following logistic propensity score model:

$$\begin{aligned} \text{logit}(e_{ij}) &= -0.3 + 0.3X_{1ij}^{(1)} + 0.3X_{2ij}^{(1)} + 0.3X_{3ij}^{(1)} + 0.3X_j^{(2)} + 0.3(X_{3ij}^{(1)})^2 - 0.3X_{1ij}^{(1)}(X_{2ij}^{(1)})^2 \\ &\quad + 0.3X_{2ij}^{(1)}I(0 < X_{2ij}^{(1)} < 1.5) + 0.3X_{1ij}^{(1)}I(X_{3ij}^{(1)} < 0.2) + 0.3U_j^{(2)} + 0.3I(U_j^{(2)} > 0), \\ Z_{ij} &\sim \text{Bernoulli}(e_{ij}) \end{aligned}$$

Here, e_{ij} is the propensity score for individual i in cluster j . We remark that the model includes a mixture of linear, quadratic, interaction, and threshold terms.

3. Generate the potential outcomes $Y_{ij}(1), Y_{ij}(0)$ and the observed outcome Y_{ij} from the following linear regression model:

$$\begin{aligned} Y_{ij}(z) &= 70 + \beta_1 z + \beta_2 z X_{3ij}^{(1)} + 2X_{1ij}^{(1)} + 2X_{2ij}^{(1)} + 2X_j^{(2)} + 2(X_{3ij}^{(1)})^2 - 2X_{1ij}^{(1)}(X_{2ij}^{(1)})^2 \\ &\quad + 2X_{2ij}^{(1)}I(0 < X_{2ij}^{(1)} < 1.5) + 2X_{1ij}^{(1)}I(X_{3ij}^{(1)} < 0.2) + 2U_j^{(2)} + 2I(U_j^{(2)} > 0) + \epsilon_{ij}, \\ Y_{ij} &= Z_{ij}Y_{ij}(1) + (1 - Z_{ij})Y_{ij}(0), \quad \epsilon_{ij} \sim N(0, 1) \end{aligned}$$

Here, ϵ_{ij} represents the random error for individual i in cluster j . Similar to the treatment model, the outcome model includes a mixture of linear, quadratic, interaction, and threshold terms. The main and interaction treatment effects are set to $(\beta_1, \beta_2) = (2, 2)$. The population ATE τ is $E(2 + 2X_{3ij}^{(1)}) = 3$.

Figure 2 summarizes the performance of our proposed estimators compared to the DRCGEE estimator and two existing ML methods, BART and TMLE. Here, the DRCGEE estimator was implemented without specifying some higher-order terms to mimic an investigator who made a reasonable attempt to specify the models and was able to specify one quadratic term correctly, but was unable to capture other higher-order terms, say the threshold terms; note that

the DRCGEE estimator used the cluster indicators to demean the treatment and the outcome variables. Our estimators outperformed the DRCGEE estimator with respect to bias, variance, and RMSE by flexibly fitting the outcome and the propensity score models. Also, our estimators outperformed existing ML methods, BART and TMLE, which can flexibly fit these models, but cannot handle cluster-level unmeasured confounders; note that BART and TMLE do not use cluster indicators. Combined, compared to existing estimators, our estimators are more robust to two sources of bias: (a) bias from model mis-specification and (b) bias from unmeasured cluster-level confounders.

Figure 2: Performance of Proposed Methods Under Design 1. DRCGEE denotes the doubly robust conditional generalized estimating equation estimator. BART denotes Bayesian additive regression trees, and TMLE denotes targeted maximum likelihood estimation. The dashed black line indicates the true effect size of 2 or 3.

In Appendix B, we provide additional simulation results for BART and TMLE that include cluster indicators as predictors. We observed that including an additional set of cluster dummies did not always help eliminate bias and RMSE, and BART and TMLE with cluster dummies still yielded higher bias and RMSE compared to our proposed estimators.

5.1.2 Additional Evaluation of the Proposed Estimators

In this section, we take a deeper look into the performance of our estimators to better understand how they produce robust estimates of the treatment effects. First, Table 1 provides a peek into the performance of the ensemble learners inside the PR, DD, and DDP estimators. Specifically, we replace the ensemble learners with the two individual predictions that make up the ensemble learner, a generalized linear model and a neural network. As discussed in Section 4.4, prior works have shown that under some assumptions, combining individual predictions using an ensemble learning algorithm are asymptotically no worse than those based on using the predictions separately. Our numerical experiments confirm this theoretical result where the three estimators based on ensemble learners performed as well as those based on a neural network, which was the best learner for this particular design. Also, the DD estimator with an ensemble learner had

lower variance and lower RMSE than those based on a single learner. Given the superiority of the ensemble learner, we will use ensemble learners for the rest of the simulation study.

Table 1: Performance of Proposed Methods Under Different Estimators of the Propensity Score Model and the Outcome Model in Design 1

Estimator		β_1			β_2			τ		
		Bias	SD	RMSE	Bias	SD	RMSE	Bias	SD	RMSE
PR	glm	0.496	0.165	0.522	0.034	0.281	0.283	0.479	0.083	0.486
	neural network	0.032	0.092	0.097	0.020	0.145	0.146	0.022	0.059	0.063
	ensemble learning	0.033	0.092	0.097	0.020	0.144	0.145	0.023	0.059	0.063
DD	glm	0.452	0.129	0.470	0.048	0.224	0.229	0.428	0.063	0.433
	neural network	0.013	0.136	0.137	0.013	0.243	0.243	0.007	0.062	0.062
	ensemble learning	0.014	0.122	0.123	0.012	0.216	0.216	0.008	0.058	0.058
DDPR	glm	0.441	0.129	0.459	0.037	0.222	0.225	0.422	0.063	0.427
	neural network	0.014	0.117	0.117	0.017	0.198	0.199	0.006	0.057	0.058
	ensemble learning	0.014	0.117	0.117	0.017	0.198	0.198	0.006	0.057	0.058

Note: glm denotes generalized linear models. The true values of the main effect β_1 and the interaction effect β_2 are 2, and the population ATE τ is 3.

Second, Table 2 examines the performance of the PR, DD, and DDPR estimators along with the DRCGEE estimator when we, despite our best efforts, mis-specified the outcome model or the propensity score. The table replicates the simulation experiment done by Zetterqvist et al. (2016) for their DRCGEE estimator to numerically assess the double robustness property of their estimators. For our setup, we explore four cases of model specification under Design 1: (1) the outcome and the propensity score models are correctly specified, (2) only the propensity score model is correctly specified, (3) only the outcome model is correctly specified, and (4) neither models are correctly specified. Here, a correctly specified model means that we fit a generalized linear model for either the treatment Z_{ij} or the outcome Y_{ij} exactly based on the above data generating model. An incorrectly specified model is a generalized linear model that includes main effects terms only and excludes the higher-order terms like the quadratic term, the interaction term, and the threshold terms; this model is likely the “default” model that most

Table 2: Performance of Estimators Under Different Specifications of the Propensity Score Model and the Outcome Model in Design 1

		β_1			β_2			τ		
	Estimator	Bias	SD	RMSE	Bias	SD	RMSE	Bias	SD	RMSE
Both models correct	PR	0.002	0.059	0.059	0.003	0.105	0.104	0.001	0.033	0.033
	DD	0.002	0.058	0.058	0.002	0.102	0.102	0.001	0.032	0.032
	DDPR	0.002	0.058	0.058	0.002	0.102	0.101	0.001	0.032	0.032
	DRCGEE	0.002	0.058	0.058	0.002	0.102	0.102	0.001	0.032	0.032
Outcome model incorrect	PR	0.001	0.128	0.128	0.011	0.242	0.242	0.004	0.044	0.044
	DD	0.011	0.104	0.104	0.019	0.200	0.201	0.001	0.032	0.032
	DDPR	0.003	0.103	0.103	0.007	0.198	0.198	0.001	0.032	0.032
	DRCGEE	0.011	0.104	0.104	0.019	0.200	0.201	0.001	0.032	0.032
Treatment model incorrect	PR	0.002	0.059	0.059	0.002	0.103	0.103	0.001	0.033	0.032
	DD	0.002	0.058	0.058	0.002	0.101	0.101	0.001	0.032	0.032
	DDPR	0.002	0.058	0.057	0.002	0.101	0.101	0.001	0.032	0.032
	DRCGEE	0.002	0.058	0.058	0.002	0.101	0.101	0.001	0.032	0.032
Both models incorrect	PR	0.555	0.152	0.576	0.088	0.263	0.277	0.512	0.074	0.517
	DD	0.521	0.133	0.538	0.115	0.227	0.254	0.464	0.064	0.468
	DDPR	0.517	0.132	0.533	0.115	0.226	0.253	0.459	0.064	0.464
	DRCGEE	0.521	0.133	0.538	0.115	0.227	0.254	0.464	0.064	0.468

Note: The true values of the main effect β_1 and the interaction effect β_2 are 2, and the population ATE τ is 3.

investigators use as part of their preliminary analysis of the data. As Table 2 suggests, when either the treatment model or the outcome model is correctly specified, our proposed estimators produced nearly unbiased estimates of β_v and τ , suggesting a double robust property; indeed, this is not surprising as all three of our proposed estimators were initially based on existing doubly robust estimators. Also, mis-specifying the outcome model produced estimates with higher variance compared to mis-specifying the treatment model, in line with prior results on double-robustness (Bang & Robins, 2005). Finally, when both the treatment and the outcome models are misspecified, all the estimators were biased. We remark that the results of the DD estimator were numerically equivalent to those from the DRCGEE estimator since we only used linear models to estimate h_x and g_x .

Third, in Appendix C, we numerically assess the convergence properties of our estimators by increasing the sample size and measuring the decrease in RMSE under different specifications of the propensity score model and the outcome model. In summary, the RMSE of our estimators decreases as the sample size increases, suggesting that our estimators are consistent under different specifications of the two models.

5.2 Design 2: Two-level Data, Cross-level Interaction, and Large Cluster Size

Design 2 extends Design 1 by allowing a cross-level interaction between a measured individual-level confounder $X_{3ij}^{(1)}$ and an unmeasured cluster-level confounder $U_j^{(2)}$ in both the propensity score model and the outcome model. Specifically, we changed the propensity score model and the outcome model to include the cross-interaction term $X_{3ij}^{(1)}I(U_j^{(2)} > 0)$.

$$\begin{aligned} \text{logit}(e_{ij}) &= -0.5 + 0.3X_{1ij}^{(1)} + 0.3X_{2ij}^{(1)} + 0.3X_{3ij}^{(1)} + 0.3X_j^{(2)} + 0.3X_{1ij}^{(1)}X_{2ij}^{(1)} + 0.3(X_{3ij}^{(1)})^2 \\ &\quad + 0.3X_{2ij}^{(1)}I(0 < X_{2ij}^{(1)} < 1.5) + 0.3X_{1ij}^{(1)}I(X_{3ij}^{(1)} < 0.2) + 0.3U_j^{(2)} + 0.3I(U_j^{(2)} > 0) \\ &\quad - 2X_{3ij}^{(1)}I(U_j^{(2)} > 0) \\ Y_{ij}(z) &= 70 + \beta_1 \cdot z + \beta_2 \cdot zX_{3ij}^{(1)} + 2X_{1ij}^{(1)} + 2X_{2ij}^{(1)} + 2X_j^{(2)} + 2(X_{3ij}^{(1)})^2 - 2X_{1ij}^{(1)}(X_{2ij}^{(1)})^2 \\ &\quad + 2X_{2ij}^{(1)}I(0 < X_{2ij}^{(1)} < 1.5) + 2X_{1ij}^{(1)}I(X_{3ij}^{(1)} < 0.2) + 2U_j^{(2)} + 2I(U_j^{(2)} > 0) \\ &\quad - 2X_{3ij}^{(1)}I(U_j^{(2)} > 0) + \epsilon_{ij} \end{aligned}$$

Figure 3 replicates Figure 2 under Design 2. As discussed in Section 4.2, the DD estimator can be biased in this setup due to the cross-level interactions and the non-separability of the unmeasured cluster-level confounder and the individual-level confounder. As such, we only use the other two estimators, the PR estimator and the DDPR estimator, which allow for cross-level interactions. Similar to Design 1, our proposed estimators outperformed the DRCGEE estimator and existing two ML methods, BART and TMLE, with respect to bias, variance, and RMSE. In short, our proposed estimators are robust to bias arising from not only model mis-specification but also different kinds of unmeasured cluster-level confounders based on cross-level interactions. In Appendix D.1, we replicate the additional evaluations done in Section 5.1.2 under Design 2 and the results are similar to those reached in Section 5.1.2.

5.3 Design 3: Two-level Data, No Cross-level Interaction, and Small Cluster Size

Design 3 is a variation on Design 1 where we explore the performance of our methods under small cluster size. Specifically, we use the same models as Design 1, but we use four sample size conditions of (J, n_j) : (100, 5), (100, 10), (100, 20), and (100, 30).

Figure 3: Performance of Proposed Methods Under Design 2. DRCGEE denotes the doubly robust conditional generalized estimating equation estimator. BART denotes Bayesian additive regression trees, and TMLE denotes targeted maximum likelihood estimation. The dashed black line indicates the true effect size of 2 or 3.

Figure 4: Performance of Proposed Methods Under Varying Cluster Size in Design 3

Figure 4 summarizes the performance of our proposed estimators with varying cluster sizes. To make the bias more comparable between sample size and to keep the scaling of the y-axis constant, we use percent bias ($= |\text{Bias}|/\theta \cdot 100$); Appendix D.2 provides different metrics to compare performance. As the cluster sizes increased from 5 to 30, the three proposed estimators produced smaller biases. Specifically, when the cluster size was 5, the estimates from PR had 5%-15% biases, the estimates from DD had 2%-15% biases, and the estimates from DDPR had 3%-7% biases. However, when the cluster size increased to 10, the estimates from DD and DDPR had less than 2.5% bias. When the cluster sizes were 30, all the estimates had less than 2.5% bias. Overall, based on the simulation result, we recommend that investigators use our proposed estimators when the cluster size is at least 10. Appendix D.2 conducts additional analysis where we vary the type of learner used in each of the three estimators and assess the finite-sample performance of our estimators under different learners.

5.4 Design 4: Cross-classified Data, No Cross-level Interaction, and Large Cluster Size

Our data generating models for cross-classified data follow Meyers and Beretvas (2006) and Suk et al. (2020). We also use the triple subscript notation to follow the traditional notation used in cross-classified studies.

1. For each factor-1 cluster $j = 1, \dots, 25$ (e.g. schools), generate the number of individuals (e.g., students) $n_j = 300$.
2. Create factor-2 cluster labels $k = 1, \dots, 25$ (e.g., neighborhoods) with two independent unobserved cluster-factor confounders (i.e., $U_j^{(2)}, U_k^{(2)}$). For every factor-1 cluster except the first and last clusters, there are three feeders where each factor-1 cluster sends 50% of its individuals to the most adjacent factor-2 cluster, 25% to the next closest factor-2 cluster, and the rest to the next factor-2 factor. For the first and last factor-1 clusters, there are two feeders where 60% of individuals in a cluster goes to the most adjacent factor-2 cluster, and the rest goes to the next closest factor-2 cluster; see Meyers and Beretvas (2006) for details. In the end, this will create roughly 100 study units within subgroups determined by factor-1 and factor-2 cluster labels.
3. For each individual $i = 1, \dots, n_{jk}$ in both factor-1 cluster j and factor-2 cluster k , generate individual-level confounders $\mathbf{X}_{ijk}^{(1)} = (X_{1ijk}^{(1)}, X_{2ijk}^{(1)}, X_{3ijk}^{(1)})$, factor-1 measured and unmeasured confounders $X_j^{(2)}, U_j^{(2)}$, and factor-2 measured and unmeasured confounders $X_k^{(2)}, U_k^{(2)}$ as follows:

$$\begin{aligned} X_{1ijk}^{(1)} &\sim \text{Unif}(-1, 1), & X_{2ijk}^{(1)} &\sim N(0, 1), & X_{3ijk}^{(1)} &\sim \text{Unif}(0, 1), \\ X_j^{(2)} &\sim N(0, 1), & U_j^{(2)} &\sim N(0, 1), & X_k^{(2)} &\sim N(0, 1), & U_k^{(2)} &\sim N(0, 1) \end{aligned}$$

4. Generate individual treatment status Z_{ijk} from the following logistic propensity score model:

$$\begin{aligned} \text{logit}(e_{ijk}) &= -0.5 + 0.3X_{1ijk}^{(1)} + 0.3X_{2ijk}^{(1)} + 0.3X_{3ijk}^{(1)} + 0.3X_j^{(2)} + 0.3X_k^{(2)} + 0.3(X_{3ijk}^{(1)})^2 \\ &\quad - 0.3X_{1ijk}^{(1)}(X_{2ijk}^{(1)})^2 + 0.3X_{2ijk}^{(1)}I(0 < X_{2ijk}^{(1)} < 1.5) + 0.3X_{1ijk}^{(1)}I(X_{3ijk}^{(1)} < 0.2) + 0.6U_j^{(2)} \\ &\quad + 0.6U_k^{(2)} + 0.6U_k^{(2)}I(U_j^{(2)} < 1), \\ Z_{ijk} &\sim \text{Bernoulli}(e_{ijk}) \end{aligned}$$

Here, e_{ijk} is the propensity score for individual i in factor-1 cluster j and factor-2 cluster k . We remark that the model includes a mixture of linear, quadratic, interaction, and threshold terms.

5. Generate the potential outcomes $Y_{ijk}(1), Y_{ijk}(0)$ and observed outcome Y_{ijk} from the following linear regression model:

$$\begin{aligned}
Y_{ijk}(z) = & 70 + \beta_1 \cdot z + \beta_2 \cdot z X_{3ijk}^{(1)} + 2X_{1ijk}^{(1)} + 2X_{2ijk}^{(1)} + 2X_{3ijk}^{(1)} + 2X_j^{(2)} + 2X_k^{(2)} + 2(X_{3ijk}^{(1)})^2 \\
& - 2X_{1ijk}^{(1)}(X_{2ijk}^{(1)})^2 + 2X_{2ijk}^{(1)}I(0 < X_{2ijk}^{(1)} < 1.5) + 2X_{1ijk}^{(1)}I(X_{3ijk}^{(1)} < 0.2) + 5U_j^{(2)} + 5U_k^{(2)} \\
& + 5U_k^{(2)}I(U_j^{(2)} < 1) + \epsilon_{ijk}, \\
Y_{ijk} = & Z_{ijk}Y_{ijk}(1) + (1 - Z_{ijk})Y_{ijk}(0), \quad \epsilon_{ijk} \sim N(0, 1)
\end{aligned}$$

Here, ϵ_{ijk} represents the random error for individual i in factor-1 cluster j and factor-2 cluster k . Similar to the treatment model, the outcome model includes a mixture of linear, quadratic, interaction, and threshold terms. The main and interaction effects are set to $(\beta_1, \beta_2) = (2, 2)$.

The above models allow for cluster-level unmeasured confounders that are not nested within each other. That is, a student, say Sally, who goes to school $j = 1$ and lives in neighborhood $k = 1$ can share an unmeasured school-level confounder, say $U_{j=1}^{(2)}$, with students who go to the same school as Sally, but live in different neighborhoods. Sally can also share another, unique unmeasured confounder, say $U_{k=1}^{(2)}$, with students who don't go to school $j = 1$, but live in the same neighborhood $k = 1$ as Sally. Sally can also share a different unmeasured confounder with students who go to the same school and live in the same neighborhood as Sally; this could be an interaction of $U_{j=1}^{(2)}$ and $U_{k=1}^{(2)}$ and we use the interaction term of the form $U_k^{(2)}I(U_j^{(2)} < 1)$ in our models above. We define the cluster indicators by the combination of two groups, say schools and neighborhoods.

If the investigator can assume that there are no unmeasured school-by-neighborhood confounders but there are unmeasured school-level confounders and unmeasured neighborhood-level confounders, we can either continue to use the cluster indicators as stated in the previous paragraph or use the school and neighborhood indicators separately to remove biases arising from unmeasured cluster-level confounding. The latter approach can lead to better finite sample performance as it incorporates an investigator's prior knowledge about the structure of cluster-level unmeasured confounding. But the approach may be biased if the investigator's assumption about the unmeasured confounding is incorrect.

Figure 5 summarizes the performance of the proposed estimators, the DRCGEE estimator and two other ML methods, BART and TMLE; we provide more detailed results of our estimators in Appendix D.3. Similar to the two-level setting, our proposed estimators outperformed the DRCGEE estimator, BART, and TMLE in terms of bias, variance, and RMSE except in two cases (SD of $\hat{\beta}_2$ from TMLE and SD of $\hat{\tau}$ from DRCGEE). In particular, the DRCGEE estimator, on average, yielded the largest bias for β_1 and τ . BART and TMLE, which can flexibly estimate the propensity score model and the outcome model, produced more extreme estimates and resulted in high variance compared to our proposed estimators. Our proposed estimators had the lowest bias across existing methods, showing that they are robust against both model-specification and cluster-robust unmeasured confounders in complex multilevel observational studies.

In Appendix B, we conducted additional simulations for BART and TMLE with two different sets of cluster indicators as predictors: (i) school-by-neighborhood cluster indicators and (ii) school and neighborhood indicator variables. Overall, we found that using indicator variables based on (i) or (ii) in BART or TMLE generally did not remove more bias than the ones presented here that do not include these indicator variables.

Figure 5: Performance of Proposed Methods Under Design 4. DRCGEE denotes the doubly robust conditional generalized estimating equation estimator. BART denotes Bayesian additive regression trees, and TMLE denotes targeted maximum likelihood estimation. The dashed black line indicates the true effect size of 2 or 3.

6 Real Data Study

We demonstrated our proposed methods by estimating the effects of taking an algebra or a higher level math course on math achievement scores from ECLS-K. We compared the results of our methods with those from the DRCGEE estimator and two pre-existing ML methods, BART and TMLE. We conclude by conducting a sensitivity analysis when assumption (A3) is violated and an unmeasured individual-level confounder is present.

6.1 Data and Methods

ECLS-K is a national longitudinal study about school achievement and student experience and is sponsored by the National Center for Education Statistics (NCES) within the Institute of Education Sciences (IES). The study recruited a nationally representative sample of kindergarteners during the Fall of 1998 using a dual-frame multistage sampling design (Walston & McCarroll, 2010) and followed them through the spring of 2007 when most of the students were in eighth grade. For more details about the ECLS-K study, see the ECLS-K website .

Using the ECLS-K data, we examined the effect of students taking an algebra or a higher level math course in eighth grade on their math achievement scores. Specifically, for each student (i.e., study unit) in a particular school (i.e., cluster) during the spring of 2007, we define a binary treatment variable with 1 indicating that the student took the algebra or a higher-level math course and 0 otherwise (i.e., general mathematics or introduction to algebra/pre-algebra). We define the outcome as the student’s math achievement scores (C7R4MSCL), ranging from 67.75 to 172.20 (mean 145.78, standard deviation 19.73). Following Walston and McCarroll (2010) and Rickles (2013), we use 11 confounders that are thought to influence the treatment and outcome variables. Eight of the eleven confounders are student-level confounders (i.e., individual-level

confounders) and the rest are school-level confounders (i.e., cluster-level confounders). Student-level confounders included prior math achievement scores (C6R4MSCL), socio-economic status measure (WKSESL), poverty level (WKPOV_R), mother’s educational level (WKMOMED), expected degree of child (P6EXPECT), family type (P7HFAMIL; living with one parent or two parents), gender (GENDER), and race (RACE). School-level confounders included public or private school (S7PUPRI), school location (R7URBAN), and region (R7REGION). We used data from individuals with no missing treatment and outcome values. Also, we imputed any missing values in the confounders with predictive mean matching (White, Royston, & Wood, 2011) implemented via the `mice` package (van Buuren & Groothuis-Oudshoorn, 2011).

We further restricted the sample by excluding students who answered “don’t know” in P6EXPECT and answered “other” in P7HFAMIL because they made up a small proportion in the study population (1.58% and 1.63%, respectively). We also only included schools with at least 10 students with at least three treated and control students. At the end, our analytic sample consists of $n = 2496$ students from $J = 164$ schools. We remark that our results may not generalize to the entire population of eighth graders in the U.S. and that our empirical analysis is designed to assess different methods’ ability to remove cluster-level confounding.

We ran the proposed methods using the `CURobustML` R package developed for this paper. We compare our estimates to those from the DRCGEE estimator implemented in the `drgee` R package (Zetterqvist & Sjölander, 2015; Zetterqvist et al., 2016). Both our methods and the DRCGEE estimator are robust to cluster-level unmeasured confounders, but our methods are also robust to potential mis-specification of the propensity score or the outcome regression by using ensemble learning. In addition, we computed the effect estimates from BART and TMLE via the `bartCause` R package (Dorie & Hill, 2019) and `tmle` R package (Gruber & van der Laan, 2012), respectively. BART and TMLE reduce concerns from misspecifying the propensity score or the outcome regression, but they are not robust to unmeasured cluster-level confounders. Data and R codes for analysis are available in the supplementary materials.

Finally, we conducted a sensitivity analysis to evaluate the sensitivity of our estimates when an individual-level unmeasured confounder is present. Specifically, we used a sensitivity analysis method from Hong and Raudenbush (2006) based on a hypothetical individual-level unmeasured confounder and reported the plausible confidence interval when such a confounder is present.

6.2 Results

Table 3 summarizes the covariate balance of the measured student-level and school-level confounders; see Appendix E for a visual representation of covariate balance. Except for gender and race, the means of all confounders differed significantly between the treatment and the control group at the $\alpha = 0.05$ level. In particular, every school-level characteristic differed between the treatment group (i.e., students who took an algebra or higher-level math course) and the control group (i.e., students who did not take an algebra class or higher-level math course) where students in the control group were more likely to go to public schools in small towns from the South or the Midwest. There may also be unmeasured cluster-level confounders, such as school’s emphasis on math education and the number of math teachers teaching algebra courses. These may play an important role in students’ decision to take math courses, and not accounting for them may induce bias in the estimates of treatment effects. Thankfully, our proposed methods are robust against biases arising from cluster-level unmeasured confounders.

Table 4 reports the ATE using different estimators. The *prima facie* effect is the unadjusted mean difference in students’ math achievement scores when students took the algebra or higher-level math class and was 13.44 points. This effect was reduced to 2.19, 3.02 and 3.15 points after applying our PR, DD, and DDPR estimators, respectively. The ATE estimate from the DRCGEE estimator was 2.79, and it was in between the estimates based on our proposed methods. BART and TMLE, which are not robust to cluster-level unmeasured confounders, also produced

Table 3: Descriptive Statistics of Confounders Between the Control and Treated Groups

	Control Group (N = 1311)		Treated Group (N = 1185)	
	Percent	Std Dev	Percent	Std Dev
<i>Student-level Variables</i>				
Prior Score	121.41	21.85	137.41	20.19
SES	-0.01	0.72	0.31	0.73
Poverty level: Below	12.5%		8.7%	
Mom Edu: High school	33.3%		23.3%	
Mom Edu: College or voca.	34.4%		33.2%	
Mom Edu: Bachelor’s or higher	23.4%		38.7%	
Expected Degree: 4-year college	49.2%		57.8%	
Expected Degree: More than	14.8%		25.3%	
Family Type: Two parents	83.0%		87.4%	
Gender: Male	50.3%		50.6%	
Race: White	76.0%		79.8%	
Race: Hispanic	12.8%		13.0%	
<i>School-level Variables</i>				
School Type: Public	89.7%		84.5%	
Location: Suburb	39.1%		43.9%	
Location: Small town	43.4%		38.1%	
Region: Northeast	20.9%		25.7%	
Region: South	23.3%		18.5%	
Region: Midwest	42.6%		40.8%	

Note: Values in bold are significant differences between two groups at $\alpha=0.05$.

estimates that were between the estimates from our proposed estimators.

Table 4: ATE Estimates of Taking an Algebra Class on Achievement Scores

	Estimate	Standard Error
Prima facie (unadjusted)	13.44	0.74
PR	2.19	0.77
DD	3.02	0.47
DDPR	3.15	0.50
DRCGEE	2.79	0.55
BART	2.45	0.61
TMLE	2.38	0.52

Note: DRCGEE denotes the doubly robust conditional generalized estimating equation estimator. BART denotes Bayesian additive regression trees, and TMLE denotes targeted maximum likelihood estimation.

Next, we investigated effect heterogeneity, specifically the effect of taking an algebra class (or higher) on math achievement scores among white versus non-white students. Studying effect heterogeneity by race is motivated in part by works from Holloway (2004), Wenglinsky (2004), and Noguera and Wing (2008) that discussed the existing racial achievement gap in math and pointed out that narrowing the gap in math is critical to reducing the overall racial gap in learning. Table 5 summarizes the results. Across all of our estimators, we found that the effect among white students was higher than that from non-white students; the former effect was significantly positive, while the latter effect was not. But group differences from our proposed estimators (i.e., PR, DD, DDPR) were larger than those from the DRCGEE estimator, BART and TMLE.

Finally, we examined whether our conclusions about the ATE would be altered by bias arising from an individual-level unmeasured confounder $U_{ij}^{(1)}$. Specifically, for PR, DD, and DDPR, we followed Hong and Raudenbush (2006) and estimated a new treatment effect based on the formula

Table 5: CATE Estimates of Taking an Algebra Class on Achievement Scores: White vs. Non-White

	Non-White		White	
	Estimate	Std. Error	Estimate	Std. Error
PR	1.24	1.64	2.46	0.88
DD	1.85	1.01	3.35	0.53
DDPR	2.16	1.09	3.43	0.57
DRCGEE	2.25	1.51	2.94	0.59
BART	1.74	0.93	2.66	0.65
TMLE	2.00	1.05	2.49	0.59

Note: DRCGEE denotes the doubly robust conditional generalized estimating equation estimator. BART denotes Bayesian additive regression trees, and TMLE denotes targeted maximum likelihood estimation. Standard errors for TMLE were estimated using cluster bootstrap sampling with 2000 replicates.

from their work: $\tau^\dagger = \tau + \zeta^{(1)}(E[U_{ij}^{(1)} | Z_{ij} = 1] - E[U_{ij}^{(1)} | Z_{ij} = 0])$. Here, τ is taken to be the lower and upper 95% confidence intervals (CIs) of the ATE based on the PR, DD, and DDPR estimators; note that because the estimated lower CIs were strictly positive, the formula helps us gauge whether our conclusions would be changed by *negative* bias arising from individual-level unmeasured confounders. Following Hong and Raudenbush (2006), we used the value of 3.538 for $\zeta^{(1)}$, which was the second largest estimated coefficient of student-level covariates that was obtained from running a regression between the outcome and the confounders; we chose the second largest value because it is hard to expect that there is an unmeasured individual-level confounder that has the strongest impact on the outcome as prior achievement scores, which were measured in our observed data and are generally considered to be the most important predictor for post-treatment achievement scores. To specify $(E[U_{ij}^{(1)} | Z_{ij} = 1] - E[U_{ij}^{(1)} | Z_{ij} = 0])$, we also used the second largest treatment-minus-control difference that was obtained from running a regression that regresses each student-level covariate on treatment and all covariates, and it was 0.097. Thus, the plausible negative bias amounted to 0.344 ($= 3.538 \cdot 0.097$).

Using the above sensitivity analysis, our adjusted 95% CIs for the ATE are: (0.33, 3.36) for the PR estimator, (1.76, 3.59) for the DD estimator, (1.83, 3.78) for the DDPR estimator, and (1.36, 3.52) for DRCGEE. All the CIs did not contain 0. Given the results, we are reasonably confident that the treatment effect exists, is significant, and is robust to unmeasured individual-level confounding.

Similarly, we also conducted a sensitivity analysis for estimates from BART and TMLE using the same methodology in Hong and Raudenbush (2006). Specifically, we use their formula $\tau^\ddagger = \tau + \zeta^{(1)}(E[U_{ij}^{(1)} | Z_{ij} = 1] - E[U_{ij}^{(1)} | Z_{ij} = 0]) + \zeta^{(2)}(E[U_j^{(2)} | Z_{ij} = 1] - E[U_j^{(2)} | Z_{ij} = 0])$ that accounted for bias arising from both individual-level and cluster-level unmeasured confounders. We specified $\zeta^{(1)}$ and $(E[U_{ij}^{(1)} | Z_{ij} = 1] - E[U_{ij}^{(1)} | Z_{ij} = 0])$ as before. For specifying $\zeta^{(2)}$, we used the value of 4.159, which was the largest estimated coefficient among school-level covariates from running a regression between the outcome and the confounders. For $(E[U_j^{(2)} | Z_{ij} = 1] - E[U_j^{(2)} | Z_{ij} = 0])$, we used the same empirical value of 0.097. Using these values, we obtained a negative bias of 0.748 ($= 3.538 \cdot 0.097 + 4.159 \cdot 0.097$) and found the 95% CIs for the new ATE estimates to be (0.74, 3.07) for BART and (0.59, 2.68) for TMLE. The two CIs did not contained 0, and these two estimates are still insensitive to unmeasured confounders although they have to deal with both individual-level and cluster-level unmeasured confounders.

7 Discussion and Conclusions

This paper proposed three different approaches to estimate treatment effects in multilevel observational studies that are robust to both model mis-specification and cluster-level unmeasured confounders. We achieve this by using different approaches to account for cluster-level unmeasured confounding and using ensemble ML methods for modeling. We showed through a simulation study that our methods perform well in various multilevel observational studies, specifically under both two-level and cross-classified settings, and almost always outperform existing methods. Our methods have lower bias and variance than the DRCGEE estimator, which is robust to cluster-level unmeasured confounders, but requires parametric models. Also, our methods have lower bias and variance than BART and TMLE, which are robust to model mis-specification, but are not robust to cluster-level unmeasured confounders. In our data analysis, all estimators generally revealed that the overall effect of taking an algebra course on math achievement scores was positive, though each estimator resulted in slightly different point estimates due to differences in the underlying assumptions that support each methodology.

There are some limitations of the proposed methods. First, we assumed SUTVA in multilevel designs where there are no spillover effects through interference. Second, we did not establish whether confidence intervals from our estimators achieve nominal coverage; our preliminary numerical analysis suggests that the reported confidence intervals have the desired coverage and we plan to theoretically establish their coverage properties in a future work. Third, our study only used an ensemble of a generalized linear model and a neural network based on a few number of fine-tuned parameters. However, it is possible that adding more ML estimators or fine-tuning more parameters inside the neural network can improve the overall performance of our estimators. Despite these limitations, our proposed approaches provide estimates of treatment effects that are robust to cluster-level unmeasured confounders and model mis-specification and provide investigators with additional tools to robustly study causal effects in multilevel observational studies.

References

- Arkhangelsky, D., & Imbens, G. (2019). *The role of the propensity score in fixed effect models*. arXiv. Retrieved from <https://arxiv.org/abs/1807.02099> doi: 10.3386/w24814
- Arpino, B., & Cannas, M. (2016). Propensity score matching with clustered data. an application to the estimation of the impact of caesarean section on the apgar score. *Statistics in Medicine*, 35(12), 2074–2091. doi: doi.org/10.1002/sim.6880
- Arpino, B., & Mealli, F. (2011). The specification of the propensity score in multilevel observational studies. *Computational Statistics & Data Analysis*, 55(4), 1770–1780. doi: 10.1016/j.csda.2010.11.008
- Athey, S., & Imbens, G. (2016). Recursive partitioning for heterogeneous causal effects. *Proceedings of the National Academy of Sciences*, 113(27), 7353–7360. doi: 10.1073/pnas.1510489113
- Austin, P. C. (2011). An introduction to propensity score methods for reducing the effects of confounding in observational studies. *Multivariate Behavioral Research*, 46, 399–424. doi: 10.1080/00273171.2011.568786
- Bang, H., & Robins, J. M. (2005). Doubly robust estimation in missing data and causal inference models. *Biometrics*, 61(4), 962–973. doi: 10.1111/j.1541-0420.2005.00377.x
- Bates, D., Mächler, M., Bolker, B., & Walker, S. (2015). Fitting linear mixed-effects models using lme4. *Journal of Statistical Software*, 67(1), 1–48. doi: 10.18637/jss.v067.i01
- Carvalho, C., Feller, A., Murray, J., Woody, S., & Yeager, D. (2019). Assessing treatment effect variation in observational studies: Results from a data challenge. *Observational Studies*, 5, 21–35. doi: 10.1353/obs.2019.0000
- Chernozhukov, V., Chetverikov, D., Demirer, M., Duflo, E., Hansen, C., Newey, W., & Robins, J. (2018). Double/debiased machine learning for treatment and structural parameters. *The Econometrics Journal*, 21(1), C1–C68. doi: 10.1111/ectj.12097
- Ding, P., Feller, A., & Miratrix, L. (2019). Decomposing treatment effect variation. *Journal of the American Statistical Association*, 114(525), 304–317. doi: 10.1080/01621459.2017.1407322
- Dorie, V., & Hill, J. (2019). bartcause: Causal inference using bayesian additive regression trees [Computer software manual]. Retrieved from <https://github.com/vdorie/bartCause> (R package version 1.0-0)
- Dorie, V., Hill, J., Shalit, U., Scott, M., & Cervone, D. (2019, 02). Automated versus do-it-yourself methods for causal inference: Lessons learned from a data analysis competition. *Statist. Sci.*, 34(1), 43–68. doi: 10.1214/18-STS667
- Evdokimov, K. (2010). *Identification and estimation of a nonparametric panel data model with unobserved heterogeneity*. Working paper, Princeton University.
- Firebaugh, G., Warner, C., & Massoglia, M. (2013). Fixed effects, random effects, and hybrid models for causal analysis. In S. L. Morgan (Ed.), *Handbook of causal analysis for social research* (pp. 113–132). Springer. doi: 10.1007/978-94-007-6094-3_7
- Glynn, A. N., & Quinn, K. M. (2010). An introduction to the augmented inverse propensity weighted estimator. *Political Analysis*, 18(1), 36–56. doi: 10.1093/pan/mpp036
- Gruber, S., & van der Laan, M. J. (2012). tmle: An R package for targeted maximum likelihood estimation. *Journal of Statistical Software*, 51(13), 1–35. Retrieved from <http://www.jstatsoft.org/v51/i13/> doi: 10.18637/jss.v051.i13
- Hansen, L. P. (1982). Large sample properties of generalized method of moments estimators. *Econometrica: Journal of the Econometric Society*, 1029–1054. doi: 10.2307/1912775
- He, Z. (2018). *Inverse conditional probability weighting with clustered data in causal inference*. arXiv. Retrieved from <https://arxiv.org/abs/1808.01647>
- Henderson, D. J., Carroll, R. J., & Li, Q. (2008). Nonparametric estimation and testing of fixed effects panel data models. *Journal of Econometrics*, 144(1), 257–275. doi: 10.1016/

- j.jeconom.2008.01.005
- Hernan, M. A., & Robins, J. M. (2020). *Causal inference: What if*. Boca Raton: Chapman & Hall/CRC.
- Hill, J. L. (2011). Bayesian nonparametric modeling for causal inference. *Journal of Computational and Graphical Statistics*, 20(1), 217–240. doi: 10.1198/jcgs.2010.08162
- Holloway, J. H. (2004). Closing the minority achievement gap in math. *Educational Leadership*, 61(5), 84.
- Hong, G., & Hong, Y. (2009). Reading instruction time and homogeneous grouping in kindergarten: An application of marginal mean weighting through stratification. *Educational Evaluation and Policy Analysis*, 31(1), 54–81. doi: 10.3102/0162373708328259
- Hong, G., & Raudenbush, S. W. (2006). Evaluating kindergarten retention policy: A case study of causal inference for multilevel observational data. *Journal of the American Statistical Association*, 101, 901–910. doi: 10.1198/016214506000000447
- Hong, G., & Raudenbush, S. W. (2013). Heterogeneous agents, social interactions, and causal inference. In S. L. Morgan (Ed.), *Handbook of causal analysis for social research* (pp. 331–352). Springer. doi: 10.1007/978-94-007-6094-3_16
- Imai, K., & Kim, I. S. (2019). When should we use unit fixed effects regression models for causal inference with longitudinal data? *American Journal of Political Science*, 63(2), 467–490. doi: 10.1111/ajps.12417
- Kim, J.-S., & Frees, E. W. (2006). Omitted variables in multilevel models. *Psychometrika*, 71(4), 659. doi: 10.1007/s11336-005-1283-0
- Kim, J.-S., & Frees, E. W. (2007). Multilevel modeling with correlated effects. *Psychometrika*, 72(4), 505–533. doi: 10.1007/s11336-007-9008-1
- Künzel, S. R., Sekhon, J. S., Bickel, P. J., & Yu, B. (2019). Metalearners for estimating heterogeneous treatment effects using machine learning. *Proceedings of the National Academy of Sciences*, 116(10), 4156–4165. doi: 10.1073/pnas.1804597116
- LeDell, E., Gill, N., Aiello, S., Fu, A., Candell, A., Click, C., ... Malohlava, M. (2020). h2o: R interface for the 'h2o' scalable machine learning platform [Computer software manual]. Retrieved from <https://github.com/h2oai/h2o-3> (R package version 3.30.1.1)
- Lee, Y., Nguyen, T. Q., & Stuart, E. A. (2019). *Partially pooled propensity score models for average treatment effect estimation with multilevel data*. arXiv. Retrieved from <https://arxiv.org/abs/1910.05600>
- Li, F., Zaslavsky, A. M., & Landrum, M. B. (2013). Propensity score weighting with multilevel data. *Statistics in Medicine*, 32(19), 3373–3387. doi: 10.1002/sim.5786
- Li, Y., Lee, Y., Port, F. K., & Robinson, B. M. (2020). The impact of unmeasured within-and between-cluster confounding on the bias of effect estimators of a continuous exposure. *Statistical Methods in Medical Research*, 29(8), 2119–2139. doi: 10.1177/0962280219883323
- Lin, Z., Li, Q., & Sun, Y. (2014). A consistent nonparametric test of parametric regression functional form in fixed effects panel data models. *Journal of Econometrics*, 178, 167–179. doi: 10.1016/j.jeconom.2013.08.014
- McCaffrey, D. F., Ridgeway, G., & Morral, A. R. (2004). Propensity score estimation with boosted regression for evaluating causal effects in observational studies. *Psychological Methods*, 9(4), 403–425. doi: 10.1037/1082-989X.9.4.403
- Meyers, J. L., & Beretvas, S. N. (2006). The impact of inappropriate modeling of cross-classified data structures. *Multivariate Behavioral Research*, 41(4), 473–497. doi: 10.1207/s15327906mbr4104_3
- Mullen, K. M., & van Stokkum, I. H. M. (2012). nnls: The lawson-hanson algorithm for non-negative least squares (nnls) [Computer software manual]. Retrieved from <https://CRAN.R-project.org/package=nnls> (R package version 1.4)
- Neyman, J. S. (1923). On the application of probability theory to agricultural experiments: essay on principles. section 9 (with discussion). *Statistical Science*, 4, 465–480.

-
- Noguera, P. A., & Wing, J. Y. (2008). *Unfinished business: Closing the racial achievement gap in our schools*. John Wiley & Sons.
- Polley, E. C., & van der Laan, M. J. (2010). *Super learner in prediction*. U.C. Berkeley Division of Biostatistics Working Paper Series. Paper 226. doi: 10.1007/978-1-4419-9782-1_3
- R Core Team. (2020). R: A language and environment for statistical computing [Computer software manual]. Vienna, Austria. Retrieved from <https://www.R-project.org/>
- Raudenbush, S. W., & Bryk, A. S. (2002). *Hierarchical linear models: Applications and data analysis methods* (Vol. 1). Sage.
- Rickles, J. H. (2013). Examining heterogeneity in the effect of taking algebra in eighth grade. *The Journal of Educational Research*, 106(4), 251–268. doi: 10.1080/00220671.2012.692731
- Rosenbaum, P. R., & Rubin, D. B. (1983). The central role of the propensity score in observational studies for causal effects. *Biometrika*, 70, 41–55. doi: 10.1093/biomet/70.1.41
- Rubin, D. B. (1974). Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of Educational Psychology*, 66(5), 688–701. doi: 10.1080/01621459.1986.10478355
- Rubin, D. B. (1986). Comment: Which ifs have causal answers. *Journal of the American Statistical Association*, 81(396), 961–962. doi: 10.2307/2289065
- Rubin, D. B. (2001). Using propensity scores to help design observational studies: application to the tobacco litigation. *Health Services and Outcomes Research Methodology*, 2(3-4), 169–188.
- Schafer, J. L., & Kang, J. (2008). Average causal effects from nonrandomized studies: a practical guide and simulated example. *Psychological Methods*, 13(4), 279–313. doi: 10.1037/a0014268
- Semenova, V., & Chernozhukov, V. (2020). *Estimation and inference about conditional average treatment effect and other structural functions*. arXiv. Retrieved from <https://arxiv.org/abs/1702.06240>
- Shadish, W. R., Clark, M. H., & Steiner, P. M. (2008). Can nonrandomized experiments yield accurate answers? a randomized experiment comparing random and nonrandom assignments. *Journal of the American Statistical Association*, 103(484), 1334–1344. doi: 10.1198/016214508000000733
- Steiner, P. M., & Cook, D. (2013). Matching and propensity scores. In T. Little (Ed.), *The oxford handbook of quantitative methods* (p. 236–258). New York, NY: Oxford University Press. doi: 10.1093/oxfordhb/9780199934874.013.0013
- Steiner, P. M., Cook, T. D., Shadish, W. R., & Clark, M. H. (2010). The importance of covariate selection in controlling for selection bias in observational studies. *Psychological Methods*, 15(3), 250. doi: 10.1037/a0018719
- Su, X., Tsai, C.-L., Wang, H., Nickerson, D. M., & Li, B. (2009). Subgroup analysis via recursive partitioning. *Journal of Machine Learning Research*, 10(2), 141–158. doi: 10.2139/ssrn.1341380
- Suk, Y., Kang, H., & Kim, J.-S. (2020). Random forests approach for causal inference with clustered observational data. *Multivariate Behavioral Research*. doi: 10.1080/00273171.2020.1808437
- Sun, Y., Carroll, R. J., & Li, D. (2009). Semiparametric estimation of fixed-effects panel data varying coefficient models. In Q. Li & J. S. Racine (Eds.), *Nonparametric econometric methods* (pp. 101–130). Emerald Group Publishing Limited. doi: 10.1108/S0731-9053(2009)0000025006
- van Buuren, S., & Groothuis-Oudshoorn, K. (2011). mice: Multivariate imputation by chained equations in r. *Journal of Statistical Software*, 45(3), 1–67. doi: 10.18637/jss.v045.i03
- van der Laan, M. J., Polley, E. C., & Hubbard, A. E. (2007). Super learner. *Statistical Applications in Genetics and Molecular Biology*, 6(1). doi: 10.2202/1544-6115.1309
- van der Laan, M. J., & Rose, S. (2011). *Targeted learning: causal inference for observational*

-
- and experimental data. Springer Science & Business Media.
- Wager, S., & Athey, S. (2018). Estimation and inference of heterogeneous treatment effects using random forests. *Journal of the American Statistical Association*, 113(523), 1228–1242. doi: 10.1080/01621459.2017.1319839
- Walston, J., & McCarroll, J. C. (2010). Eighth-grade algebra: Findings from the eighth-grade round of the early childhood longitudinal study, kindergarten class of 1998-99 (ECLS-K). statistics in brief. nces 2010-016. *National Center for Education Statistics*.
- Wenglinsky, H. (2004). Closing the racial achievement gap: The role of reforming instructional practices. *Education Policy Analysis Archives*, 12, 64. doi: 10.14507/epaa.v12n64.2004
- Westreich, D., Lessler, J., & Funk, M. J. (2010). Propensity score estimation: neural networks, support vector machines, decision trees (cart), and meta-classifiers as alternatives to logistic regression. *Journal of Clinical Epidemiology*, 63(8), 826–833. doi: 10.1016/j.jclinepi.2009.11.020
- White, I. R., Royston, P., & Wood, A. M. (2011). Multiple imputation using chained equations: issues and guidance for practice. *Statistics in Medicine*, 30(4), 377–399. doi: 10.1002/sim.4067
- Wooldridge, J. M. (2010). *Econometric analysis of cross section and panel data*. The MIT press.
- Wooldridge, J. M. (2012). *Introductory econometrics: A modern approach*. South-Western Cengage Learning.
- Yang, S. (2018). Propensity score weighting for causal inference with clustered data. *Journal of Causal Inference*, 6(2). doi: 10.1515/jci-2017-0027
- Zetterqvist, J., & Sjölander, A. (2015). Doubly robust estimation with the R package drgee. *Epidemiologic Methods*, 4(1), 69–86. doi: 10.1515/em-2014-0021
- Zetterqvist, J., Vansteelandt, S., Pawitan, Y., & Sjölander, A. (2016). Doubly robust methods for handling confounding by cluster. *Biostatistics*, 17(2), 264–276. doi: 10.1093/biostatistics/kxv041

A Ensemble Learning Algorithms

We provide additional details on the learning algorithms used in our three approaches. The first and third approach, denoted as PR and DDPR in the main paper, use the same learning algorithms. Their learning algorithms, denoted as Algorithms 2 and 3, are detailed below.

Algorithm 2 Ensemble Machine Learning for $\tilde{m}_1$

Input: Outcome: Y_{ij}

Input: Predictors: $Z_{ij}, \mathbf{X}_{ij}, \mathbf{D}_{ij}$

- 1: Predict the outcome based on linear regression with Y_{ij} as the outcome and $Z_{ij}, \mathbf{X}_{ij}, \mathbf{D}_{ij}$ as the predictors. Denote the predictions under $Z_{ij} = 1$ as $\hat{Y}_{ij}^{\text{GLM}}$.
- 2: Predict outcome based on a neural network. Denote the predictions under $Z_{ij} = 1$ as $\hat{Y}_{ij}^{\text{NN}}$.
- 3: Combine predictions using non-negative least squares from different ML methods.

$$\begin{aligned}\hat{\zeta}_1, \hat{\zeta}_2 &= \underset{\zeta_1, \zeta_2}{\operatorname{argmin}} \sum_{ij} \left(Y_{ij} - \zeta_1 \hat{Y}_{ij}^{\text{GLM}} - \zeta_2 \hat{Y}_{ij}^{\text{NN}} \right)^2 \\ \zeta_1 + \zeta_2 &= 1, \quad \zeta_k \geq 0 \\ \hat{Y}_{ij} &= \hat{\zeta}_1 \hat{Y}_{ij}^{\text{GLM}} + \hat{\zeta}_2 \hat{Y}_{ij}^{\text{NN}}\end{aligned}$$

Output: Predictions from $\tilde{m}_1$, i.e., $\hat{Y}_{ij}$.

Algorithm 3 Ensemble Machine Learning for $\tilde{e}_1$

Input: Outcome: Z_{ij}

Input: Predictors: $\mathbf{X}_{ij}, \mathbf{D}_{ij}, \bar{\mathbf{X}}_j$

- 1: Predict the outcome based on linear regression with Z_{ij} as the outcome and $\mathbf{X}_{ij}, \mathbf{D}_{ij}$ as the predictors. Denote the predictions as $\hat{Z}_{ij}^{\text{GLM}}$.
- 2: Predict outcome based on a neural network. Denote the predictions as $\hat{Z}_{ij}^{\text{NN}}$.
- 3: Combine predictions using non-negative least squares from different ML methods.

$$\begin{aligned}\hat{\zeta}_1, \hat{\zeta}_2 &= \underset{\zeta_1, \zeta_2}{\operatorname{argmin}} \sum_{ij} \left(Z_{ij} - \zeta_1 \hat{Z}_{ij}^{\text{GLM}} - \zeta_2 \hat{Z}_{ij}^{\text{NN}} \right)^2 \\ \zeta_1 + \zeta_2 &= 1, \quad \zeta_k \geq 0 \\ \hat{Z}_{ij} &= \hat{\zeta}_1 \hat{Z}_{ij}^{\text{GLM}} + \hat{\zeta}_2 \hat{Z}_{ij}^{\text{NN}}\end{aligned}$$

Output: Predictions from $\tilde{e}_1$, i.e., $\hat{Z}_{ij}$.

For the second approach, denoted as DD in the main paper, we estimate g_x^* and h_x^* by using Algorithm 4 and Algorithm 5, respectively. Here, we implement a single generalized linear model because we use a heuristic suggested in the main text that includes both the demeaned variables $\mathbf{X}_{ij}^*$ and variable means $\bar{\mathbf{X}}_j$.

Algorithm 4 Ensemble Machine Learning for g_x^*

Input: Outcome: Y_{ij}^*

Input: Predictors: $Z_{ij}^*, \mathbf{X}_{ij}^*, \bar{\mathbf{X}}_j$

- 1: Predict the outcome based on linear regression with Y_{ij}^* as the outcome and $Z_{ij}^*, \mathbf{X}_{ij}^*, \bar{\mathbf{X}}_j$ as the predictors. Denote the predictions under $-\bar{Z}_j$ as $\hat{Y}_{ij}^{*GLM}$.
- 2: Predict outcome based on a neural network with Y_{ij}^* as the outcome and $Z_{ij}^*, \mathbf{X}_{ij}^*, \bar{\mathbf{X}}_j$ as the predictors. Denote the predictions under $-\bar{Z}_j$ as $\hat{Y}_{ij}^{*NN}$.
- 3: Combine predictions using non-negative least squares from different ML methods.

$$\begin{aligned} \hat{\zeta}_1, \hat{\zeta}_2 &= \operatorname{argmin}_{\zeta_1, \zeta_2} \sum_{ij} \left(Y_{ij} - \zeta_1 \hat{Y}_{ij}^{*GLM} - \zeta_2 \hat{Y}_{ij}^{*NN} \right) \\ \zeta_1 + \zeta_2 &= 1, \quad \zeta_k \geq 0 \\ \hat{Y}_{ij}^* &= \hat{\zeta}_1 \hat{Y}_{ij}^{*GLM} + \hat{\zeta}_2 \hat{Y}_{ij}^{*NN} \end{aligned}$$

Output: Predictions from g_x^* , i.e., $\hat{Y}_{ij}^*$.

Algorithm 5 Ensemble Machine Learning for h_x^*

Input: Outcome: Z_{ij}^*

Input: Predictors: $\mathbf{X}_{ij}^*, \bar{\mathbf{X}}_j$

- 1: Predict the outcome based on linear regression with Z_{ij}^* as the outcome and $\mathbf{X}_{ij}^*, \bar{\mathbf{X}}_j$ as the predictors. Denote the predictions as $\hat{Z}_{ij}^{*GLM}$.
- 2: Predict outcome based on a neural network with Z_{ij}^* as the outcome and $\mathbf{X}_{ij}^*, \bar{\mathbf{X}}_j$ as the predictors. Denote the predictions as $\hat{Z}_{ij}^{*NN}$.
- 3: Combine predictions using non-negative least squares from different ML methods.

$$\begin{aligned} \hat{\zeta}_1, \hat{\zeta}_2 &= \operatorname{argmin}_{\zeta_1, \zeta_2} \sum_{ij} \left(Z_{ij} - \zeta_1 \hat{Z}_{ij}^{*GLM} - \zeta_2 \hat{Z}_{ij}^{*NN} \right) \\ \zeta_1 + \zeta_2 &= 1, \quad \zeta_k \geq 0 \\ \hat{Z}_{ij}^* &= \hat{\zeta}_1 \hat{Z}_{ij}^{*GLM} + \hat{\zeta}_2 \hat{Z}_{ij}^{*NN} \end{aligned}$$

Output: Predictions from h_x^* , i.e., $\hat{Z}_{ij}^*$.

B Simulation Results for BART and TMLE with Cluster Dummies

In this section, we performed additional simulations with BART and TMLE under Designs 1 and 4. For Design 1, we included cluster dummies (denoted as id) as additional predictors into BART and TMLE. For Design 4, we included two different sets of cluster indicators into BART and TMLE: (i) factor-1-by-factor-2 dummies (denoted as f12id) and (ii) factor-1 dummies and factor-2 dummies (denoted as f1id and f12, respectively).

Figures 6 and 7 summarize the results under Design 1 and Design 4, respectively. From Figure 6, we observed that including an additional set of cluster dummies did not always help remove bias and RMSE for β , and both BART and TMLE with cluster dummies (i.e., BART+id, TMLE+id) still yielded higher bias and RMSE for β and τ , compared to our proposed estimators. Similarly, as shown in Figure 7, including either cluster indicators in BART or TMLE generally did not always remove additional biases from the effect estimates and their biases were still worse than our proposed estimators.

Figure 6: Performance of Bayesian Additive Regression Trees (BART) and Targeted Maximum Likelihood Estimation (TMLE) With and Without Cluster Dummies Under Design 1. id represents cluster dummies. The dashed black line indicates the true effect size of 2 or 3.

Figure 7: Performance of Bayesian Additive Regression Trees (BART) and Targeted Maximum Likelihood Estimation (TMLE) With and Without Cluster Dummies Under Design 4. f12id represents factor-1-by-factor-2 cluster dummies. f1id and f2id represent factor-1 dummies and factor-2 dummies, respectively. The label BART+f12id represents BART with cluster dummies f12id included as predictors. The dashed black line indicates the true effect size of 2 or 3.

C Convergence Properties of Proposed Methods Under Design 1

Figure 8 shows the RMSE of our proposed estimators under Design 1 when we increase the cluster size and vary the outcome and the propensity score model specifications. Overall, we see that the RMSE from all three estimators decrease as the cluster size increases so long as the outcome model or the propensity score model is correct, numerically suggesting the double robustness property of our proposed estimators.

Figure 8: RMSE of Proposed Methods With Increasing Cluster Sizes Under Design 1

D Simulation Results for Designs 2, 3, and 4

In the following subsections, we provide Tables 6, 7, 8, 9, 10, 11, 12, and 13. Tables 6, 8, 9, 10, 11, and 12 show the performance of our estimators with different estimators for the outcome model and the propensity score. Overall, the results confirm that using an ensemble algorithm generally performed the best compared to using the individual learner inside the ensemble algorithm separately.

Tables 7 and 13 replicate Table 2 for Design 2 and Design 4 in order to evaluate the robustness of our methods under different specifications of the propensity score and the outcome model if our estimators for them were unfortunately mis-specified. Overall, as long as we correctly specified either the treatment model or the outcome model, the effect estimates of β_v and τ from our proposed estimators had almost no bias. In contrast, when both the treatment and outcome models were mis-specified, all the estimators were biased. Also, we observed that mis-specifying the outcome model led to estimates with higher variance compared to mis-specifying the treatment model.

D.1 Design 2: Two-level Data, Cross-level Interaction, and Large Cluster Size

Table 6: Performance of Proposed Methods Under Different Estimators of the Propensity Score Model and the Outcome Model in Design 2

	Estimator	Bias	β_1	RMSE	Bias	β_2	RMSE	Bias	τ	RMSE
			SD			SD			SD	
PR	glm	0.490	0.158	0.515	0.041	0.274	0.277	0.470	0.078	0.477
	neural network	0.035	0.090	0.096	0.029	0.153	0.155	0.020	0.054	0.057
	ensemble learning	0.036	0.090	0.096	0.029	0.153	0.155	0.021	0.054	0.058
DDPR	glm	0.433	0.128	0.452	0.029	0.225	0.226	0.419	0.061	0.423
	neural network	0.015	0.110	0.111	0.018	0.197	0.198	0.006	0.055	0.056
	ensemble learning	0.015	0.110	0.111	0.018	0.197	0.198	0.006	0.055	0.056

Note: glm denotes generalized linear models. The true values of the main effect β_1 and the interaction effect β_2 are 2, and the population ATE, τ is 3.

Table 7: Performance of Estimators Under Different Specifications of the Propensity Score Model and the Outcome Model in Design 2

	Estimator	Bias	β_1	RMSE	Bias	β_2	RMSE	Bias	τ	RMSE
			SD			SD			SD	
Both models correct	PR	0.001	0.059	0.058	0.002	0.107	0.107	0.000	0.033	0.033
	DDPR	0.001	0.058	0.058	0.001	0.105	0.105	0.000	0.033	0.032
Outcome model incorrect	PR	0.000	0.128	0.128	0.009	0.256	0.256	0.004	0.047	0.047
	DDPR	0.002	0.102	0.101	0.004	0.201	0.200	0.000	0.033	0.033
Treatment model incorrect	PR	0.001	0.058	0.058	0.001	0.106	0.106	0.000	0.033	0.033
	DDPR	0.001	0.058	0.058	0.001	0.105	0.105	0.000	0.032	0.032
Both models incorrect	PR	0.552	0.148	0.571	0.100	0.265	0.283	0.502	0.072	0.507
	DDPR	0.512	0.133	0.529	0.114	0.230	0.256	0.456	0.063	0.460

Note: The true values of the main effect β_1 and the interaction effect β_2 are 2, and the population ATE, τ is 3.

D.2 Design 3: Two-level Data, No Cross-level Interaction, and Small Cluster Size

Table 8: Performance of Proposed Methods Under Different Estimators of the Propensity Score Model and the Outcome Model in Design 3: 100 Clusters with Cluster Sizes of 30

	Estimator	Bias	β_1	RMSE	Bias	β_2	RMSE	Bias	τ	RMSE
			SD			SD			SD	
PR	glm	0.429	0.197	0.472	0.039	0.348	0.350	0.411	0.101	0.423
	neural network	0.052	0.117	0.128	0.017	0.189	0.189	0.044	0.071	0.083
	ensemble learning	0.054	0.117	0.129	0.017	0.189	0.189	0.046	0.071	0.084
DD	glm	0.418	0.168	0.451	0.062	0.296	0.302	0.388	0.084	0.397
	neural network	0.026	0.289	0.290	0.022	0.538	0.538	0.015	0.131	0.132
	ensemble learning	0.033	0.200	0.203	0.024	0.366	0.367	0.022	0.096	0.099
DDPR	glm	0.404	0.168	0.437	0.044	0.297	0.299	0.382	0.084	0.391
	neural network	0.030	0.236	0.237	0.021	0.398	0.399	0.020	0.115	0.117
	ensemble learning	0.030	0.236	0.237	0.021	0.398	0.398	0.020	0.115	0.117

Note: glm denotes generalized linear models. The true values of the main effect β_1 and the interaction effect β_2 are 2, and the population ATE, τ is 3.

Table 9: Performance of Proposed Methods Under Different Estimators of the Propensity Score Model and the Outcome Model in Design 3: 100 Clusters with Cluster Sizes of 20

	Estimator	Bias	β_1	RMSE	Bias	β_2	RMSE	Bias	τ	RMSE
			SD			SD			SD	
PR	glm	0.403	0.232	0.464	0.059	0.409	0.413	0.373	0.115	0.391
	neural network	0.071	0.147	0.163	0.024	0.251	0.252	0.059	0.084	0.103
	ensemble learning	0.083	0.145	0.167	0.027	0.247	0.248	0.070	0.084	0.109
DD	glm	0.410	0.214	0.462	0.081	0.385	0.393	0.370	0.099	0.383
	neural network	0.033	0.527	0.528	0.052	1.004	1.004	0.007	0.238	0.238
	ensemble learning	0.047	0.300	0.304	0.051	0.522	0.524	0.022	0.140	0.142
DDPR	glm	0.387	0.216	0.443	0.050	0.389	0.392	0.362	0.100	0.376
	neural network	0.015	0.369	0.368	0.018	0.641	0.641	0.006	0.175	0.175
	ensemble learning	0.015	0.366	0.366	0.017	0.638	0.637	0.006	0.174	0.174

Note: glm denotes generalized linear models. The true values of the main effect β_1 and the interaction effect β_2 are 2, and the population ATE, τ is 3.

Table 10: Performance of Proposed Methods Under Different Estimators of the Propensity Score Model and the Outcome Model in Design 3: 100 Clusters with Cluster Sizes of 10

	Estimator	Bias	β_1			β_2			τ		
			SD	RMSE	Bias	SD	RMSE	Bias	SD	RMSE	
PR	glm	0.356	0.330	0.485	0.081	0.592	0.597	0.311	0.159	0.349	
	neural network	0.131	0.202	0.240	0.055	0.345	0.349	0.105	0.105	0.148	
	ensemble learning	0.137	0.199	0.242	0.055	0.346	0.350	0.111	0.105	0.152	
DD	glm	0.386	0.328	0.507	0.138	0.586	0.602	0.318	0.155	0.354	
	neural network	0.468	10.058	10.059	0.348	13.394	13.385	0.294	8.790	8.786	
	ensemble learning	0.013	0.639	0.639	0.035	1.094	1.093	0.031	0.307	0.308	
DDPR	glm	0.343	0.353	0.492	0.061	0.629	0.631	0.313	0.165	0.354	
	neural network	0.043	0.486	0.488	0.055	0.835	0.836	0.016	0.228	0.228	
	ensemble learning	0.041	0.486	0.487	0.053	0.834	0.835	0.015	0.226	0.227	

Note: glm denotes generalized linear models. The true values of the main effect β_1 and the interaction effect β_2 are 2, and the population ATE, τ is 3.

Table 11: Performance of Proposed Methods Under Different Estimators of the Propensity Score Model and the Outcome Model in Design 3: 100 Clusters with Cluster Sizes of 5

			β_1			β_2			τ		
	Estimator	Bias	SD	RMSE	Bias	SD	RMSE	Bias	SD	RMSE	
PR	glm	0.300	0.731	0.789	0.110	1.332	1.335	0.235	0.346	0.418	
	neural network	0.281	0.304	0.414	0.109	0.515	0.526	0.226	0.163	0.278	
	ensemble learning	0.282	0.323	0.429	0.105	0.559	0.569	0.229	0.166	0.283	
DD	glm	0.504	0.640	0.815	0.522	1.111	1.226	0.244	0.347	0.424	
	neural network	1.492	26.060	26.077	2.700	43.489	43.529	0.120	6.396	6.391	
	ensemble learning	0.205	1.103	1.121	0.302	2.116	2.135	0.052	1.184	1.184	
DDPR	glm	0.317	0.897	0.951	0.138	1.566	1.571	0.248	0.432	0.498	
	neural network	0.119	0.627	0.638	0.088	1.116	1.118	0.074	0.341	0.349	
	ensemble learning	0.130	0.539	0.554	0.064	0.922	0.923	0.098	0.262	0.279	

Note: glm denotes generalized linear models. The true values of the main effect β_1 and the interaction effect β_2 are 2, and the population ATE, τ is 3.

D.3 Design 4: Cross-classified Data, No Cross-level Interaction, and Large Cluster Size

Table 12: Performance of Proposed Methods Under Different Estimators of the Propensity Score Model and the Outcome Model in Design 4

	Estimator	Bias	β_1	RMSE	Bias	β_2	RMSE	Bias	τ	RMSE
			SD			SD			SD	
PR	glm	0.486	0.154	0.510	0.050	0.266	0.271	0.460	0.081	0.467
	neural network	0.083	0.113	0.141	0.047	0.160	0.166	0.059	0.080	0.099
	ensemble learning	0.085	0.112	0.141	0.048	0.159	0.166	0.061	0.079	0.100
DD	glm	0.455	0.116	0.469	0.039	0.195	0.199	0.435	0.058	0.439
	neural network	0.012	0.109	0.110	0.008	0.186	0.186	0.007	0.051	0.051
	ensemble learning	0.012	0.108	0.108	0.008	0.183	0.183	0.008	0.050	0.051
DDPR	glm	0.447	0.115	0.462	0.032	0.193	0.195	0.431	0.058	0.434
	neural network	0.019	0.104	0.106	0.017	0.182	0.182	0.010	0.051	0.052
	ensemble learning	0.019	0.104	0.106	0.017	0.182	0.182	0.010	0.051	0.052

Note: glm denotes generalized linear models. The true values of the main effect β_1 and the interaction effect β_2 are 2, and the population ATE, τ is 3.

Table 13: Performance of Estimators Under Different Specifications of the Propensity Score Model and the Outcome Model in Design 4

	Estimator	Bias	β_1	RMSE	Bias	β_2	RMSE	Bias	τ	RMSE
			SD			SD			SD	
Both models correct	PR	0.001	0.065	0.065	0.002	0.115	0.115	0.001	0.033	0.033
	DD	0.001	0.057	0.057	0.001	0.098	0.098	0.000	0.030	0.030
	DDPR	0.001	0.057	0.057	0.001	0.098	0.097	0.000	0.030	0.030
	DRCGEE	0.001	0.057	0.057	0.001	0.098	0.098	0.000	0.030	0.030
Outcome model incorrect	DRPR	0.001	0.150	0.150	0.004	0.298	0.298	0.002	0.065	0.064
	DD	0.003	0.098	0.098	0.008	0.184	0.184	0.000	0.030	0.030
	DDPR	0.004	0.096	0.096	0.010	0.180	0.180	0.000	0.030	0.030
	DRCGEE	0.003	0.098	0.098	0.008	0.184	0.184	0.000	0.030	0.030
Treatment model incorrect	DRPR	0.000	0.066	0.066	0.000	0.114	0.114	0.000	0.033	0.033
	DD	0.001	0.057	0.057	0.001	0.097	0.097	0.000	0.030	0.030
	DDPR	0.001	0.057	0.057	0.001	0.097	0.097	0.000	0.030	0.030
	DRCGEE	0.001	0.057	0.057	0.001	0.097	0.097	0.000	0.030	0.030
Both models incorrect	DRPR	0.552	0.170	0.577	0.070	0.314	0.321	0.516	0.091	0.524
	DD	0.523	0.117	0.536	0.106	0.196	0.223	0.470	0.058	0.473
	DDPR	0.523	0.117	0.536	0.107	0.196	0.223	0.469	0.058	0.472
	DRCGEE	0.523	0.117	0.536	0.106	0.196	0.223	0.470	0.058	0.473

Note: The true values of the main effect β_1 and the interaction effect β_2 are 2, and the population ATE, τ is 3.

E Covariance Balance

We checked covariate balance with respect to the absolute standardized mean differences and variance ratios between the treatment group and control group. As a rule of thumb, if the mean difference of each covariate is less than 0.1 standard deviation and the variance ratio lies between $4/5$ and $5/4$, we can achieve good balance of the covariates (Rubin, 2001; Shadish, Clark, & Steiner, 2008; Steiner, Cook, Shadish, & Clark, 2010). We observed that there was initial imbalance in covariates. After using the marginal mean weighting through stratification (MMW-S) (Hong & Hong, 2009) with the propensity scores obtained from PR estimator, we achieved acceptable covariate balance between the treated and untreated groups.

Figure 9: Covariate Balance Before and After Propensity Score Adjustment
(Standardized mean differences on the x-axis and variance ratios on the y-axis)